TigerGraph vs. Neo4j Scalability Overview

Unified Enterprise Schema

TigerGraph
One schema.

Neo4j 4.0 Fabric
DB administrators must manually shard the schema into different sub-schemas for each machine/database.

SOURCE

Automatic Data Partitioning

TigerGraph
Load data with a single loading job and partition across a cluster of machines automatically.

Neo4j 4.0 Fabric
Users must manually partition data and load each partition or shard separately to each machine/database.

Distributed Querying

TigerGraph
Write the query once and use it for 1, 10, 100, or 1,000 machines.

Neo4j 4.0 Fabric
Users must design multi-stage queries to manually query each machine/database and then stitch results together.

SOURCE

ACID Transactions Across the Cluster

TigerGraph
Full ACID transactions across the cluster with distributed graph.

Neo4j 4.0 Fabric
ACID transactions are not valid in a fabric. Transactions are only within a *single graph*.

SOURCE

Graph Algorithm Execution Across the Cluster

TigerGraph
Seamless execution of graph algorithms such as community detection and PageRank across the cluster.

Neo4j 4.0 Fabric
Can't run graph algorithms across the shards.

SOURCE

More Details in Our Blog Post. Read HERE.