MySQL Conference & Expo

Presented by MySQL Conference & Expo O'REILLY

MySQL Replication Tutorial

Lars Thalmann
Technical lead
Replication, Backup, and Engine Technology

Mats Kindahl
Lead Developer
Replication Technology

MySQL Conference and Expo 2008

MySQL®
Conference & Expo

Presented by

MySQL Conference & Expo

O'REILLY

Concepts

MySQL Replication

Why?

1. **High Availability**
   Possibility of fail-over
2. **Load-balancing/Scale-out**
   Query multiple servers
3. **Off-site processing**
   Don't disturb master

How?

Snapshots (Backup)

1. **Client program mysqldump**
   With log coordinates
2. **Using backup**
   InnoDB, NDB

Binary log

1. **Replication**
   Asynchronous pushing to slave
2. **Point-in-time recovery**
   Roll-forward

Terminology

Master MySQL Server

- Changes data
- Has binlog turned on
- Pushes binlog events to slave after slave has requested them

Slave MySQL Server

- Main control point of replication
- Asks master for replication log
- Gets binlog event from master

Binary log

- Log of everything executed
- Divided into transactional components
- Used for replication and point-in-time recovery

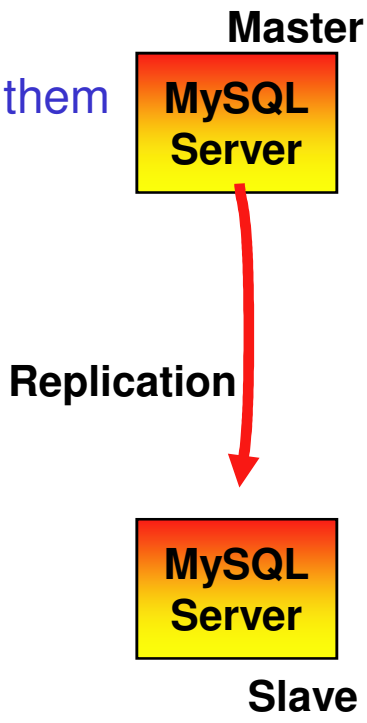

Terminology

Synchronous replication

- A transaction is not committed until the data has been replicated (and applied)
- Safer, but slower
- This is available in MySQL Cluster

Asynchronous replication

- A transaction is replicated after it has been committed
- Faster, but you can in some cases loose transactions if master fails
- Easy to set up between MySQL servers