ArangoDB

The Leading
Native Multi-Model Database

Copyright © ArangoDB Inc., 2018 - Confidential

Company at a Glance

$7.2m Funding

48 Employees

7.2m Downloads

95+ External code contributors

Numbers in parentheses are from January 2017

335+ (180) Production installations

Fortune 10 and **500** customers

6,100+ (3,000) Stargazers

Multi-Model Pioneers

Mentor · cisco · BARCLAYS · THOMSON REUTERS · NASA · nurocor · Kabbage

Copyright © ArangoDB Inc., 2016 - Confidential

Performance: Single Instance (execution time, less is better)

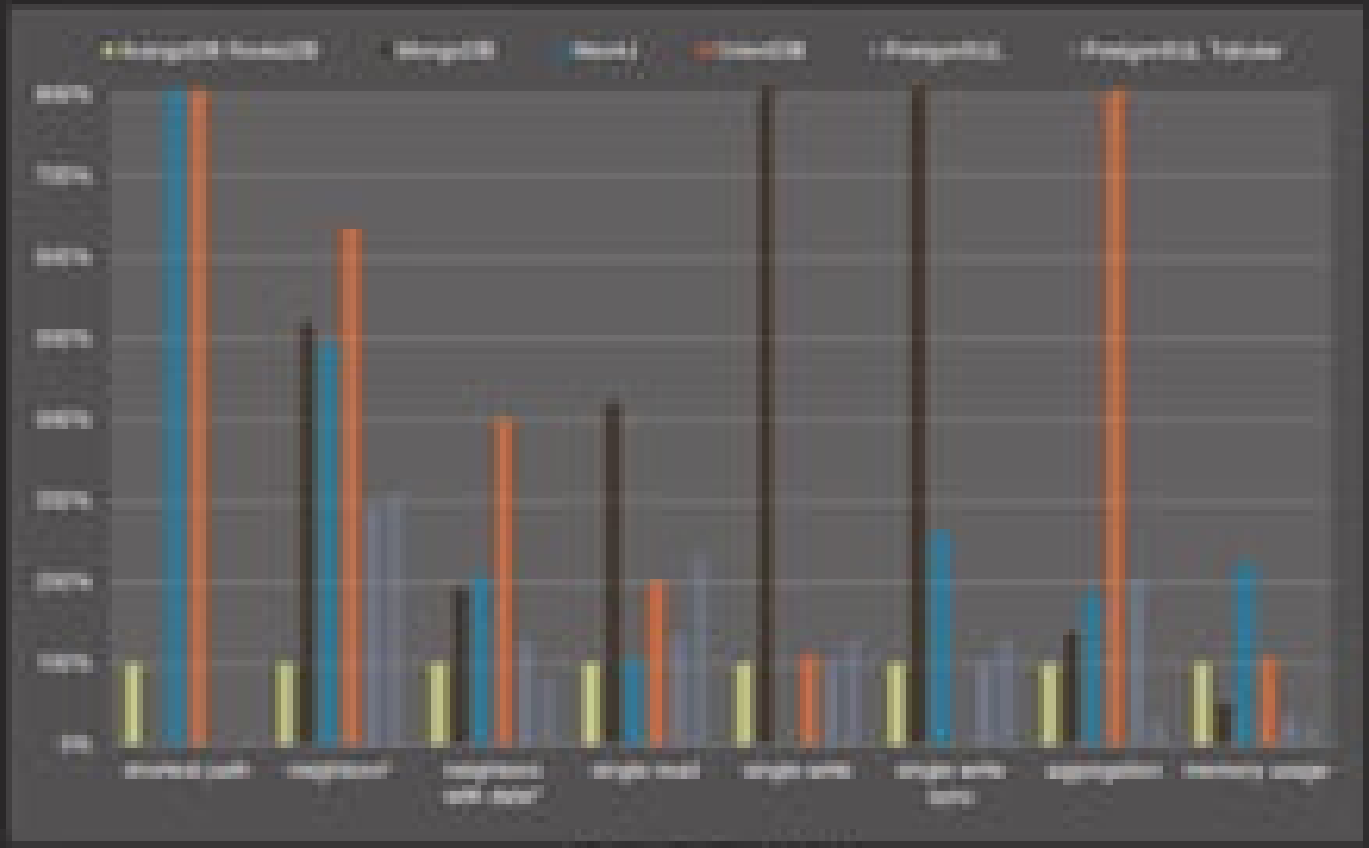

Cluster (more is better)

