



## TIGERGRAPH OUTPERFORMS NEO4J BY ORDERS OF MAGNITUDE

A comprehensive graph data management study by researchers at the University of California Merced found that TigerGraph consistently outperformed Neo4j, more than 100 times faster in some cases, with that gap increasing with the size of data. The study is the [Social Network Benchmark \(SNB\) of the Linked Data Benchmark Council \(LDBC\)](#) and considered to be the reference standard for evaluating graph databases.

*Disclaimer: These benchmark results are not official LDBC benchmark results, as they have not been audited, and the experiments do not follow all rules necessary for such an audit, as detailed in <https://ldbouncil.org/benchmarks/snb/retrospective-report-tigergraph.pdf>.*

Some of the key findings of the study were:

- **TigerGraph is the first scalable, distributed graph database to demonstrate both analytics and transactional processing capability.** In the graph database market, vendors can handle either OLTP or OLAP, but not both – until now. As the first evaluation of performance with both interactive and sophisticated query workloads, the study demonstrated TigerGraph’s unprecedented ability to scale to ever-larger data sets when performing complex queries, even with graph pattern matching queries, which are considered more challenging than relational joins.
- **Only TigerGraph was able to scale to 1TB.** TigerGraph was shown to **consistently outperform Neo4j on the majority of queries by two or more orders of magnitude (100x factor)** on certain interactive complex and business intelligence queries, and up to four orders of magnitude on others. On the larger datasets, Neo4j often timed out, unable to complete the query within five hours.
- **TigerGraph shows significantly compressed storage on typed property graph data as compared to Neo4j.** While bulk loading time is comparable for these two databases, the storage sizes are significantly different, with Neo4j’s being approximately four times that of TigerGraph in all settings.

The study is the first complete test of graph database vendors’ performance with intensive analytical and transactional workloads. In addition to thoroughly sizing up the performance of the 46 queries on four data input scale factors, from 1GB to 1TB, the study also measured bulk loading time and storage size. As such, the UC Merced study is a unique assessment of graph analytics platforms’ ability to handle real-world data challenges in real time, regardless of how large or complex the data set is.

The power to execute increasingly arduous computations in real time is crucial for many of today’s most important applications, such as fraud and money laundering detection, customer 360, security analytics, hyper-personalized recommendation engines, artificial intelligence and machine learning.



The study compared graph database systems from TigerGraph and Neo4j across a range of short-running (OLTP) and long-running (OLAP) inquiries:

**1. Interactive Short (IS) Workload** - this type of query begins at a single source and requires graph traversals of two or fewer hops. Sample results are shown in figure 1.




Figure 1: TigerGraph completed IS queries much faster than Neo4j in many instances

TigerGraph was faster than Neo4j in multiple instances in this section of the study although, because only two or fewer graph traversals were required, the difference between the two systems was small.

**2. Interactive Complex (IC) Workload** - this type of query begins at a single source and requires graph traversals which exceed two hops. Sample results are shown in figure 2.




Figure 2: TigerGraph completed IC queries significantly faster than Neo4j - and Neo4j's timed out in some instances



TigerGraph was considerably faster than Neo4j in the vast majority of tests in this section of the study—**sometimes by as much as four orders of magnitude**. While TigerGraph finished the queries in tens of milliseconds, Neo4j required more than a thousand seconds.

**3. Business Intelligence (BI) Workload** - this type of query begins from multiple points in the graph and requires extensive optimizations across all the layers of the system—not only graph traversal. Sample results are shown in figure 3.




Figure 3: TigerGraph completed BI queries faster than Neo4j

TigerGraph clearly outperforms Neo4j across almost all queries in this section of the study and, on average, is close to **10X faster than Neo4j** across many queries. Moreover, as the size of the data set increases to terabytes, **fewer and fewer queries were performed in the allocated time by Neo4j**.

## ABOUT THE STUDY

The Linked Data Benchmark Council is an effort to establish benchmarking practices for evaluating graph data management systems. The main objectives of LDBC are to design benchmark specifications and procedures, and publish benchmarking results. The LDBC Social Network Benchmark (SNB) is the first result of this effort and was applied by researchers at the University of California Merced to evaluate Neo4j and TigerGraph in an independent study. You can read the full report here: <https://arxiv.org/pdf/1907.07405.pdf>



**TigerGraph**

## **The Social Network Benchmark of the Linked Data Benchmark Council**

### **About TigerGraph**

Get started today with TigerGraph Cloud ([www.tigergraph.com/cloud](http://www.tigergraph.com/cloud)) which contains multiple starter kits demonstrating the seven key capabilities delivered by a native parallel graph. It's FREE to sign up and takes only a few minutes.

TigerGraph is the only scalable graph database for the enterprise. Based on the industry's first native parallel graph technology, TigerGraph unleashes the power of interconnected data, offering organizations deeper insights and better outcomes. TigerGraph fulfills the true promise and benefits of the graph platform by tackling the toughest data challenges in real time, no matter how large or complex the dataset. TigerGraph's proven technology supports applications such as fraud detection, customer 360, MDM, IoT, AI and machine learning to make sense of ever-changing big data, and is used by customers including Amgen, China Mobile, Intuit, Wish and Zillow, along with some of the world's largest healthcare, entertainment and financial institutions. The company is headquartered in Redwood City, CA. Follow TigerGraph on Twitter at @TigerGraphDB or visit [www.tigergraph.com](http://www.tigergraph.com).