



MySQL HA vs. HA

DOAG Konferenz 2016, Nürnberg

Oli Sennhauser

Senior MySQL Consultant, FromDual GmbH

oli.sennhauser@fromdual.com



www.fromdual.com

Über FromDual GmbH

Support



Beratung



remote-DBA



Schulung

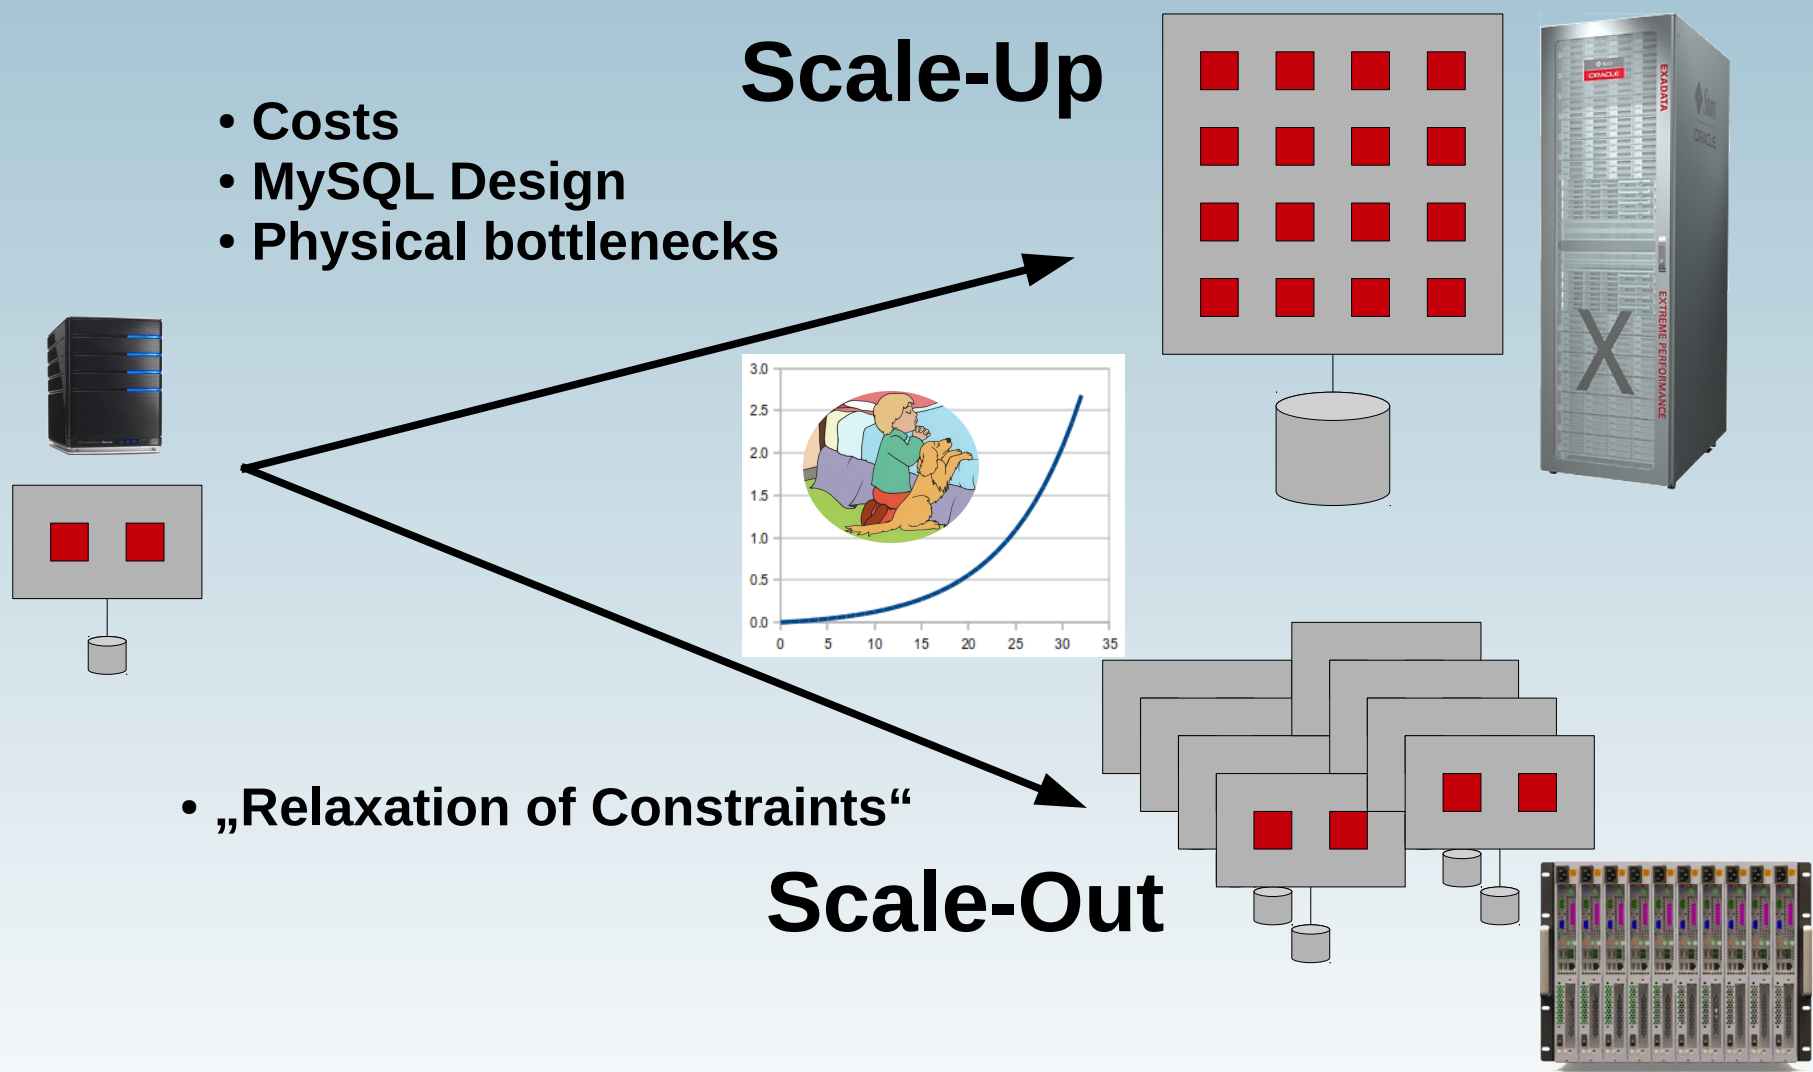


Contents

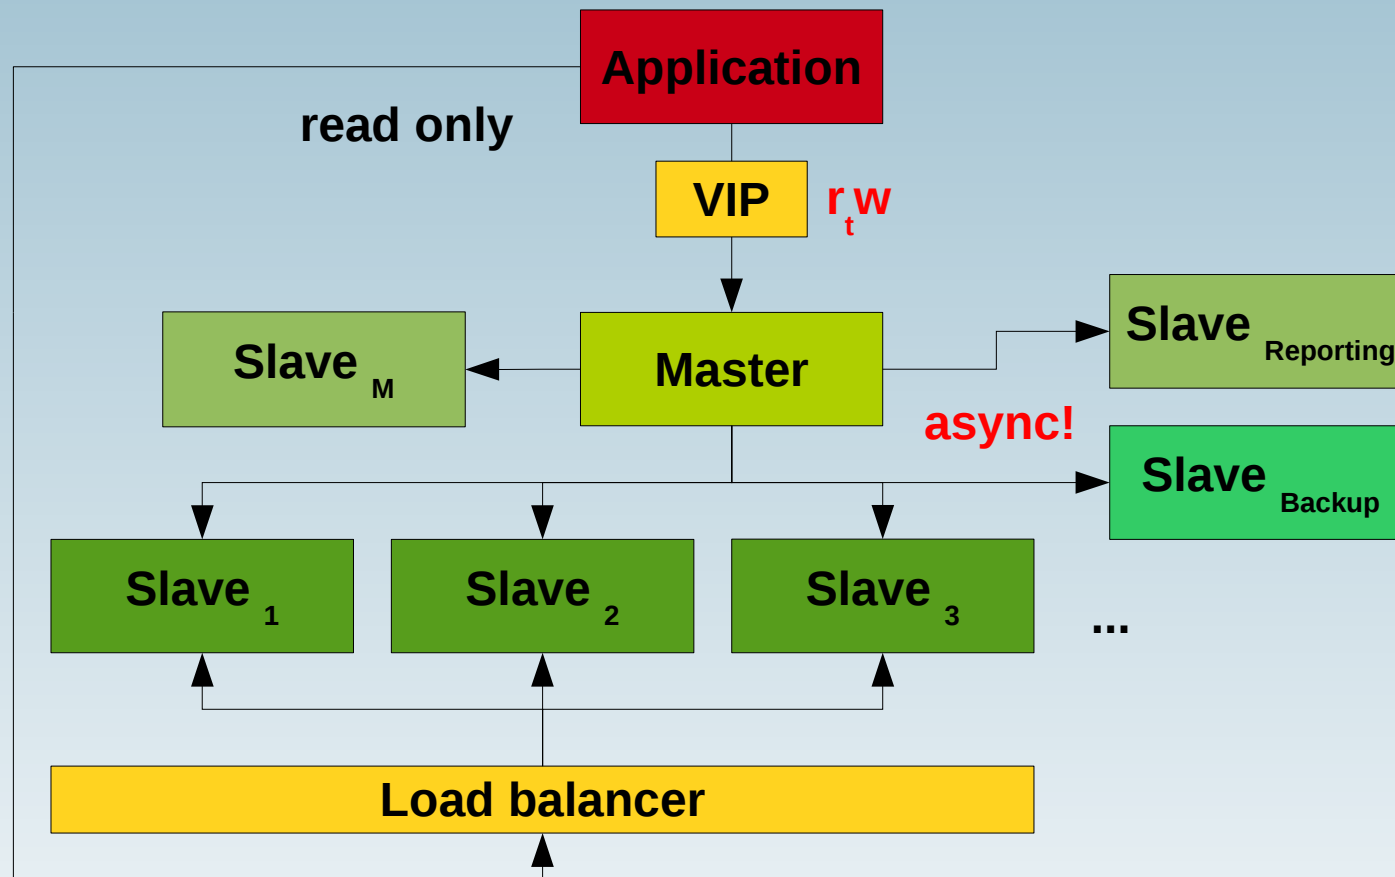
Various MySQL High Availability (HA) Solutions

- **MySQL Replication / MySQL Scale-Out**
- **High-Availability with Replication**
- **Master-Master Replication**
- **Active/passive fail-over with SAN**
- **Active/passive fail-over with DRBD**
- **Galera (synchronous) Replication**

