



Master Thesis

2015-11-05

MySQL Replication using Software-defined Networking

Context

Modern applications are more and more driven by data that needs to be stored and handled properly by database systems, e.g. to be available and reliable. MySQL, as a popular candidate of relational, sql-based database systems, provides data replication mechanisms to realise a backup or even a failover strategy if one MySQL instance crashes. Further, extensions like MySQL Cluster or third-party tools like Galera offer clustering support to MySQL, allowing horizontal scalability and improved data replication. MySQL server and client are communicating using a proprietary tcp based MySQL protocol.

On the underlying infrastructure, where e.g. a MySQL instance is hosted, modern techniques like software-defined networking (SDN) become increasingly important and offer new ways to control data flows in a computer network, even context and application aware.

Scope of the Thesis

So why not using a sdn-enabled network infrastructure to replicate a MySQL server? This thesis is about setting up a MySQL replication, using SDN to control the Client/Server interaction. A setup should consist of at least two MySQL instances, one sdn controller and a (software or hardware) sdn switch. For replicating traffic to those MySQL instances and for dealing with forwarding only one (merged or validated) reply to the caller, an existing sdn controller must be extended to interpret the MySQL protocol and take care of a deterministic behaviour of both database instances. The resulting setup should be evaluated regarding response times and failover times, compared to existing replication techniques like Galera. Based on a working replication setup, an approach for horizontal scaling MySQL with SDN can be sketched.

Requirements and Comments

If this thesis achieves good progress and outcome, its results are to be integrated in the CACTOS research project (<http://www.cactosfp7.eu>) which is released under an OpenSource license. For that reason, we appreciate if you are ready to OpenSource your results.

We focus on Linux based server operating systems. For that reason basic knowledge and experience with Linux systems is recommended (but no must).

If you are interested in this or similar theses, please contact Christopher Hauser either by mail or directly in his office.

mail: christopher.hauser@uni-ulm.de
office: Uni West, 43.2.209

**Faculty of Engineering and
Computer Science**

Institute of Information
Resource Management