Database	Document writes per vCPU/second
ArangoDB	1,730
Aerospike	2,300
Cassandra	965
CouchBase	1,375
FoundationDB	758

Scales Linearly

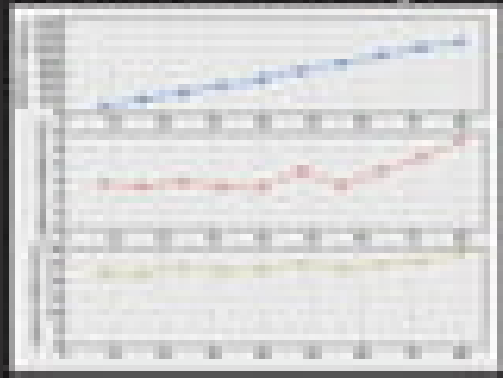

Copyright © ArangoDB Inc., 2016 - Confidential

ArangoDB

The Leading Native Multi-Model Database

Copyright © ArangoDB Inc., 2019 - Confidential

Databases, a Crowded Market

200X

- New challenges (unstructured data, data volume, velocity) led to the development of new technologies.
- The architectural pattern Polyglot Persistence emerged.
  - —> choose the right database / data model for the job.
    recommendations -> graphs
    metadata -> document
    sensor data -> time series
- The overall system complexity rises.

201X

- To keep operational and development efforts (TCO) within limits, companies are trying to narrow their tech stack.

2017

- Multi-model - offering support for different data models - is the key to succeed for DB vendors in the coming consolidation process. ArangoDB develops a **native** multi-model approach that allows developers to select and combine data models as they see fit.

Copyright © ArangoDB Inc., 2019 - Confidential

Databases, a Crowded Market Continued

- In the last few years, the decision-making process was more driven by solving *new* problems with new technologies.
- In this process, an architectural pattern has emerged that will not disappear.
  - Polyglot Persistence: Choose the database for the part of your system that solves the problem best.
    But it isn't about choosing the right database it's more about choosing the right data model. Recommendations -> graphs, Metadata -> document, sensor data -> times series, and so on
- Nowadays, companies moving back more to TCO and flexibility for their decisions because they have very large databases stacks to handle (>20). That's an issue on the operational side but even more on the development side.
- Multi-model is the main driver for that because different data models will also needed in the future. The support of different data models is the key for a modern database.
  - Layered approach
  - Native multi-model
- In addition, there are the "standard" criteria for databases, including:
  - Performance
  - Open Source

Copyright © ArangoDB Inc., 2019 - Confidential

Database Market

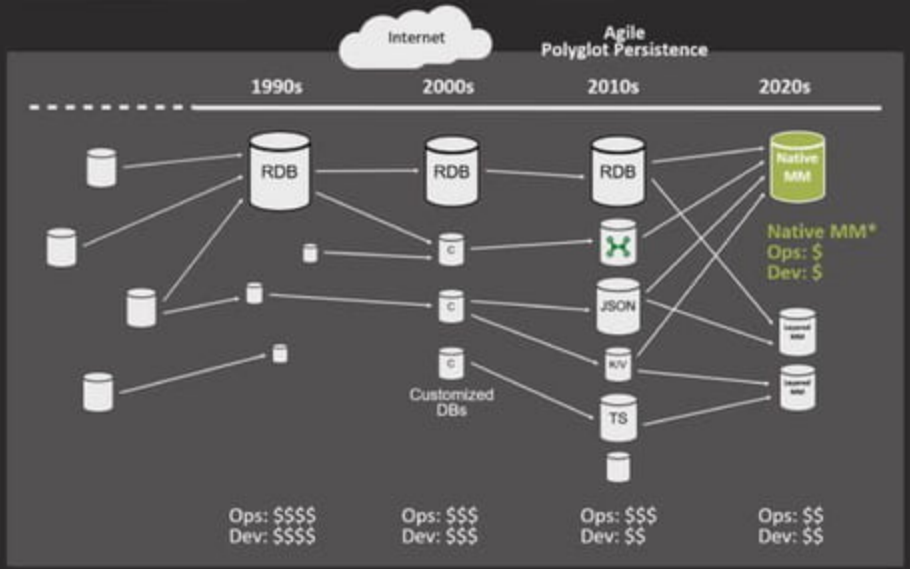

- '90: Consolidation to Relational DBs
- '00: Internet boom, new customized DBs.
- '10: Polyglot Persistence, raise of noSQL & newSQL
- '20: Consolidation to Multi-Model DBs

