DATA SHEET

RedisGraph

Connecting data in real time to deliver contextual intelligence

One of the most pressing issues with big data is the ability to store, manage and examine hundreds of millions of data points that exist beyond traditional boundaries, determine their relationship with each other and deliver new insights to existing context. The speed with which distributed big data can be analyzed is often critical in making sense of the various datasets.

RedisGraph leverages complex and dynamic relationships in highly connected data to deliver new insights and intelligence across a variety of different use cases, including real-time recommendation engines, personalization, fraud detection, cyber security, master data management, social networking, 360-degree customer view and many more.

Why RedisGraph?

Graph databases leverage complex and dynamic relationships in highly connected data to understand relationships between various datasets, using two common data representation and traversal approaches—adjacency lists and adjacency matrices. Both these approaches have significant drawbacks, with notably slow processing and a large memory and processing overhead.

RedisGraph implements an enhanced matrix traversal methodology representing connected data as sparse adjacency matrices and adopts a standardized engine, GraphBLAS, that uses linear algebra and compressed matrix representation to overcome the performance and scale challenges. RedisGraph simplifies the traversal of highly connected, variable data to answer complex questions and deliver contextual insights.

Benefits of RedisGraph

Rapidly integrate and analyze connected data to drive rich user experiences

Collect, process and analyze complex, connected data and understand data relationships 10 - 600 times faster using the industry-leading solution with record-setting performance for linear computations. Improve customer experience by using native graph support to integrate into or migrate from your current data store

Design for efficiency while delivering lower cost of ownership

Scale easily to meet the needs of your fast-moving organization. Cut historical data storage costs by up to 60% with the optimized graphical representation of large datasets.

Multi-model data platform, in any cloud of your choice

Eliminate operational complexity and improve productivity to unlock the full potential of your connected data. Address multiple use cases and extend to custom data structures using a single multi-model data platform. Proven fault-tolerance and linear scale on commodity hardware or any cloud infrastructure make it the perfect platform for your mission-critical data.