MySQL Scale-Out vs Scale-Up www.fromdual.com

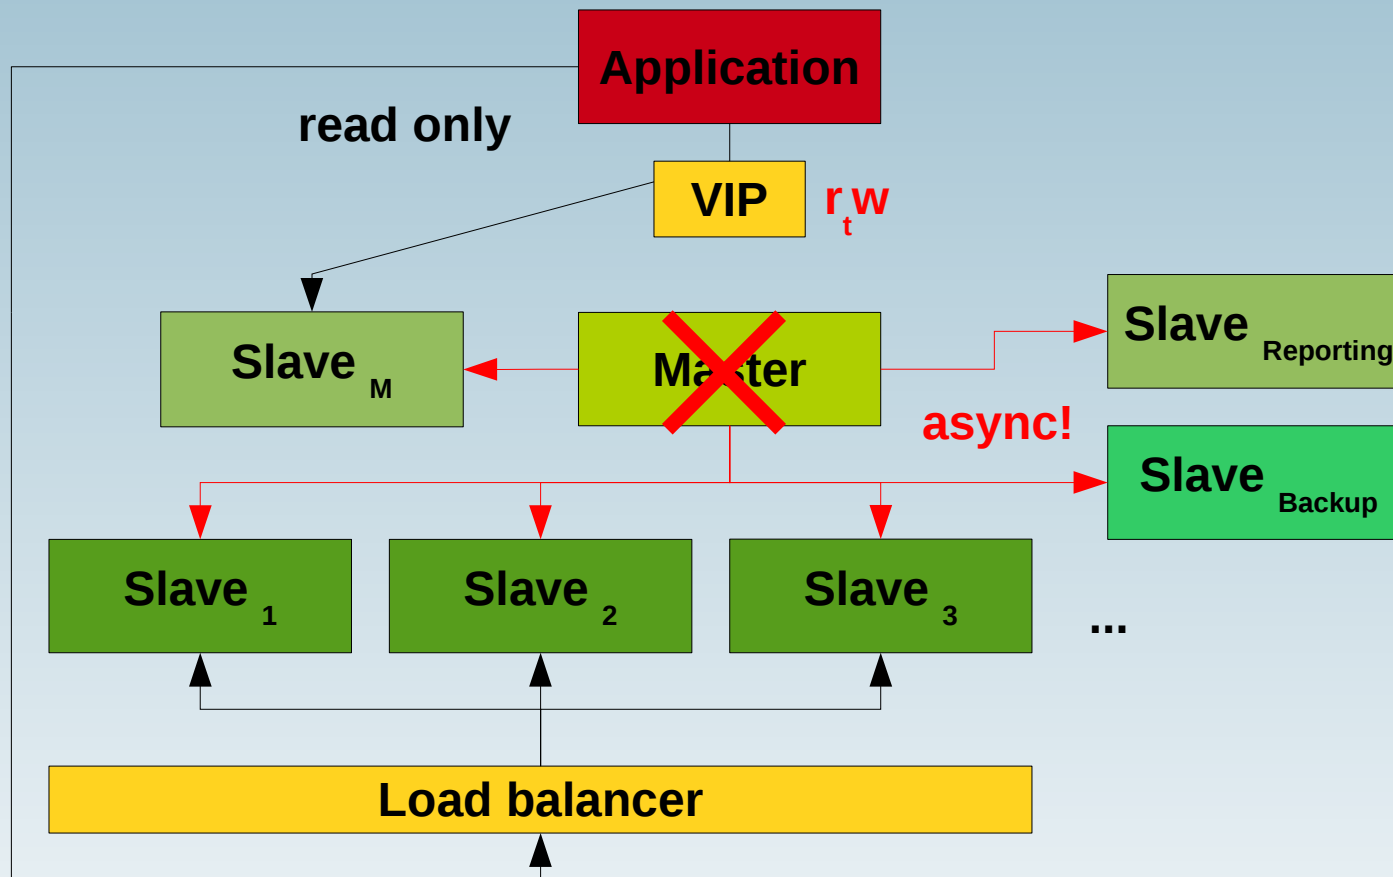


High-Availability with Replication



- **Fail-over?**

Replication fail-over



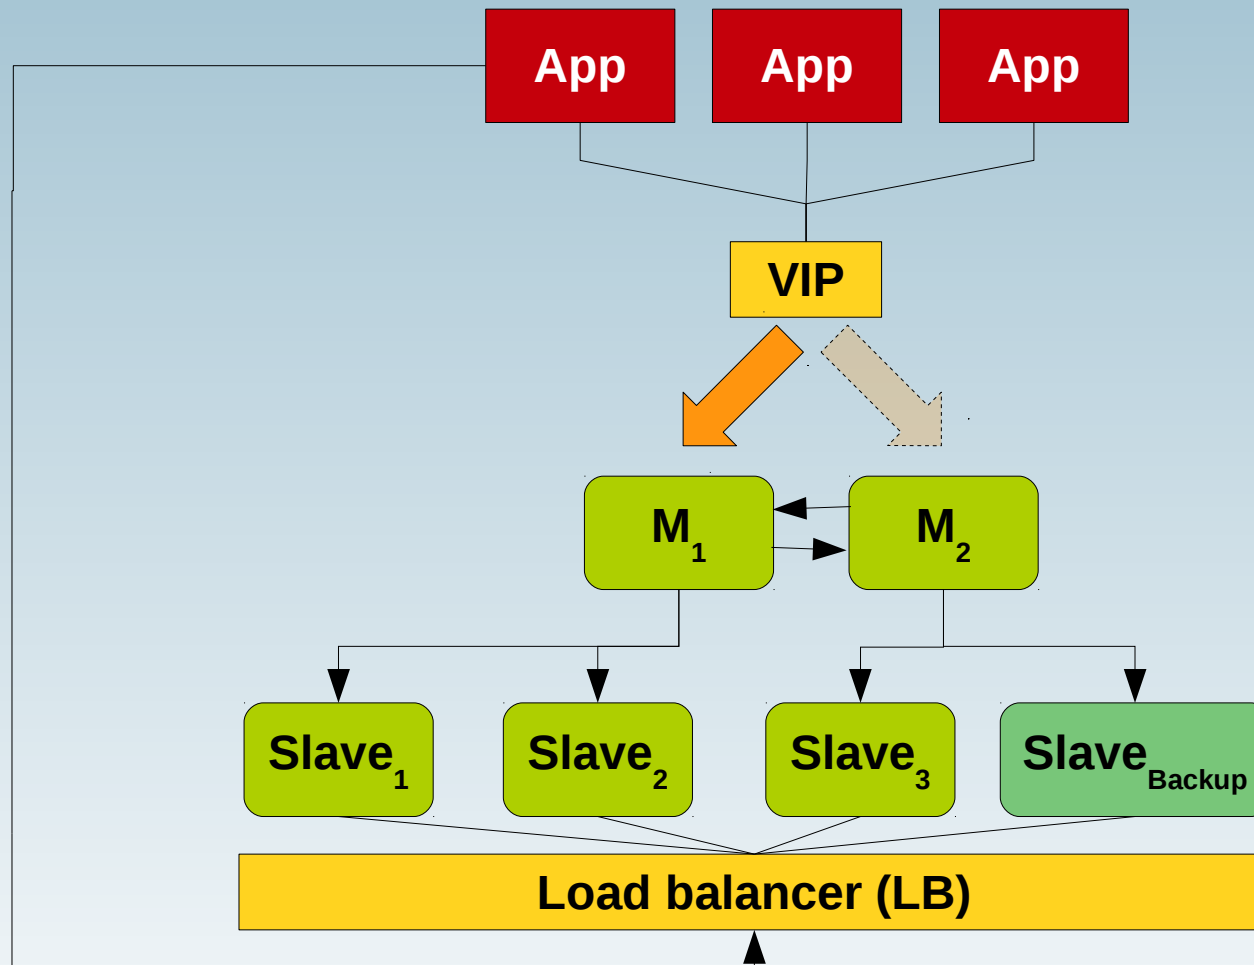
Advantages / Disadvantages

- + Simple „standard“ Set-up
- + Works very well if $r \gg w$