*) native multi-model: (one core, several data models); layered multi-model (several data models, no common core)

Copyright © ArangoDB Inc., 2019 - Confidential

Native Multi-Model Approach

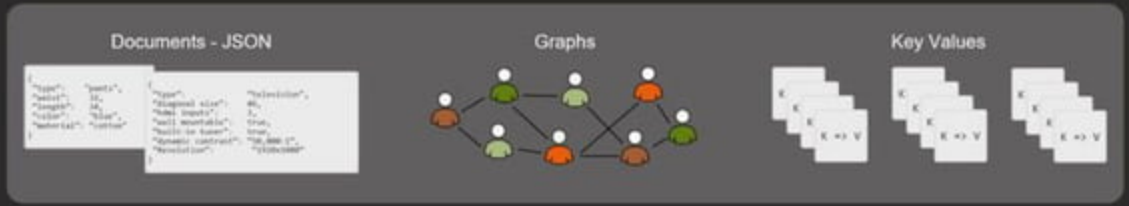

One Engine, One Query Language.
Multiple Data Models.

Copyright © ArangoDB Inc., 2019 - Confidential

Product Positioning

Copyright © ArangoDB Inc., 2019 - Confidential

Tame Complexity (e.g. Modern Ecommerce App)

- Different Query Languages to Learn
- Many Databases to Administer
- Complex Code Base
- Increased Costs
- Hindered Productivity

- One Query Language to Learn
- One Database to Administer
- Streamlined Code Base
- Lower Total Cost of Ownership
- Dramatically Improved Productivity

Copyright © ArangoDB Inc., 2019 - Confidential

Company at a Glance

$7.2m Funding

48 Employees

7.2m Downloads

95+ External code contributors

Numbers in parentheses are from January 2017

335+ (180) Production installations

Fortune 10 and **500** customers

6,100+ (3,000) Stargazers

Multi-Model Pioneers

Mentor
CISCO
BARCLAYS
THOMSON REUTERS
NASA JPL
nurocor
bringing order to clinical data
Kabbage