10x-600x

faster than other graph databases

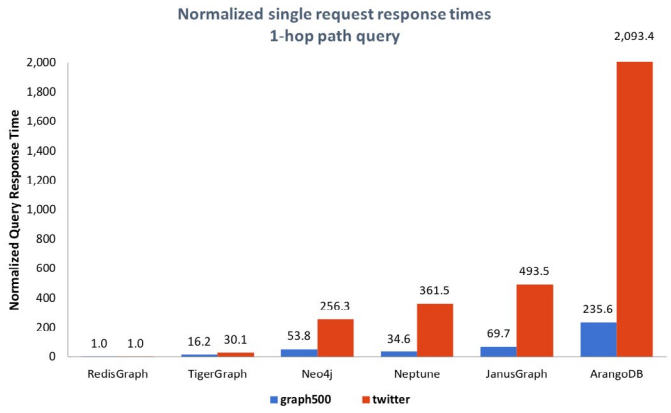

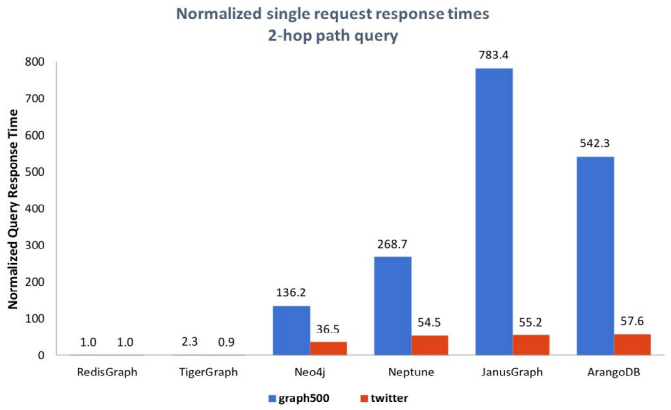

Graph500: 2.4M nodes and 64M edges | Twitter: 41.6M nodes and 1.47B edges

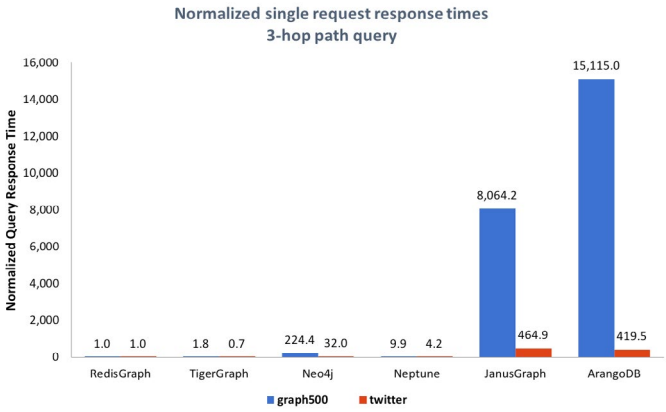

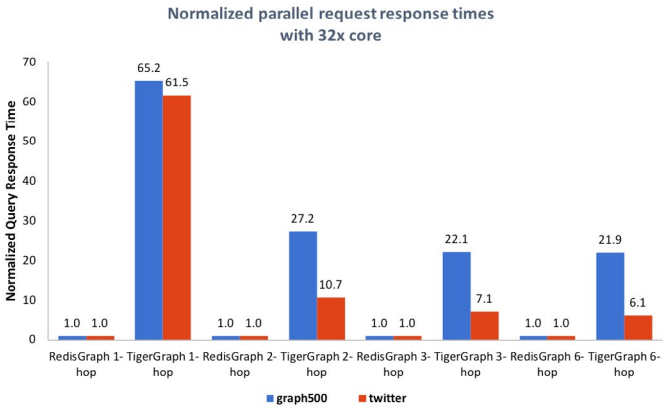

RedisGraph: Under the hood

Memory efficiency:

RedisGraph stores adjacency matrices in Compressed Sparse Column (CSC) format, eliminating the need to track every single entry within the adjacency matrix and instead track only non-zero values. This reduces the amount of memory required to store large graphs, delivering cost savings and turning these matrices into a functional format.

Standardization of query implementation:

RedisGraph adopts the GraphBLAS engine - an open source effort that provides standardized building blocks of graph algorithms based on linear algebra to find the nearest neighbors for data computation.

The combination of sparse matrix representation and linear algebra optimizes and simplifies many different graph queries and algorithms and is computationally simpler.

Improving developer productivity:

RedisGraph implements and supports the declarative and widely adopted Cypher query language for graph databases, and automatically translates them into linear algebraic expressions, enabling developers to adopt open standards to quickly develop applications.

Multi-model data platform for all your data needs

Combine the benefits of multiple data models into one platform and provide a unified interface to search across all your connected data easily. Eliminate your data silos, enable your developers to build applications faster, reduce your data storage costs and ensure faster and improved data governance. Future-proof your investment with continuous innovation from the vibrant open source community and back your deployment with enterprise-grade functionality and 24/7 support.

- Amplify performance with the distinctive shared-nothing architecture, linear scale and the most efficient use of your hardware resource.
- Eliminate downtime with seamless failure detection, failsafe high availability, instant failover and zero data loss.
- Simplify development of globally distributed applications with versatile data structures and support for Active-Active and consensus-free conflict resolution*.
- Reduce operational overhead by fully automating configuration, provisioning, deployment, and continuous monitoring.
- Secure your business with advanced security and certifications as well as multi-level authentication.
- Deliver more savings with tiered memory storage, using Flash memory as a RAM extender and built-in multi-tenancy*.
- Gain portability across all major public clouds and platforms and on-premises for hybrid/multi-cloud deployment.

* Coming soon

"We've been using Redis for over a decade and it has been great to see the success our teams are achieving with it. With the simplicity and scale that RedisGraph is able to deliver with its multi-tenant architecture and the compactness of the data storage, we've been able to reduce our infrastructure footprint quite considerably. With support for the Cypher Query Language, the switch from our existing graph implementation in favor of RedisGraph was straightforward. Now, with the expected increase of the speed of our queries, we can get the insights we rely on faster than ever."

Emre Sokullu
CEO and Founder

GROU.PS

Get Started with RedisGraph Today!

Visit https://redis.com/redisgraph/.
Talk to a Redis Enterprise expert today. Contact expert@redis.com.

redis

700 E El Camino Real Suite 250, Mountain View, CA 94040 (415) 930 9666 redis.com