



# MySQL für Oracle DBAs

DOAG Webinar 14. Juni 2013

**Oli Sennhauser**

Senior MySQL Consultant, FromDual GmbH

[oli.sennhauser@fromdual.com](mailto:oli.sennhauser@fromdual.com)

# Über FromDual GmbH

- FromDual bietet neutral und unabhängig:
  - Beratung für MySQL
  - Support für MySQL und Galera Cluster
  - Remote-DBA Dienstleistungen für MySQL
  - MySQL Schulungen
- Oracle Silver Partner (OPN)
- Mitglied der SOUG, DOAG, /ch/open



www.fromdual.com

# Inhalt

## MySQL für Oracle DBAs

- › **Einsatz von MySQL**
- › **Installation, Konfiguration, Starten/Stoppen**
- › **Architektur, Storage Engines**
- › **InnoDB**
- › **Monitoring, Logging**
- › **Backup, Restore, Point-in-Time-Recovery**
- › **Replikation**
- › **Hochverfügbarkeit**
- › **RAC für MySQL**

# Einsatz von MySQL

- **Wo wird MySQL eingesetzt:**
  - Facebook – 1 Mia User, 72 M QPS
  - Google – Adwords, Mia Umsatz/Jahr (M→O→M→F1)
  - Wikipedia – z. Zt. #6 weltweit
  - BörseGo – Online Börsenhandel
  - Playboy – Drupal CMS
  - EMKA – ERP, 1000 MA
  - V-Zug – Hybris Webshop
  - Buch.de – #2 online Buchhändler in D
  - Kikxxl – Callcenter, 1000 MA
  - Integrics – VoIP Lösungen, 1000e Anschlüssen
  - RePower – zig 1000 Windmühlen






