

7 key reasons to replace SQL Server with Yellowbrick Data Warehouse

	Yellowbrick	SQL Server
1 150X performance with no licensing headaches	Yellowbrick customers report 150X (and beyond) performance improvement on complex queries versus SQL Server, at a fraction of the cost and without the management headaches.	SQL Server performance pales in comparison. (For example, scan speed rarely exceeds 10GB/sec, 10X slower than Yellowbrick.) Licensing is also expensive and complicated.
2 Superior scalability with a single instance	Yellowbrick Data Warehouse can scale to multiple PBs, offering a way to combine multiple SQL Server databases into a single, easy-to-manage instance.	SQL Server can't scale for bigger workloads without expensive hardware or dividing applications into multiple data marts, making licensing and management a nightmare.
3 Peak performance out of the box	Yellowbrick provides peak performance out of the box – no manually-created indexes, partitions, or query tuning required; just load and go.	SQL Server requires constant performance tuning and other optimizations, with operations like index, cube, and aggregation maintenance often requiring downtime.
4 Granular control for bigger workloads	Yellowbrick's Advanced Workload Management gives you granular control of workloads to prioritize concurrent queries across massive amounts of data.	With SQL Server, managing workloads for datasets beyond 10TB is complex and difficult, requiring specialized advanced DBA skills that can be hard to find.
5 Expand capacity easily	With Yellowbrick, capacity can be expanded easily without downtime, eliminating any impact on operations.	In comparison, expanding capacity is painful and expensive in SQL Server and requires downtime, affecting operations.
6 Faster data ingest, whether bulk or streaming	Yellowbrick supports bulk data load at near line speed (10GB/sec) and can write millions of rows/second for streaming updates. All data is immediately query-able.	SQL Server can't load data faster than 3GB/sec in bulk and writes ~1m rows/second for streaming updates. Data is not immediately query-able.
7 Backups designed for big data	Yellowbrick natively offers a full spectrum of business continuity tools, including rich backup-and-restore functionality designed for massive data sets.	Managing backup operations across multiple SQL Server instances is complicated and expensive.