Contribution ID: **64** Type: **not specified**

Database clusters for sync and share services

Monday 30 January 2017 14:00 (20 minutes)

Sync and share services like ownCloud often rely on database backends for storing metadata. Those databases should offer a high availability and performance. With a clustered database, both of these requirements can be fulfilled.

One method of running a database cluster is a master-master Galera replication. In real life, the database performs more robust, if write requests are sent to a single database node and read requests are distributed among all other nodes.

In such a setup we expect a read-only workload to scale linearly with the number of nodes. The drawback: writes have to be replicated to all participating nodes. Thus, a scalable performance can only be expected if the rate of modification is not too high. The scalability of the performance is therefore limited to low writes versus reads ratios.

This talk will address the question of how the performance of a database cluster scales with the number of nodes. We will investigate typical sync and share-workloads as well as more pathologic writes-reads-ratios. Additionally, the influence of the speed of the network interconnect will be measured. Finally, the quantitative difference between a single database server and a clustered solution is presented.

Author: Mr ANGENENT, Holger (University of Muenster)

Co-author: Dr WILMER, Andreas (University of Muenster)

Presenter: Mr ANGENENT, Holger (University of Muenster)

Session Classification: Technology