



vs.



What you can do with ArangoDB that you can not do with Neo4j

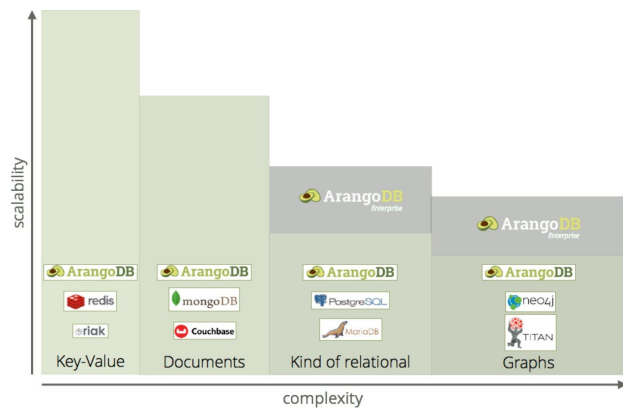
Graph technology is a powerful tool in many use cases but not everywhere the best access pattern to your data. While providing the same graph functionalities as Neo4j the superior flexibility in data modelling gives ArangoDB an edge – thanks to being C++ based and smart index usage also performance-wise.

- ▶ Multi-model: Create arbitrary ad-hoc queries on data that is stored in different data-models
E.g. filtering all parts from an aircraft that are connected or developing a web shop with recommendation engine harnessing the advantages of polyglot persistence without the cost
- ▶ Scalability: ArangoDB scales with all supported data models including graphs. Achieve single instance performance in a cluster with ArangoDB SmartGraphs.
- ▶ Extensibility: Use existing or run your own data-centric microservices in a dedicated JavaScript framework within ArangoDB
With ArangoDB's Foxx framework users can build e.g. a production ready session service within 10 minutes
- ▶ Performance: ArangoDB is pretty fast (see Benchmark results below)
- ▶ Operational costs: ArangoDB can be used for a broad range of use cases and reduces the number of storage products in your technology stack.

Scalability and ArangoDB

ArangoDB is cluster ready for document, key/value and even for graph-models. Leverage Satellite Collections and SmartGraphs for high end performance at scale.

ArangoDB is suitable for e.g. recommendation engines, personalization or other graph-related use cases. Tests showed higher performance than Neo4j.



Founders and Core Team

Both founders combined have more than 30 years of experience building high performance databases specifically designed for challenging use cases.

In projects for e.g. German Postal Service, DHL, Commerzbank and the New York Stock Exchange Frank and Claudius connected the dots to build a native multi-model database which can compete with pure document, key/value and graph database.

“ArangoDB with RocksDB support and granular access control, it blows away other proposed DBs in my environment.”

(Tanvir Mansuri, Lead Developer Thomson Reuters)

ArangoDB Inc.

548 Market St #61436,
San Francisco, CA, 94104-5401
United States

Contact: Jan Stücker

Email: jan.stuecke@arangodb.com
Mobile: +49 (0) 221-2722999-60



vs.



High-Level Overview: Neo4j vs. ArangoDB

Feature	Neo4j	ArangoDB
Initial Release	2007	2012
Data-Model	graph	multi-model (documents, graphs, key-value)
Data format	JSON	JSON/Velocypack
Written in	C++	C++
License	AGPLV3 / Commercial	Apache 2 / Commercial
Data Storage	Neo4j graph storage	MMFiles / RocksDB
Schema free	Yes schema validation	Yes schema validation via Foxx
Replication	Master/Slave	Master/Slave
Sharding	No	Auto
Relations	Edges	Edges/Joins
Declarative Query Language	Neo4j Cypher, Gremlin	AQL One sql-like query language for all data-models
Joins	No	Yes
Transactions	ACID	ACID ⁽¹⁾
Complex Domains	Labels (multiple)	Edges are documents without any restriction in complexity
Pattern Matching	Yes	Yes
Extensibility	Server-side Java plugins	Microservices framework Foxx based on Google V8 ⁽²⁾
Relational Joins	No	Yes
Cluster ready	No	Yes ⁽³⁾
Encryption at rest	No	Yes
Authentication	Yes	Yes
Role-based access control	No	Yes Attribute level via Foxx Framework ⁽⁴⁾
Auditing	No	Yes

^{1.} In single server setups, ArangoDB supports full ACID transactions for multi-document & multi-collection transactions. In a cluster setup, ArangoDB only supports multi-document & multi-collection transactions for non-sharded collections. Single document transactions are supported for sharded collections.