# Installation

## Oracle: OUI (Oracle Universal Installer)

- **MySQL: Windows: Installer**

`C:\Program files\mysql\mysql-server-5.6\`

`C:\Program files\mysql\mysql-server-5.6\data`

- **MySQL Linux:**

- **Pakete: \*.rpm, \*.deb**

`/usr/`

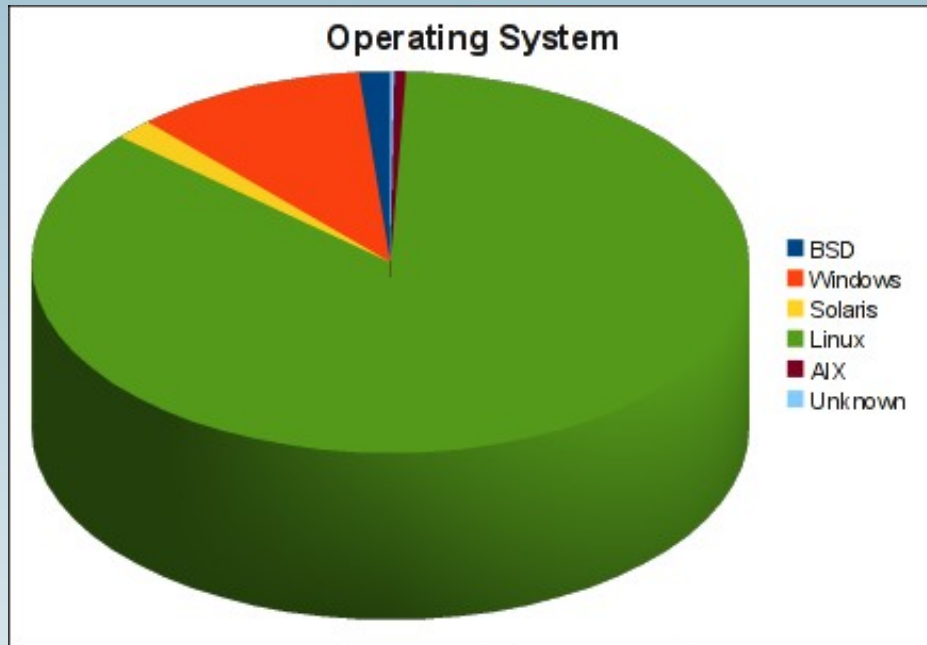
`/var/lib/mysql`

- **Binary Tar-Ball: mysql-5.7.1-linux-x86\_64.tar.gz**

- **Quellen → Kompilieren: cmake; make; make install**

- **MySQL Community vs. Enterprise, Drittanbieter**

# MySQL Plattform



- **85.7% Linux**
- **10.5% Windows**
- **1.7% Solaris**
- **1.4% BSD**
- **0.7% Others**

- **„Exotische“ Plattformen führen aus statistischen Gründen eher zu Problemen!**

# Konfiguration

Oracle: `$ORACLE_HOME/dbs/init<SID>.ora`

- MySQL:

- Windows:

- `C:\Program Files\mysql\mysql-server-5.6\my.ini`

- Linux:

- `/etc/my.cnf, /etc/mysql/my.cnf, $basedir/my.cnf`

- `!includedir /etc/mysql/conf.d/`

- `--defaults-file, --defaults-extra-file`

# Konfigurations-Parameter

- `my.cnf / my.ini`
- `mysql> SHOW GLOBAL VARIABLES;`
  - 5.1.69: 277
  - 5.5.31: 317
  - 5.6.11: 429
- `mysql> SET GLOBAL variable = value;`
- Kein Persistieren (`spfile`)



# Starten / stoppen

Oracle: `sqlplus / as sysdba STARTUP`

- **MySQL Linux:**

- Alt: `shell> /etc/init.d/mysql start|stop|restart`

- Neu: `shell> service mysql start|stop|restart`

- „von Hand“:

- `shell> bin/mysqld_safe &`

- `shell> bin/mysqld &`

- `shell> mysqladmin --user=root shutdown`

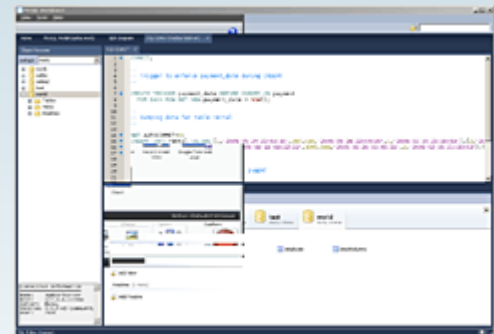
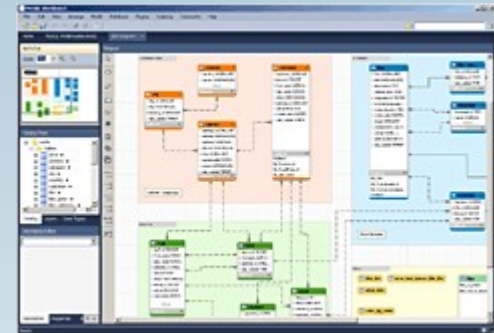
- **MySQL Windows**

- **Windows Service Utility**

- `cmd> net start|stop|restart MySQL`

# Tools

- **Tools:**
  - `sqlplus` → `mysql`
  - `srvmgr1` → `mysqladmin`
- **MySQL Workbench**
  - Admin
  - Query Browser
  - ER - Diagramme
- **Heidi SQL, phpMyAdmin**



# Prozess-Architektur

**Oracle: Multi-Prozess Modell  
Shared Memory**

**PMON, SMON, RECO, DBW0, LGWR, ARC0, ...**

- **MySQL: Multi-Thread Modell**

**mysqld**

- **Angel-Prozess**

**mysqld\_safe**

- **Vordergrund- und Hintergrund-Threads:**

# MySQL Thread Architektur

```
mysql> SELECT thread_id, name AS 'thread_name', type, processlist_user AS user
        FROM performance_schema.threads;
```

