Proven Performance. Validated by Evidence.

An Authoritative Technical Summary of Virtuoso's Multi-Model Benchmark Results

Virtuoso by OpenLink Software

A claim of high performance requires extraordinary proof.

Virtuoso is a high-performance, multi-model database supporting Relational, Graph, and RDF data models. Instead of making marketing claims, we present **objective evidence** from the industry's most rigorous, independent arbiters of database performance.

This dossier examines Virtuoso's capabilities through the lens of four key industry-standard benchmarks, demonstrating proven excellence across a wide range of demanding, real-world workloads.

TPC-H: The Gold Standard for Decision Support

LDBC: The Authority on Graph & RDF Systems

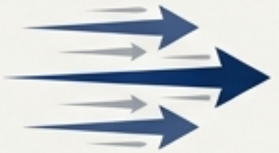

BSBM: The Measure of SPARQL Engine Throughput

DBpedia: The Test of Real-World Query Performance

Evidence Dossier 01: TPC-H

The Standard

The Gold Standard for Complex Decision Support

Maintained by the Transaction Processing Performance Council (TPC), TPC-H is the industry-standard benchmark for data warehousing. It simulates complex analytical workloads with 22 ad-hoc queries over 8 tables, making it the definitive test for evaluating performance on large-scale business intelligence and analytics.

The Findings

Demonstrated Scalability for Enterprise Workloads

100 TB+

Official and community results demonstrate Virtuoso's ability to scale on clusters and cloud instances, handling TPC-H datasets exceeding 100 terabytes.

Virtuoso uniquely treats TPC-H data as Linked Data, enabling the same dataset to be queried via both SQL and SPARQL.

Proven for enterprise-scale data warehousing and ad-hoc analytics.

Evidence Dossier 02: LDBC

The Standard

The Vendor-Neutral Authority for Graph & RDF Systems

The Linked Data Benchmark Council (LDBC) develops industry-driven benchmarks for graph and RDF data management systems. Its key benchmarks, like the Social Network Benchmark (SNB) and Graphalytics, simulate realistic interactive, business intelligence, and core graph algorithm workloads.

The Findings

Superior Performance in Audited Workloads

Merriweather Regular

In the **LDBC Social Network Benchmark (SNB)** Interactive Workload, Virtuoso demonstrated superior performance in short reads and updates on the SF300 dataset.

Virtuoso has shown superiority in formal **LDBC audits**, particularly in tests of core graph algorithms at massive scale using the **Graphalytics** benchmark.

The Verdict

Proven for high-throughput interactive graph applications and large-scale graph analytics.

Evidence Dossier 03: BSBM

The Standard

The Canonical Benchmark for SPARQL Engine Performance

The Berlin SPARQL Benchmark (BSBM) is one of the earliest and most widely used benchmarks for SPARQL query performance. It is built around a realistic e-commerce use case (products, vendors, reviews) and tests explore, business intelligence, intelligence, and concurrent update scenarios.

The Findings

World-Class Throughput at Massive Scale

1,170 QMpH

A Virtuoso 7 cluster achieved up to 1,170 Queries per Minute per Hour (QMpH) on the 1 Billion triples dataset.

Virtuoso has been benchmarked on BSBM datasets scaling from 10 million to **150 Billion triples**, demonstrating a 750x scale increase over prior benchmarks.

The Verdict

Proven for extreme-scale, high-throughput SPARQL workloads for both analytical and transactional use cases.

Evidence Dossier 04: DBpedia

The Standard

The Ultimate Test: Real Queries on Real Data

The DBpedia SPARQL Benchmark (DBPSB) moves beyond synthetic workloads. It mines and replays actual query logs from the public DBpedia endpoint—the structured version of Wikipedia—to test a system's robustness and performance against diverse, real-world query patterns.

The Findings

Robust and Efficient Knowledge Graph Serving

210s → 33s

reduction in cold query times on a 198M triple DBpedia dataset.

Showed superior real-query scalability over Sesame, Jena-TDB, and BigOWLIM in the DBPSB Performance Assessment.

59% faster

than Fuseki on FEASIBLE SPARQL Benchmark query mixes using DBpedia data.

The Verdict

Proven for robust, high-performance deployment against complex and unpredictable real-world knowledge graph queries.

An Unparalleled Breadth of Verified Capabilities

Benchmark	Primary Focus	Data Type	Key Use Cases	Maintainer
TPC-H	Decision support / data warehousing	Relational & RDF	Complex analytics, ad-hoc queries	TPC + OpenLink Software
LDBC	Graph & RDF systems	Property & RDF graphs	Interactive, BI, analytics	Linked Data Benchmark Council
BSBM	SPARQL engine performance	RDF triples	Explore, BI, updates	Freie Universität Berlin / OpenLink
DBpedia	Real-world SPARQL queries	DBpedia knowledge graph	Query-log driven testing	AKSW Group, Leipzig University

The Complete Picture: A Synthesis of Proven Performance

The evidence from these independent, industry-standard benchmarks presents a clear and coherent conclusion. Virtuoso delivers world-class, validated performance not just in one specialized area, but across the critical data models that power modern applications.