- + Fail-over site is already warm/hot!
- - Delay Master/Slave (asynchronous)
- - Slave lagging (Slave is sometimes bottleneck)
- - Data **in**-consistencies between Master and Slave?
- - If master fails → which Slave becomes new master?

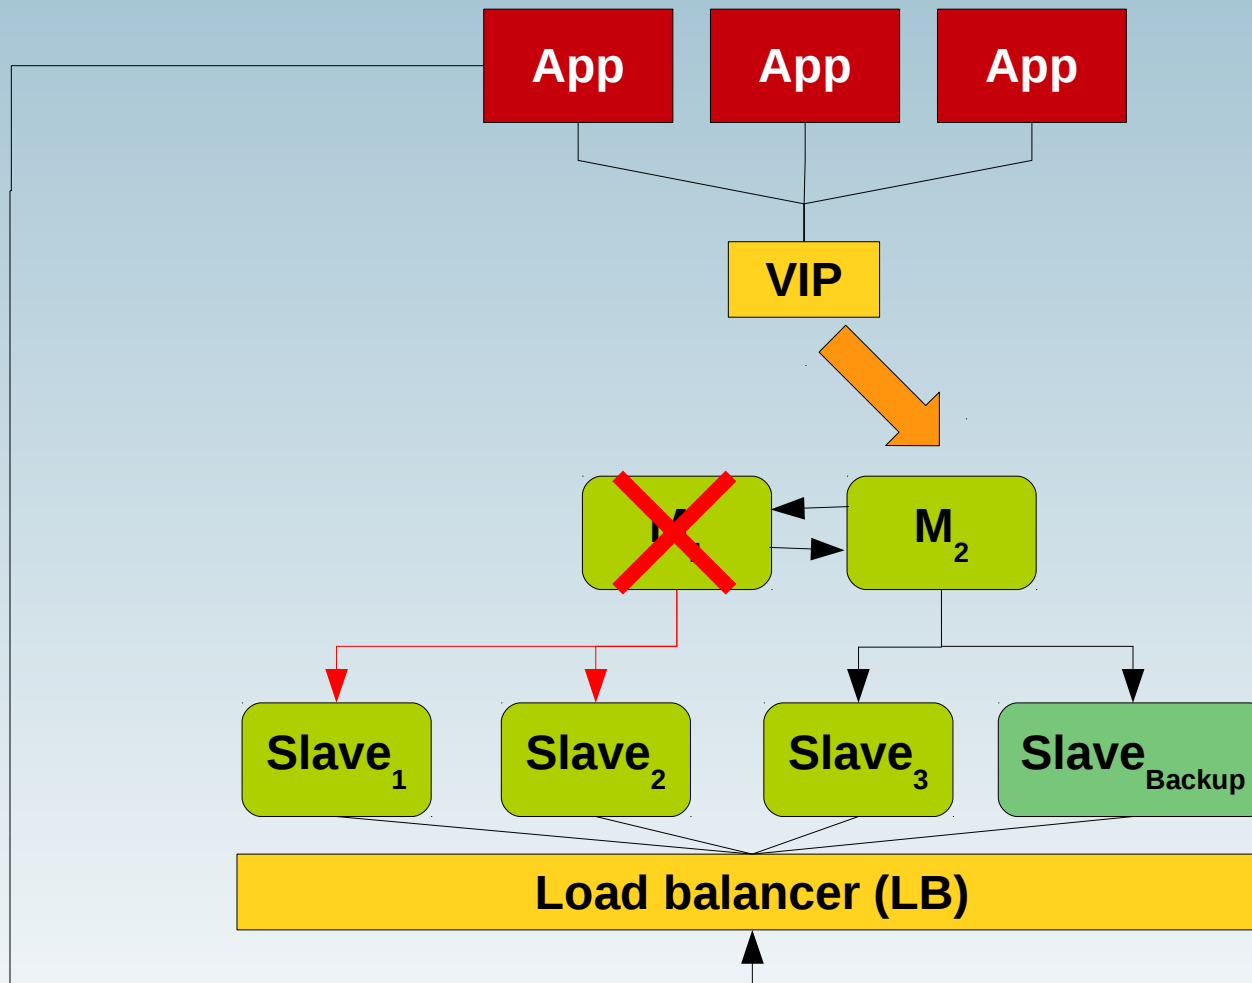
Switch → a lot of work, delicate!