thread_id	thread_name	type	user
1	thread/sql/main	BACKGROUND	NULL
2	thread/innodb/io_ibuf_thread	BACKGROUND	NULL
3	thread/innodb/io_log_thread	BACKGROUND	NULL
4	thread/innodb/io_read_thread	BACKGROUND	NULL
11	thread/innodb/io_write_thread	BACKGROUND	NULL
14	thread/innodb/srv_error_monitor_thread	BACKGROUND	NULL
15	thread/innodb/srv_lock_timeout_thread	BACKGROUND	NULL
16	thread/innodb/srv_monitor_thread	BACKGROUND	NULL
17	thread/innodb/srv_master_thread	BACKGROUND	NULL
18	thread/innodb/srv_purge_thread	BACKGROUND	NULL
19	thread/innodb/page_cleaner_thread	BACKGROUND	NULL
20	thread/sql/signal_handler	BACKGROUND	NULL
22	thread/sql/one_connection	FOREGROUND	root
28	thread/sql/one_connection	FOREGROUND	oli

```
shell> ps -efL | egrep 'mysqld|PID'
```

UID	PID	PPID	LWP	C	NLWP	STIME	TTY	TIME	CMD
mysql	3248	1	3248	0	1	16:28	pts/0	00:00:00	/bin/sh bin/mysqld_safe
mysql	3925	3248	3925	0	23	16:28	pts/0	00:00:00	bin/mysqld
...									
mysql	3925	3248	4088	0	23	16:31	pts/0	00:00:00	bin/mysqld