^{2.} Easily create a REST API for data centric use cases and add any missing functionality

^{3.} Scalability is possible with ArangoDB and its graph model but whether queries run quickly depends on the individual use case

^{4.} ArangoDB supports all basic security requirements. By using ArangoDB's Foxx microservice framework users can achieve very high security standards fitting individual needs

ArangoDB Inc. 548 Market St #61436, San Francisco, CA, 94104-5401 United States	Contact: Jan Stücker Email: jan.stuecke@arangodb.com Mobile: +49 (0) 221-2722999-60
---	--



vs.



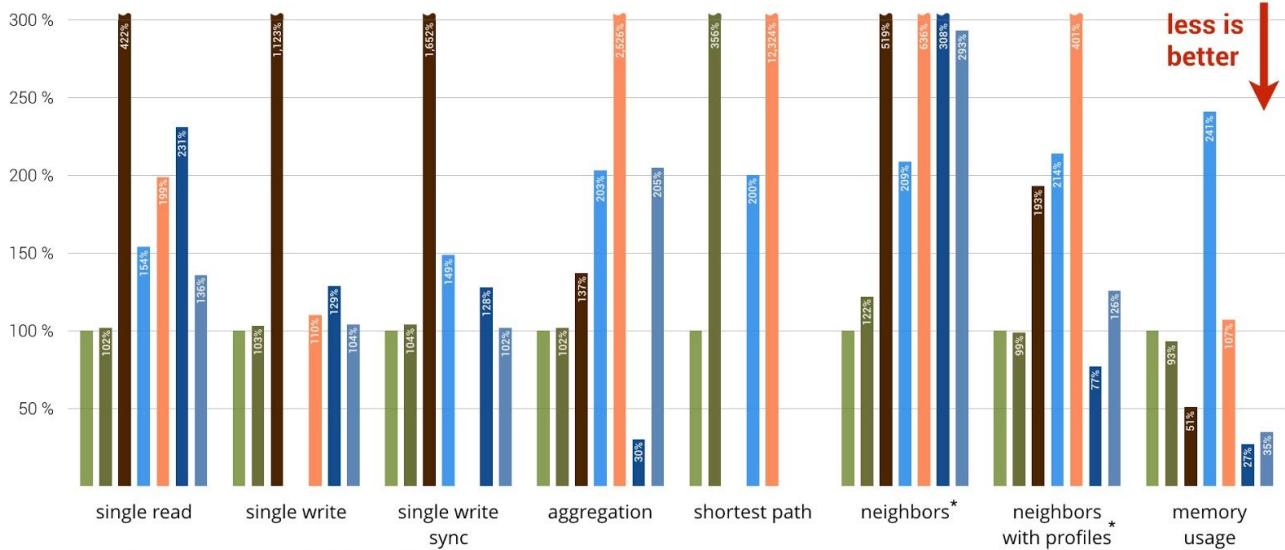
Performance Benchmark of ArangoDB

The latest performance benchmark (2018) shows that ArangoDB’s native multi-model approach can compete with leading single model vendors on their home turf. The benchmark test is completely open-source and investigated the performance of most common database operations. The benchmark test included performing single read/write, single write sync, aggregation, shortest path, neighbors, neighbors with profiles and memory usage.

Overall Results: Performance Benchmark (February 2018)

NoSQL Performance Benchmark 2018
ArangoDB, MongoDB, Neo4j, OrientDB and PostgreSQL

- ArangoDB Rocks
- ArangoDB MMfiles
- MongoDB
- Neo4j
- OrientDB
- Postgres (tab)
- Postgres (jsonb)



*) neighbors and neighbors of neighbors (distinct)

arangodb.com/performance – 2018-02-27

For further information please find the detailed article of the benchmark [here](#):

All interested parties can find the open-source performance test on Claudius Weinberger’s GitHub profile: <https://github.com/weinberger/nosql-tests>

ArangoDB Inc.

548 Market St #61436,
San Francisco, CA, 94104-5401
United States

Contact: Jan Stücker

Email: jan.stuecke@arangodb.com
Mobile: +49 (0) 221-2722999-60