Copyright © ArangoDB Inc., 2019 - Confidential

GitHub Stars % Growth over the last 2 years

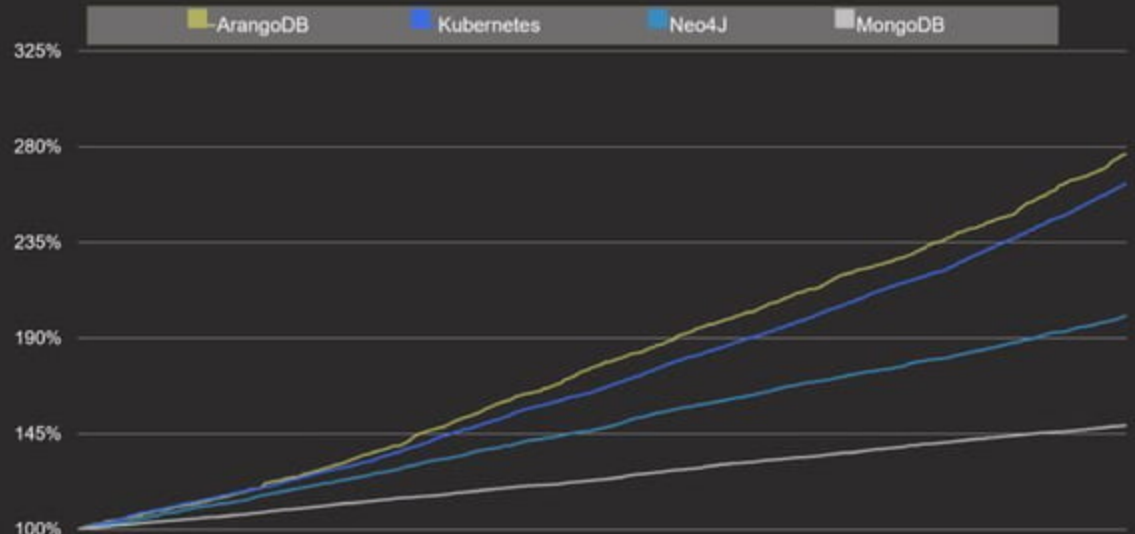

Copyright © ArangoDB Inc., 2019 - Confidential

updated: 2/14/18

The Founders

Claudius Weinberger

CEO & Co-Founder

- Serial entrepreneur: started first software company with 20 & has led teams ever since
- Master degrees in economics and computer science
- Has built and operated multitude of databases since 2000
- Lead the product development department in a financial technology company, with responsibility for 80 individuals. Drove product strategy and execution in the company's expansion from 20 to 250 FTE, including IPO and merger with a peer

Frank Celler (PhD)

CTO & Co-Founder

R&D and customer care

- 30 years of development experience in high performance databases
- PhD in Mathematics from RWTH Aachen
- Developed GAP, a programming language for group theory still in use world wide and awarded with the ACM/SIGSAM Richard Dimick Jenks Memorial Prize
- Development and operation of databases in a listed European financial information company
- Organizer of the biggest NoSQL conferences in Europe with thousands of attendees

Founders working together for 18 years.

In 2004, they started a company building bespoke databases for large enterprises like Deutsche Bank, NYSE Euronext, Commerzbank and the German Postal Service.

Copyright © ArangoDB Inc., 2019 - Confidential

Investors and Industry Advisors

Current Investors

TARGET PARTNERS is one of the leading tech investors in Germany. They invest in young technology companies in Germany, Austria and Switzerland. In 2017 Target Partners helped its portfolio of 26 active portfolio companies to raise more than €110m in 25 separate financing rounds.

Machao Holding AG is an investment vehicle of a Norwegian/Swiss family to invest in different opportunities. Machao Holding AG was founded in 1978.

Industry Advisory Council

Florian Leibert, Advisor
Florian Leibert is the CEO and Co-founder of Mesosphere. Florian helps ArangoDB to further sharpen its focus on the needs of large scale enterprises, their problems and solutions.

Ted Dunning, Advisor
Ted Dunning is Chief Application Architect at MapR and has years of experience with machine learning and other big data solutions across a range of sectors. He was the chief architect behind the MusicMatch - now Yahoo Music - and Veoh recommendation systems.

Copyright © ArangoDB Inc., 2019 - Confidential

AQL - A Query Language That Feels Like Coding

- For a native multi-model database a common query language for all data-models is crucial
- AQL aims to be human-readable
- Same language for all clients, no matter what programming language people use
- Easy to understand for anyone with an SQL background

```
SELECT customers.name
FROM customers
WHERE customers.id = ?
      and
      customers.isActive
```

```
FOR c IN customers
FILTER c.id == @id
       and
       c.isActive
RETURN c.name
```

```
SELECT
  customer.name,
  orders.date,
  orders.amount
FROM customers
JOIN orders
  ON customers.customer_id =
     orders.customer_id
WHERE customer.id = ?
```

```
FOR c IN customers
FILTER c.customer_id == @id
RETURN {
  customer: c.name,
  orders: (
    FOR o IN orders
    FILTER o.customer_id ==
           c.customer_id
    RETURN {
      date: o.date,
      amount: o.amount })
  }
```

Copyright © ArangoDB Inc., 2019 - Confidential

Integrated with leading Orchestration systems

- ArangoDB is the first fully certified database including the persistence primitives for DC/OS
- Orchestration Systems such as Mesosphere DC/OS, and Kubernetes are the future
- ArangoDB's cluster resource management offers an innovative two layer approach
- Later this year we will also support Kubernetes and Docker Swarm based on our two layer approach
- ArangoDB has been a Mesosphere partner since 2014

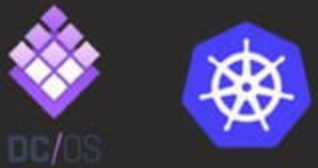

Copyright © ArangoDB Inc., 2019 - Confidential

Performance: Single Instance (execution time, less is better)

