aquenos
software & more

Latest Trends in Database Technology

NoSQL and Beyond…

Sebastian Marsching

Why we want more than SQL…

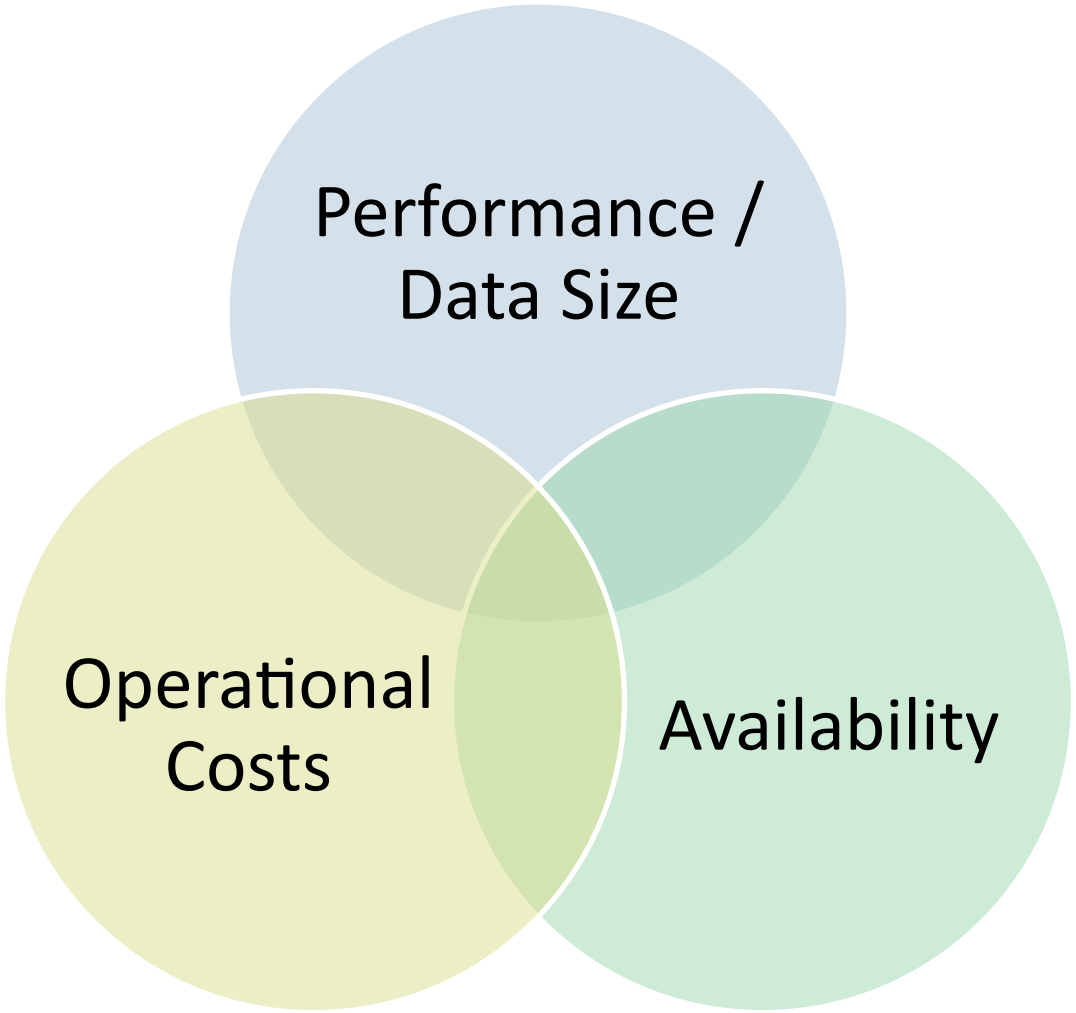

NoSQL

NoSQL ↔ Not Only SQL

NoSQL Concepts

- Document Stores
  - Store (hierarchical) documents (e.g. XML databases, MongoDB, CouchDB).
  - Allow complex queries on these documents.
- Graph Databases
  - Store relations between nodes.

NoSQL Concepts (cont.)

- Key-Value Stores
  - Typically in-memory databases (e.g. memcached, Redis).
  - Low latency.
  - Extremely simple.
- Column-Oriented Databases
  - Simple query model.
  - Typically on-disk storage.
  - Designed for high throughput and scalability.

Common Concepts in NoSQL

- Replication
  - Store same data on multiple nodes to improve availability.
  - Achieve reliability on software instead of hardware level (e.g. no RAID).
- Sharding
  - Distribute data over multiple nodes.
  - Allow for transparent (and possibly linear) scaling by adding more nodes.
  - Use many relatively small systems instead of one big system.