There are tools to help (MMM v1/v2, MHA, Tungsten, MySQL utilities, ...)

Master-Master Replikation



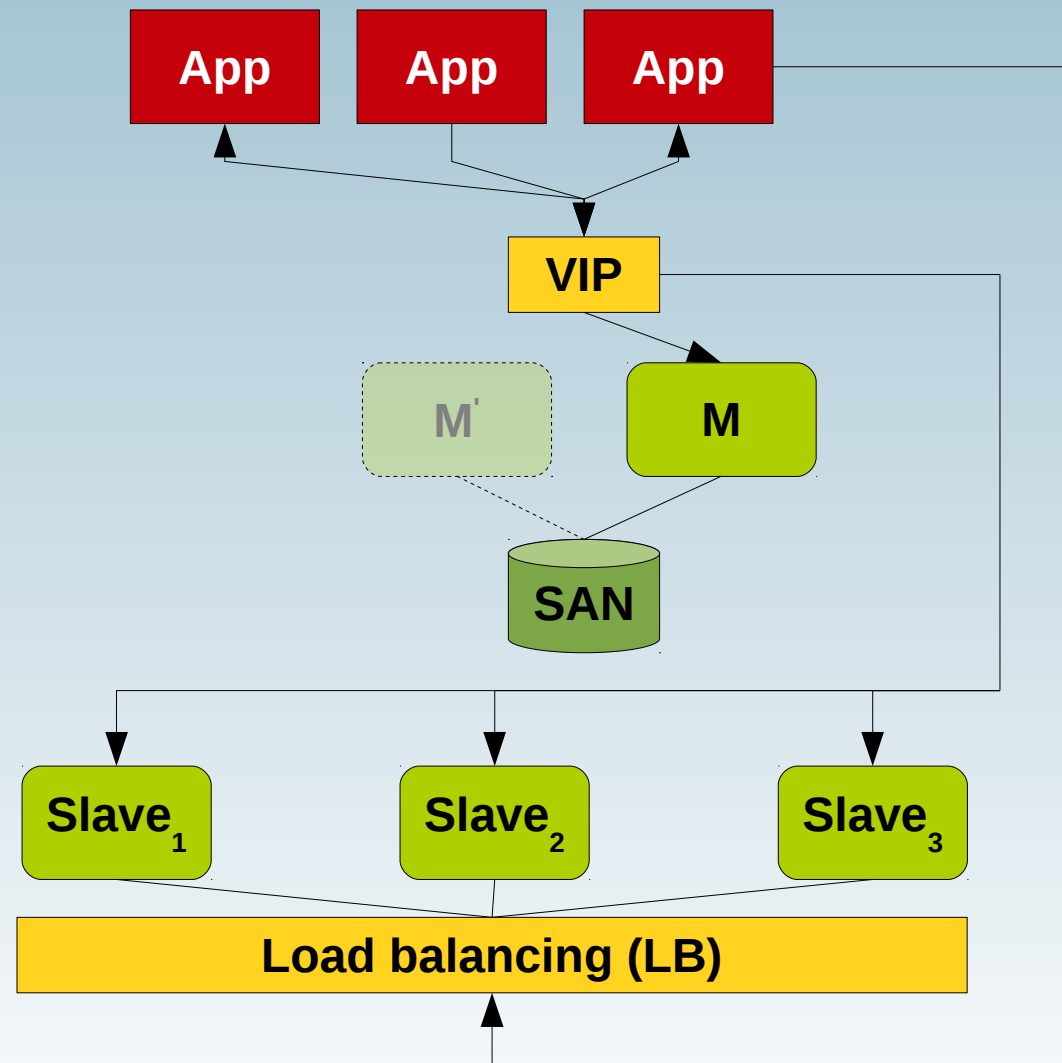
Master-Master Replication



Advantages / Disadvantages

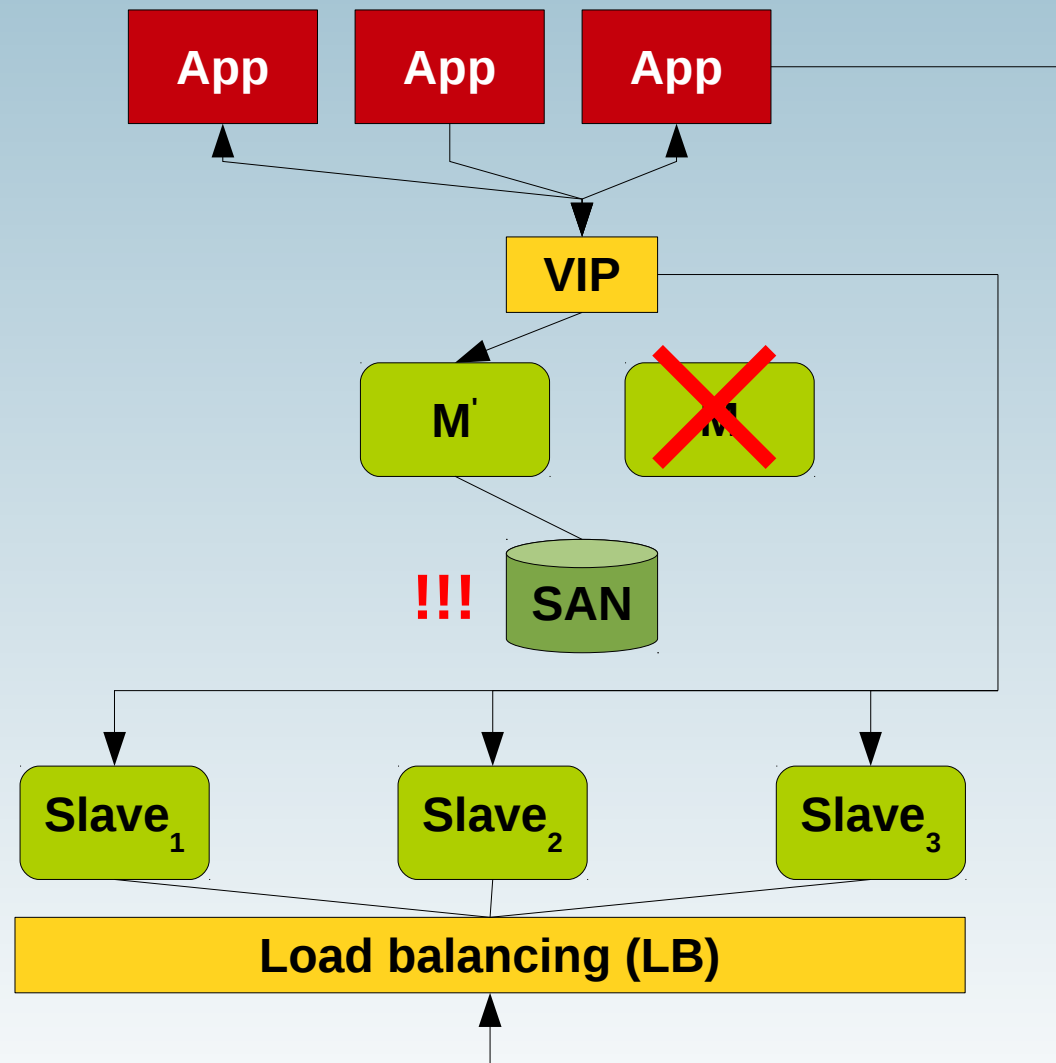
- + Only slightly more complex than Master/Slave
- + Works very well if $r \gg w$
- + Fail-over site is already warm/hot!
- - Delay Master₁/Master₂ (asynchronous)
- - Master₂ lagging (Slave is sometimes bottleneck)
- - Data **in**-consistencies between Master₁ and Master₂?
- - Careful when writing on both Masters!
 - - Data **in**-consistency possible because of asynchronous MM replication
- - You will NOT get more I/O throughput!
- - A little more complicated to (re-)set-up

Active/passive fail-over with SAN



Active/passive fail-over with SAN

- SPOF!