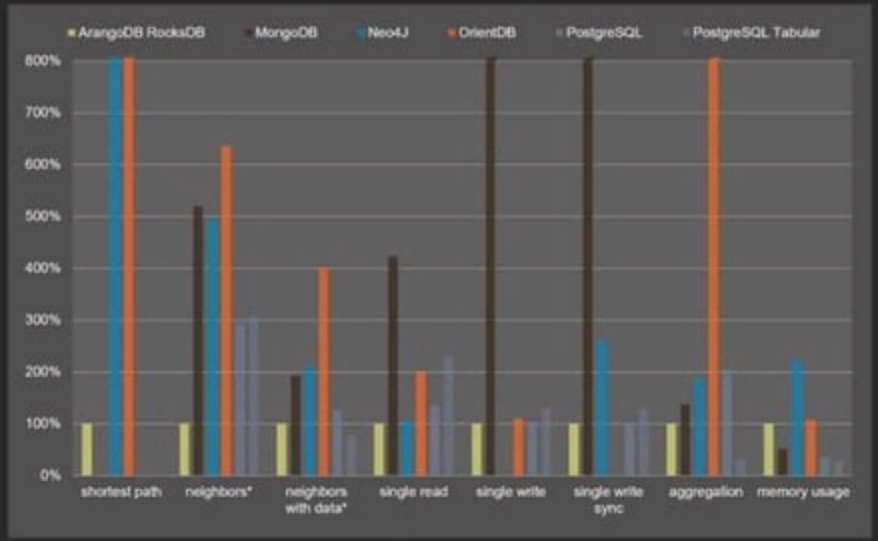

New version - released: 2/14/18

Cluster (more is better)

Database	document writes per vCPU/second
ArangoDB	1,730
Aerospike	2,500
Cassandra	965
Couchbase	1,375
FoundationDB	750

Scales Linearly

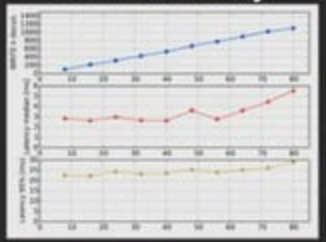

Copyright © ArangoDB Inc., 2019 - Confidential

ArangoDB Enterprise

Enterprise Features like:

- Auditing (3.1)
- Stronger security features (3.1)
- On-Disc encryption (3.2)
- Encrypted backup (3.3)
- LDAP (3.3)
- Multi data center support (3.3)
- Managed Service

In addition, we support unique features for enterprise customers like:

- Smart Graphs (3.1)
- Satellite Collections (3.2)
- OneShard Cluster
- n Dimensional Geo Index

With the Enterprise Edition (GA in November 2016, v 3.1) we started our commercialisation