Document Stores in Scientific Applications

- Often used for storing experiment meta-data and measurement data organization.
- Typically not optimized for very large amounts of data.
- MongoDB is very popular.

Key-Value Stores in Scientific Applications

- Typically hybrid solutions (combining a key-value store with a column-oriented concepts) are used.
- Often used for storing very large amounts of data (e.g. control system archive).
- Can scale to hundreds of TB.
- Apache Cassandra and Apache HBase are very popular.

Apache Cassandra

- Started at Facebook, released to the public in 2008.
- Apache top-level project since 2010.
- Data-model is similar to Google Bigtable.
- Used by
  - CERN
  - eBay
  - Hewlett Packard
  - IBM
  - Netflix
  - Spotify
  - Twitter

MySQL vs. Cassandra for Archiving

- MySQL supposedly is the fastest SQL database (using the transaction-less MyISAM engine).

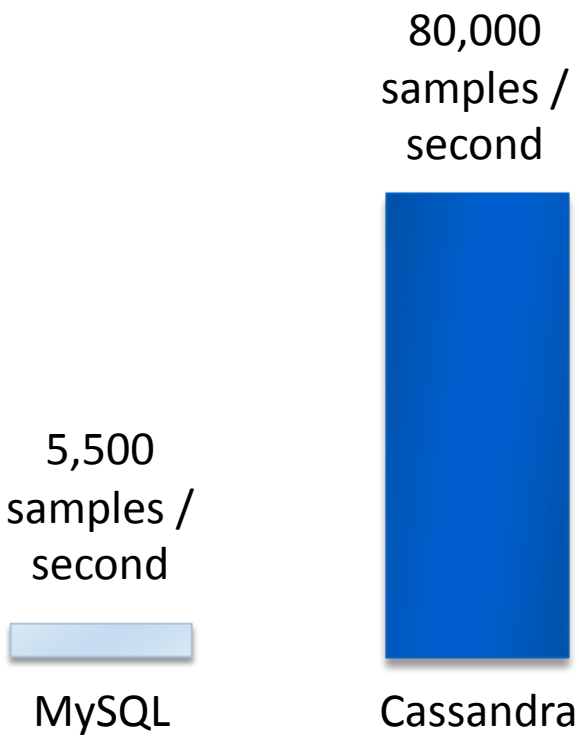

Cassandra Scalability

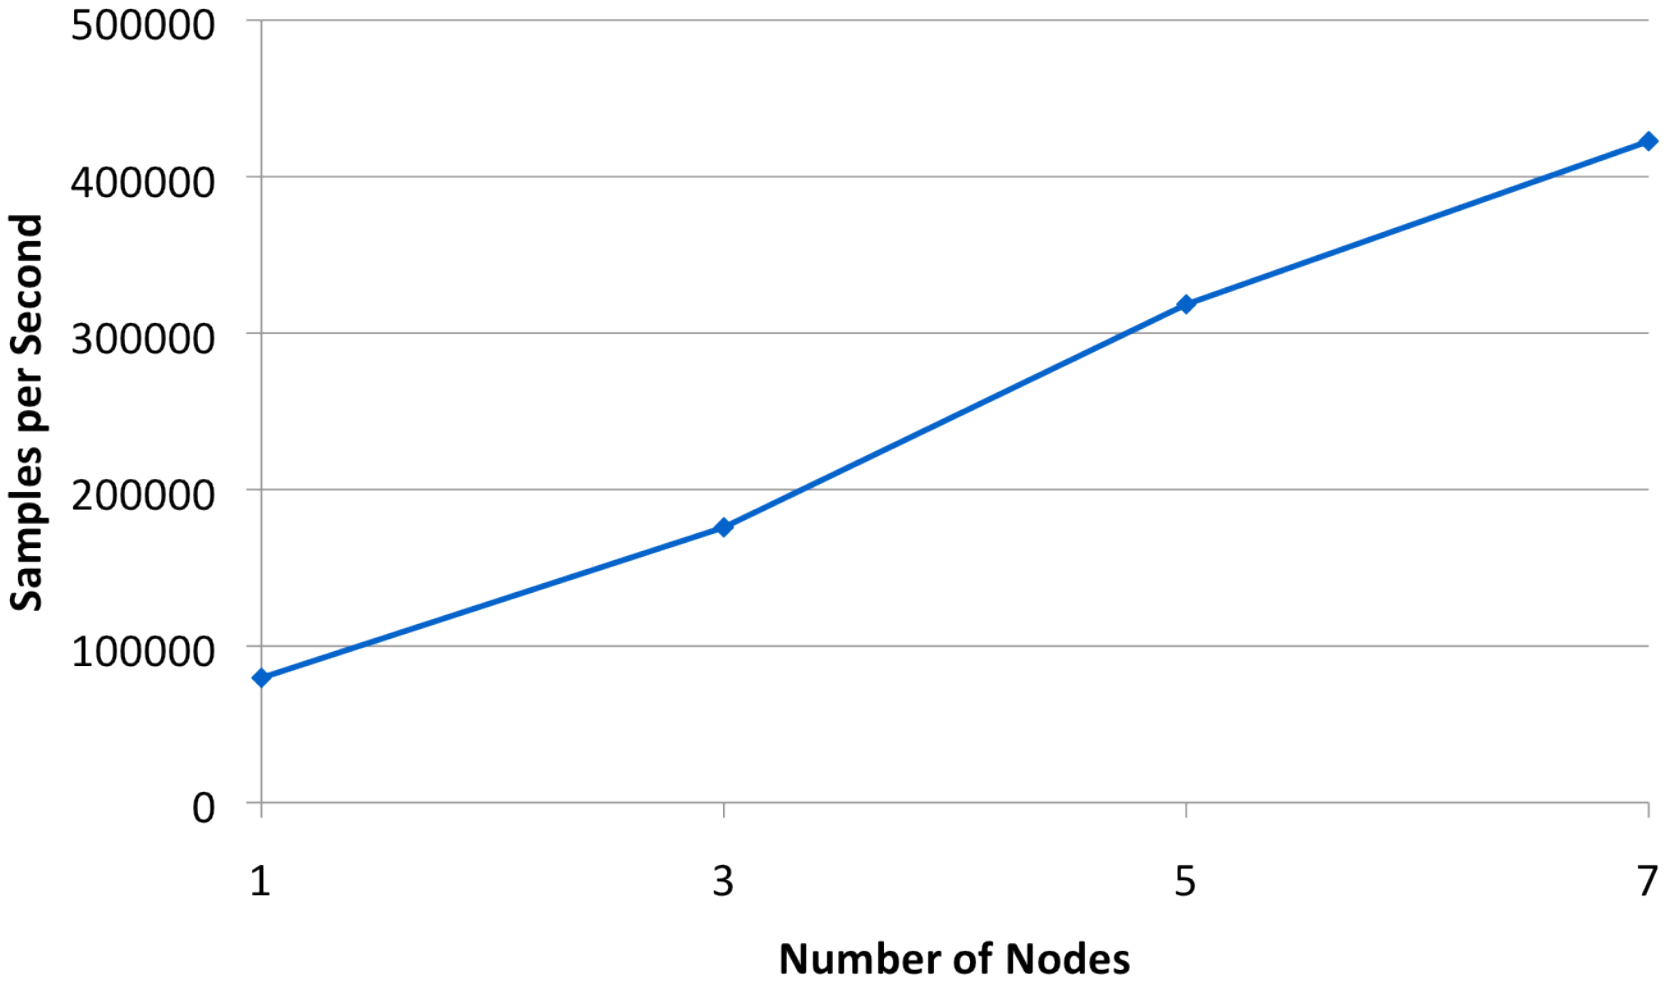

Beyond NoSQL

- Combine the best of both worlds:
  - Complex queries where needed.
  - High performance and scalability where possible.
- Google F1:
  - Consistent transactions and queries across shards.
  - Built on top of Google Spanner (descendant of Google Bigtable)
  - Unfortunately not available to the public.
- Apache Cassandra:
  - Cassandra Query Language (CQL): Not SQL, but similar.
  - Cassandra storage engine for MySQL: Use data from Cassandra like a regular table in MySQL (your mileage may vary).

Summary

- For many applications, SQL databases are not optimal and come at a high price.
- MongoDB may be interesting if you want to store meta-data.
- Apache Cassandra if good for storing large amounts of measurement data (e.g. control system archive).
- In the future, NoSQL and traditional SQL concepts might converge.

More Information

- Cassandra Archiver for CSS
  - Our free, open-source solution for archiving control-system data.
    http://oss.aquenos.com/epics/cassandra-archiver/

- Contact the speaker at
  sebastian.marsching@aquenos.com

References

- MongoDB website; http://www.mongodb.org/
- Apache Cassandra website; http://cassandra.apache.org/
- S. Marsching, “Scalable Archiving with the Cassandra Archiver for CSS”, ICALEPCS’13, San Francisco, October 2013; http://www.jacow.org/
- J. Shute et al., “F1: A Distributed SQL Database That Scales”, Proceedings of the VLDB Endowment, Vol. 6, No. 11, Trento, August 2013; http://research.google.com/pubs/pub41344.html