

Yellowbrick Data Warehouse: Data Warehousing Without Limits

While legacy platforms like IBM Netezza struggle to stay relevant, and cloud-only options like Snowflake suffer from a reliance on virtualized commodity hardware, Yellowbrick shatters ceilings on price/performance, deployment flexibility across on-premises and major clouds, and low TCO.



Get 100x Performance

Let thousands of users run interactive, ad hoc queries 10x-100x faster than with any legacy or cloud-only data warehouse, at PB scale.



Deploy Anywhere

Consume analytics from a private cloud/data center, major public cloud(s), or both, with users everywhere querying the same data simultaneously.



Save Millions

Pay a fraction of what other options cost via fixed-cost subscriptions for always-on, single-tenant instances—the more queries you run, the lower the cost per query.

Inside Yellowbrick Data Warehouse

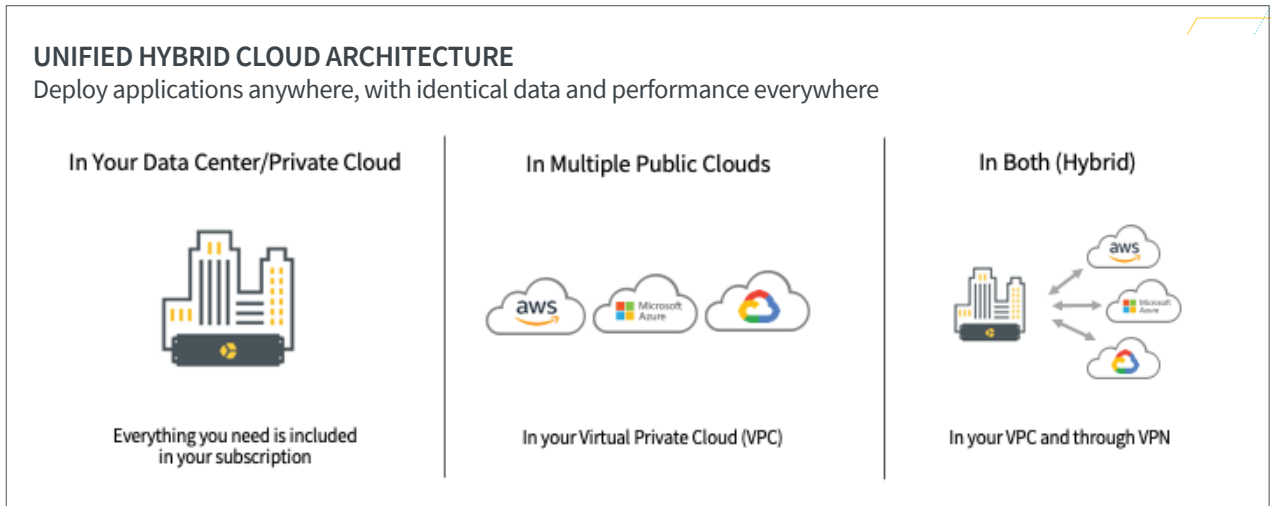
Traditional data warehouses are built on stale, aging architectures that can't keep up with the times. SQL-on-Hadoop engines never met the promise of enabling ad hoc analytics at scale on top of data lakes. And cloud-only options are either built on VMs that can never reliably deliver the performance at scale needed by business-critical applications, or lock you in with their own customized hardware.

Based on a redesign/modernization of MPP analytic databases from the ground up (including storage, networking, CPU, memory, OS, and database), Yellowbrick Data Warehouse offers price/performance for ad hoc analytics far beyond what those options can deliver, along with effortless manageability, enterprise-class security and compliance (HIPAA, SOC, PCI-DSS, FedRamp), and 24x7x365 availability for SaaS applications. And thanks to a PostgreSQL front end, the user and developer experience feels just like you're building for, and working with, the most advanced open source database in the world.

“With the power of Yellowbrick, we can produce insights that were **impossible** before.”