Copyright © ArangoDB Inc., 2019 - Confidential

SmartGraphs – Fast, distributed Graph Processing

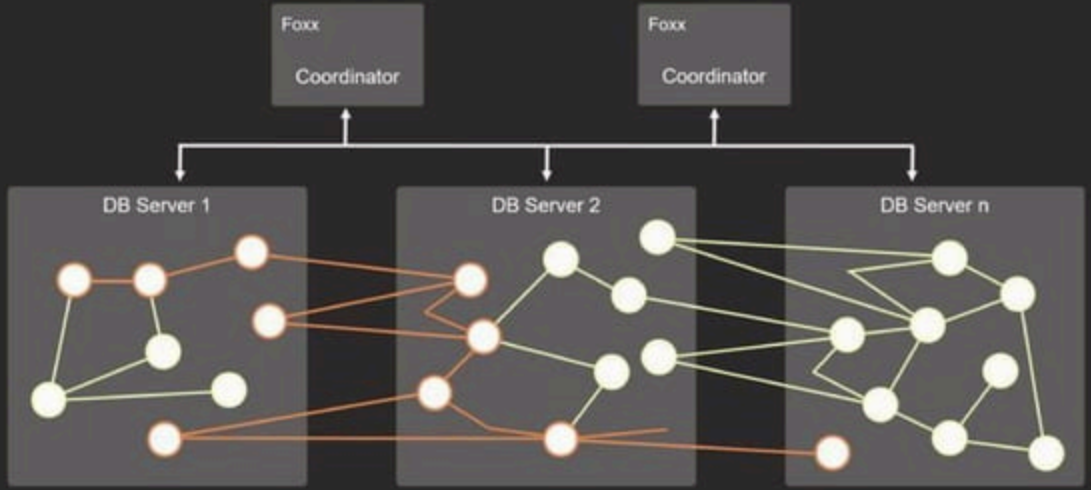

Copyright © ArangoDB Inc., 2019 - Confidential

SatelliteCollections - Fast, distributed JOINs

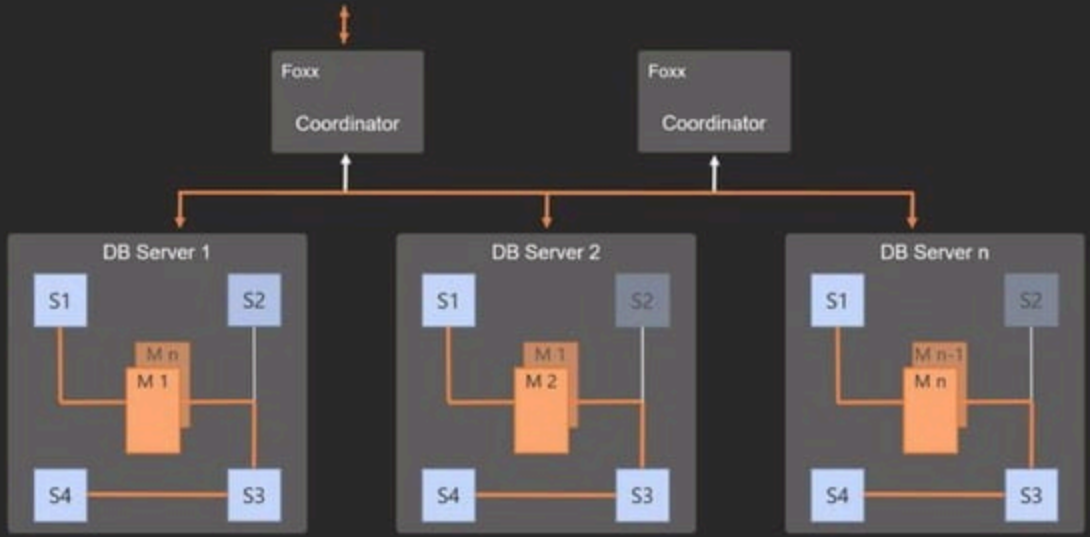

Copyright © ArangoDB Inc., 2019 - Confidential

Roadmap - Top Features for the next 18-24 Months

Community Edition

- **ArangoDB Search**
  Integration of a full text ranking engine similar to Lucene but based on C++
- **Kubernetes Integration**
  Official support for Kubernetes with an operator
- **Improved Geo Index**
  Better geo search functions based on the Google S2 lib
- **Schema Validation**
  Direct support of JSON Schema Validation on C++ level
- **Trigger**
  Cluster-wide triggers to initiate internal and external actions
- **Distributed Transactions**
  Cluster-wide multi document transaction, configurable on collation level
- **Zero Administration**
  Improve the deployment and administration capabilities

Enterprise Edition

- **Managed Service**
  A cloud based managed service on all major platforms starting with AWS
- **OneShard Database**
  Full ACID compliance with strong HA capabilities if the data set fits into one machine (many legacy use cases)
- **Improved Pregel Support**
  Improve Pregel support to cover more Big Data Graph use cases
- **Incremental Backup**
  Incremental consistent snapshot backups

Copyright © ArangoDB Inc., 2019 - Confidential

Prepare for Next Year

Prepare for next year's goals

Lay the foundation for the next stage of our marketing, sales, support and development organization, hiring seasoned executives.

Change sales approach/compensation

Switch from cash-driven TCV sales model to a renewable ACV/ARR sales model/compensation.

Reference customers

Get landmark reference customers on our website and advocating our technology.

Next 18 Months

Substantially scale the company in all areas

Make United States the main territory with the largest market share

Build our partnerships to scale non-linearly, including channels sales and partner marketing

Copyright © ArangoDB Inc., 2019 - Confidential

ORACLE

Thank you!

Copyright © ArangoDB Inc., 2019 - Confidential