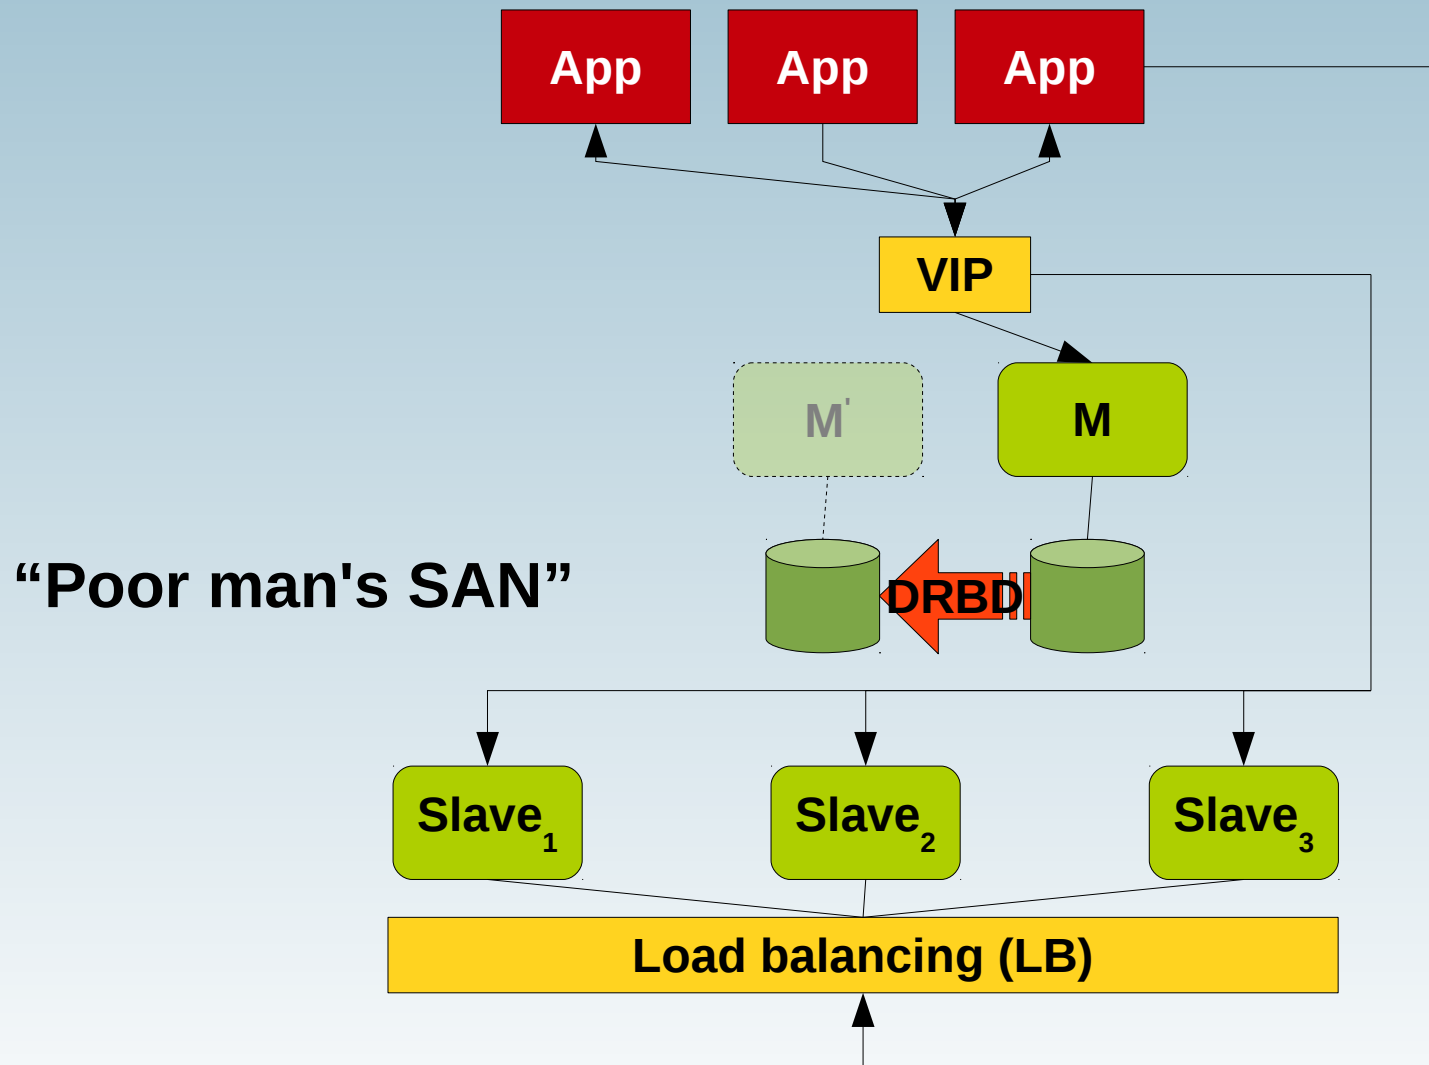




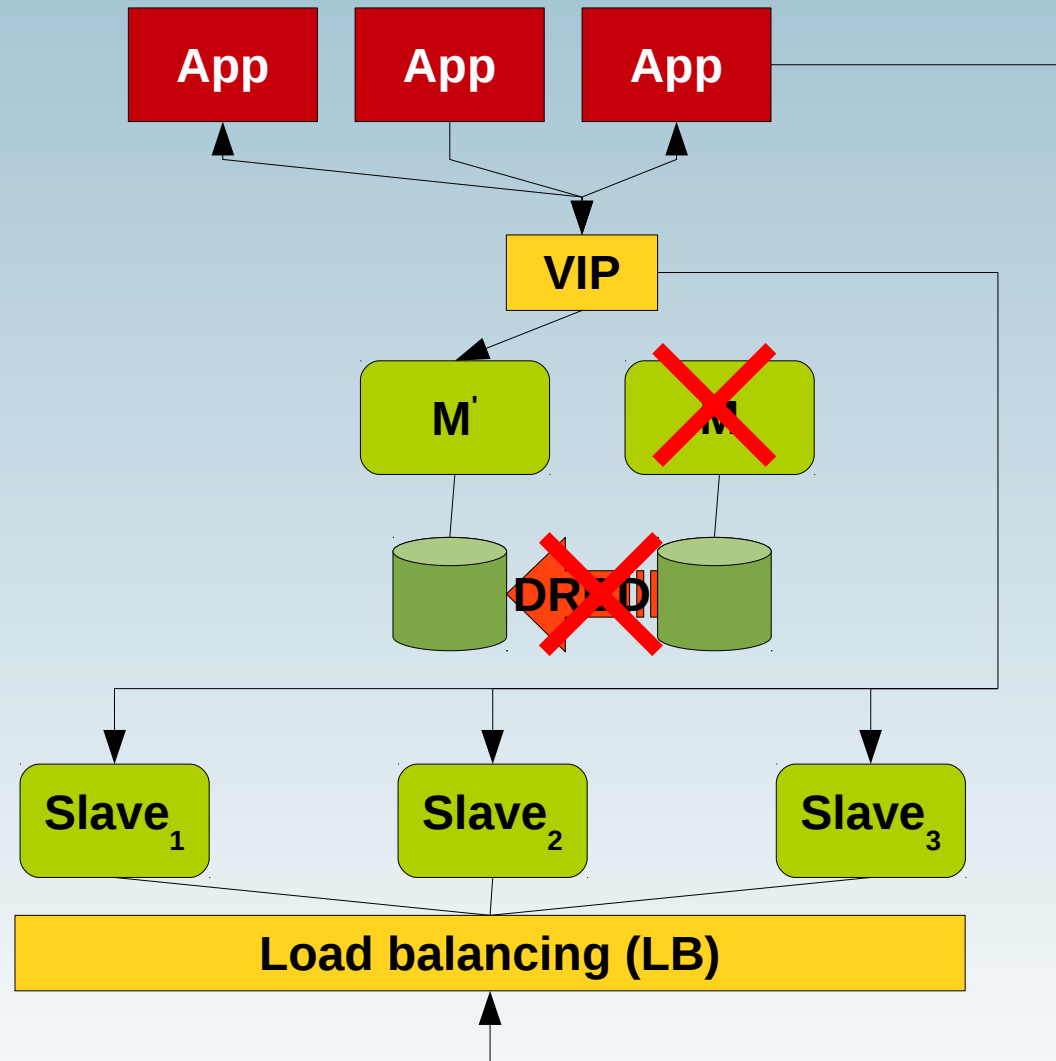
Advantages / Disadvantages

- + Synchronous “replication”
- + I/O throughput depends on SAN (I/O system)
- + No data **IN**-consistencies possible
- + Only one possible data source
- + Slaves are automatically and properly fail-overed
- - SAN and Filesystems are SpoFs!
- - Expensive if SAN are not available yet
- - SAN's are not easy to handle!
- - Other site is cold after fail-over!
- - Half of the hardware is idling
- - Far more complex to set-up
- - Need Unix know-how/root rights

Active/passive fail-over with DRBD



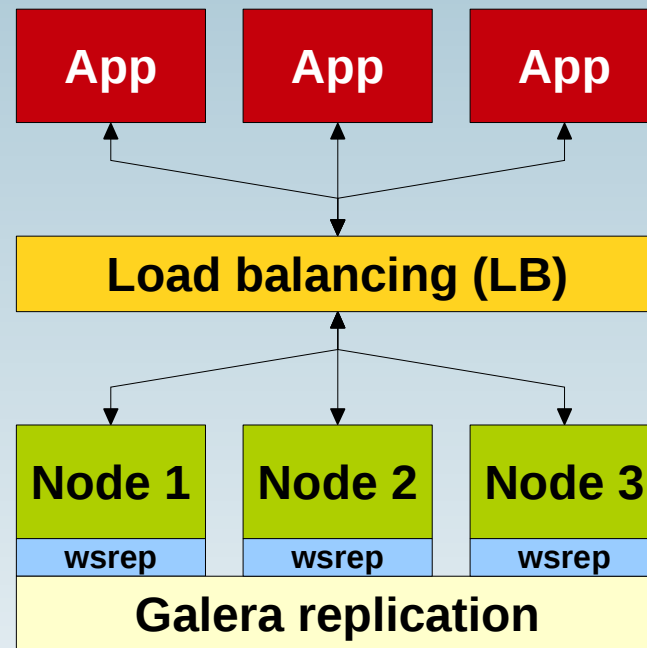
Active/passive fail-over with DRBD



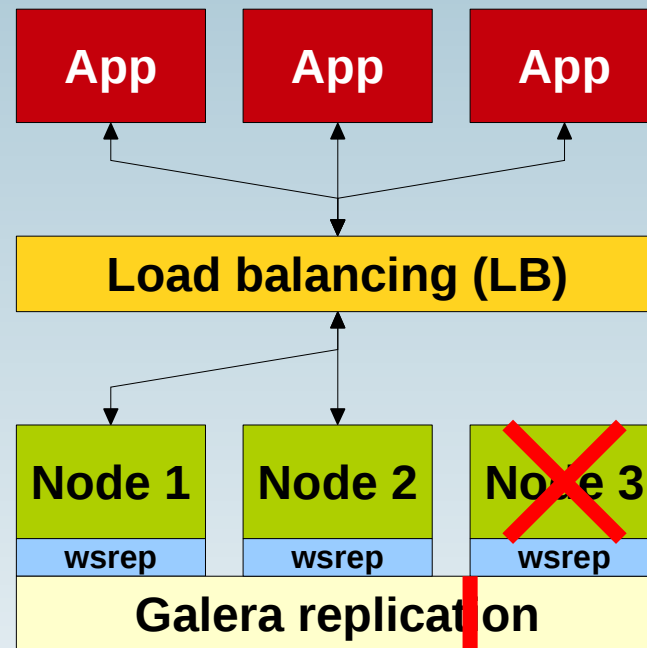
Advantages / Disadvantages

- + Synchronous replication
- + No data **IN**-consistencies possible
- + Only one possible data source
- + Slaves are automatically and properly fail-overed
- - Filesystem is SpoF!
- - I/O throughput possibly lower?
- - DRBD can break! → monitor it...
- - Other site is cold after fail-over!
- - Half of the hardware is idling
- - Far more complex to set-up
- - Need Unix know-how/root rights

Galera (synchronous) Replication



Galera (synchronous) Replication



Advantages / Disadvantages

- + Synchronous replication: No lost transaction
- + Based on InnoDB SE (other SE theoretically possible)
- + Active-active real multi-master topology: Read and write to any cluster node
- + Automatic membership control
- + True parallel replication, on row level: No slave lag
- + Read scalability (Read Scale-Out!) and write improvements (+ SSD)
- + Rolling Restart (Upgrade of Hardware, O/S, DB release, etc.)
- - No original MySQL binaries → Codership MySQL binaries
- - Be aware of Hot Spots on rows: Higher probability of deadlocks
 - - → application must be cluster aware!
- - Slowest node is pace maker

Q & A



www.fromdual.com



Fragen ?

Diskussion?

Wir haben Zeit für ein persönliches Gespräch...

- **FromDual bietet neutral und unabhängig:**
 - **Beratung**
 - **Remote-DBA**
 - **Support für MySQL, Galera, Percona Server und MariaDB**
 - **Schulung**

www.fromdual.com/presentations