ATUL JAIN
Chairman and CEO
TEOCO

- > Ingest data at line rate speed (10TB/hr) and query it immediately, with no manual data prep needed
- > ANSI SQL compliance and a PostgreSQL front end make management and migrations straightforward
- > Works with common BI, data science, and data motion tools (see yellowbrick.com/partners)



Pick a Cloud, Any Cloud

Unlike legacy or cloud-only platforms that limit your deployment options, Yellowbrick can be consumed as reserved, single-tenant, always-on instances in your data center/private cloud, in a major public cloud, or both—unlocking hybrid and multi-cloud strategies. That approach streamlines and de-risks cloud journeys by making migrations as simple as activating a cloud instance in one of our service locations; on-premises data will be replicated there asynchronously in real-time. Then, connect to your cloud instance through an AWS, Azure, or Google Cloud VPC or directly via VPN, depending on your needs.

For disaster recovery, our low-cost Cloud DR service depends on asynchronous real-time replication to the cloud to deliver multi-region fail-over and fail-back for on-premises Yellowbrick Data Warehouses. When disaster strikes, your cloud instance becomes primary; when normalcy returns, it returns to secondary status.

Yellowbrick: The Only Modern Data Warehouse for Hybrid Cloud

Yellowbrick Data provides the world's only modern data warehouse for hybrid and multi-cloud environments. Enterprises rely on Yellowbrick Data Warehouse to power critical business outcomes and 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 fixed-cost annual subscription.

Accelerate Analytics in AWS, Azure, GCP

For pure cloud deployments, our service locations have private, sub-10ms latency connections to AWS, Azure, or GCP. Yellowbrick is simply another service in your VPC for always-on analytics.

Build Your Private Cloud

Have or want your own private cloud? We'll provide an amazingly compact hardware instance (replacing racks of hardware) that you can have up and running in hours to ingest billions of rows and query them immediately.

Avoid Cloud Lock-in

Never fear cloud lock-in again. Keep your data available and updated in real-time across instances in multiple public clouds as well as on premises, with users everywhere querying the same data simultaneously.

KEY FEATURES

Category	Feature
Manageability	Web-based GUI for admins; remote CLI
Encryption	FIPS 140-2-compliant AES-256 encryption for data at rest, with encryption key rotation
Compliance	HIPAA and SOC 2 compliance; PCI-DSS and FedRamp by request
Resilience and HA	Automatic failover supported for manager node, manager node storage, switches, worker nodes, and worker storage. Yellowbrick Assisted Support provides 24x7 predictive monitoring, and we offer a 99.9 percent service uptime guarantee (under 9 total hours of down time per year).
Backup/Restore	Point-in-time snapshots to external storage
Supported Hadoop File Formats	Apache ORC, Parquet and Avro; CSV; and JSON
Supported Client Platforms	<ul style="list-style-type: none"> - Ubuntu 14.04 LTS, 16.04 LTS, and 18.04 LTS - Red Hat Enterprise Linux and CentOS 6 and 7 - IBM AIX 7.1 - MacOS 10.10 and later - Microsoft Windows Server 2012 and later (with the exception of the backup and restore tools, which are not supported on Windows)
SQL Support	SQL-2011 compliant, with most major features supported
Supported Interfaces	ODBC 3.8, JDBC 4.2, ADO.NET
Import/Export	High-speed (10TB/hr) parallel import and export with Seagate NAS and Amazon S3; Apache Spark and Kafka connectors
Database Portability	Transfer data from Teradata, Vertica, Netezza, Amazon Redshift, Microsoft SQL Server, Oracle, and others
BI Tools Supported	Tableau 10.x, MicroStrategy 10.4, SAS, Information Builders WebFOCUS Enterprise Edition 8, SAP BusinessObjects
ETL/ELT Tools Supported	Informatica 10.2HF2, Syncsort DMEExpress 9.4, IBM Infosphere DataStage