Enterprise Analytics (SQL)

Proven for massive data warehousing and complex ad-hoc queries.

(Evidence: TPC-H)

Connected Data (Graph)

Proven for high-throughput interactive applications and large-scale graph analytics.

(Evidence: LDBC)

Knowledge Representation (RDF)

Proven for extreme-scale, high-throughput SPARQL at billion-triple scale and beyond.

(Evidence: BSBM)

Real-World Robustness (KG)

Proven to excel against the diversity and complexity of real-world knowledge graph queries.

(Evidence: DBpedia)

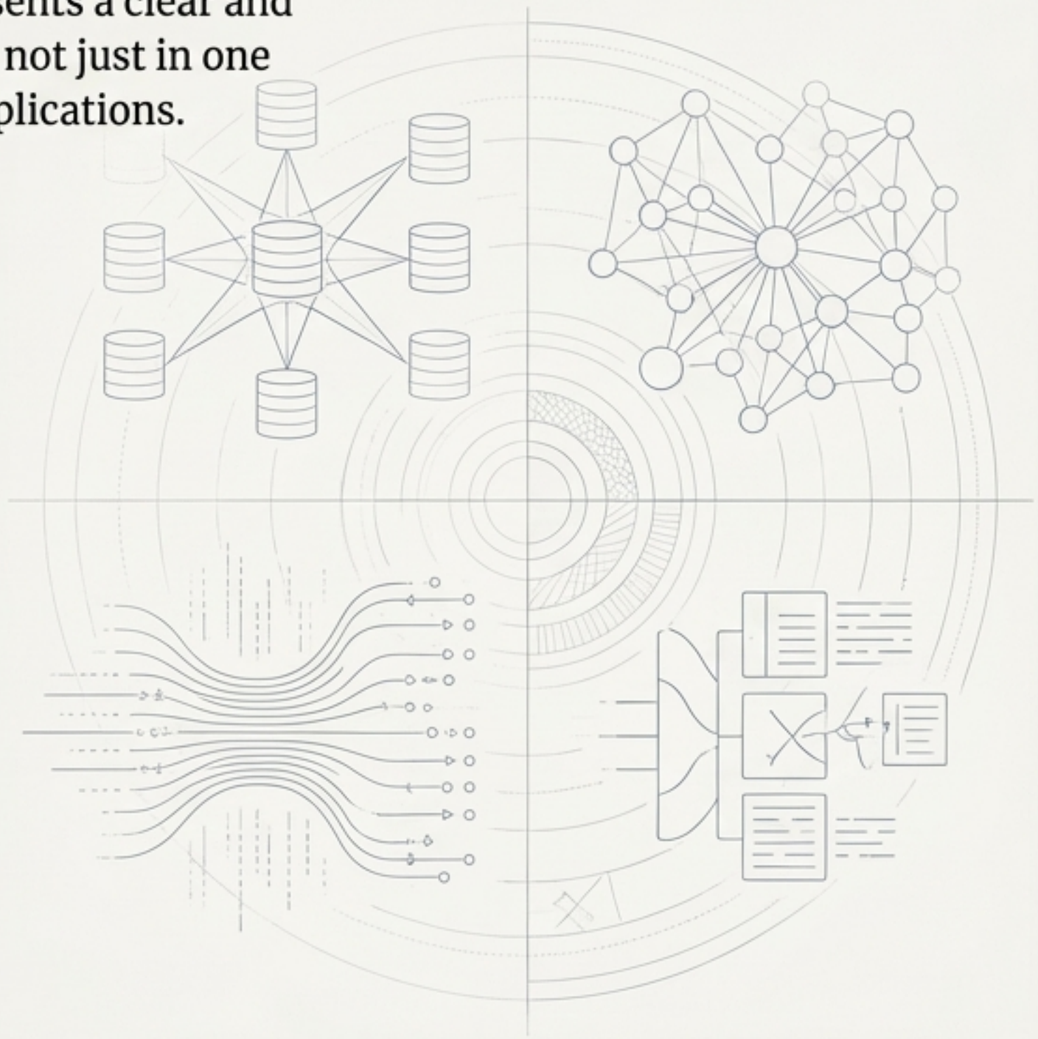

One engine, rigorously tested across the spectrum of data workloads, delivering consistent, world-class performance.