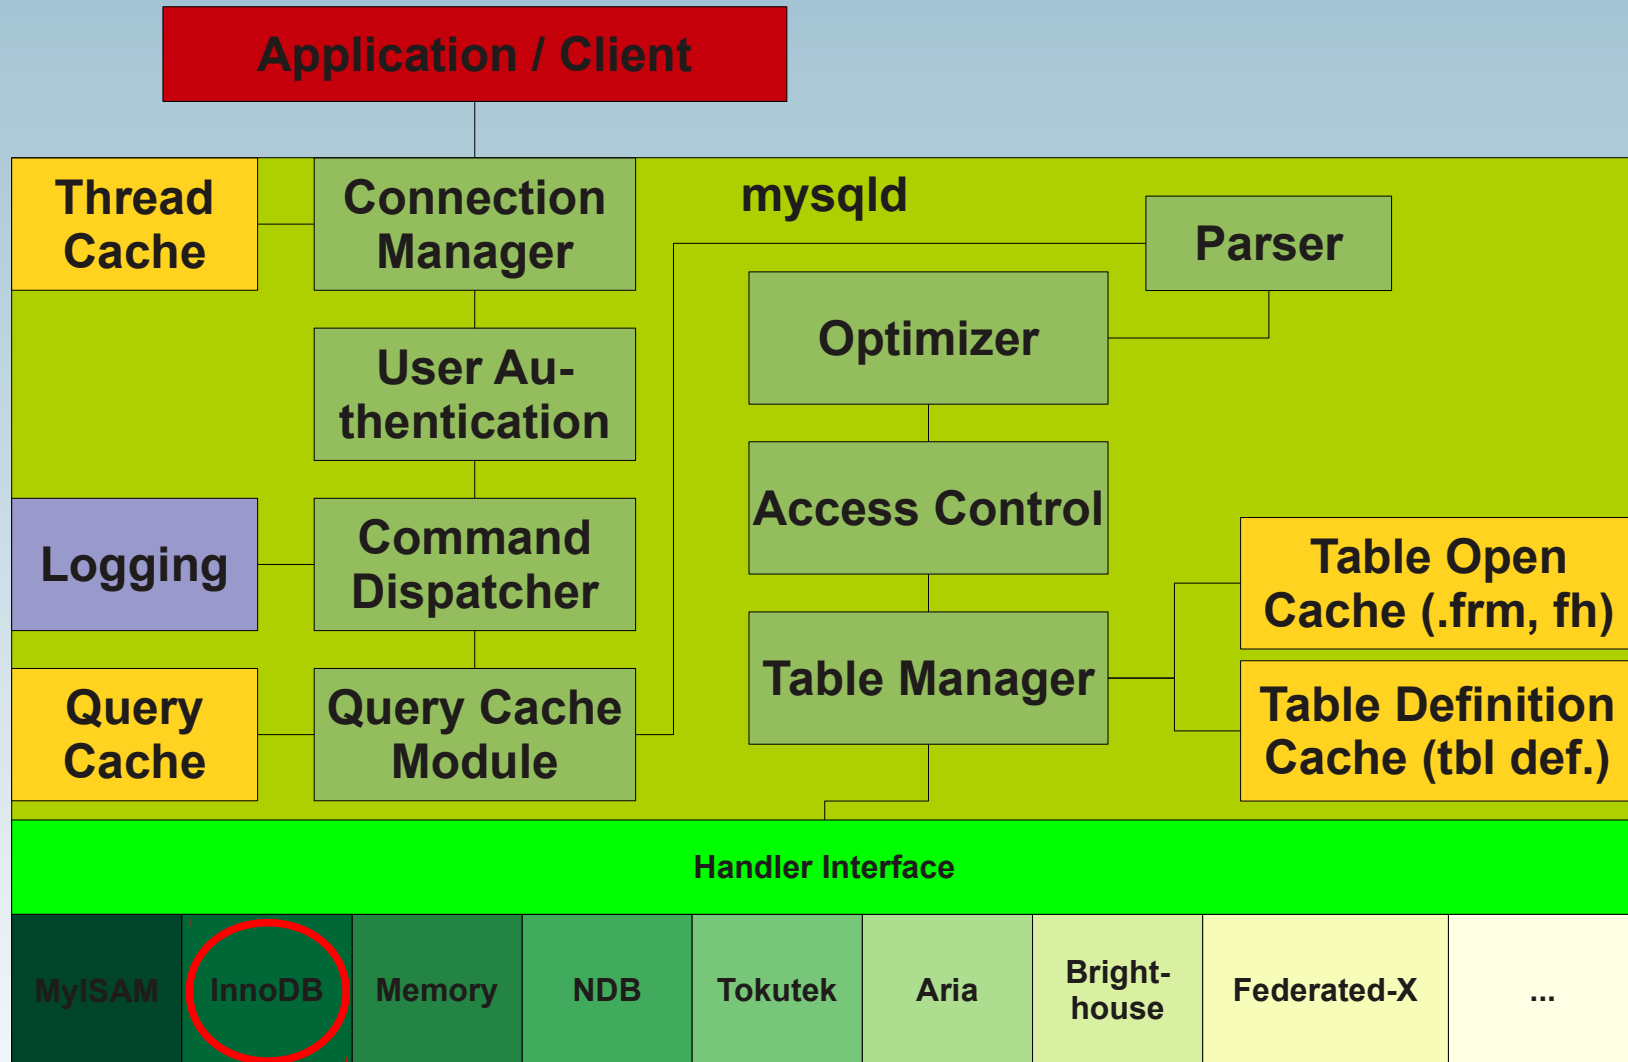
# Connections / Connectors

- **Verbindung**
  - In MySQL billig: oft KEIN Connection-Pooling
  - 1 Verbindung = 1 Thread → 1 Query → 1 Core
  - Thread Pool (1000e von Verbindungen)
- **Connectors:**
  - JDBC/ODBC
  - PHP, Perl, Python, Ruby, .NET

# User und Schema

- **User**
  - `'oli'@'localhost'` → **Unix Socket**
  - `'oli'@'127.0.0.1'` → **TCP Port**
  - `'oli'@'%'` → **TCP von überall her**
- **Privilegien**
  - **Global: \*.\***, pro Schema , pro Tabelle, pro Spalte
- **Schema (= Database)**
  - **Unabhängig vom User** (→ gehört System)

# Storage Engines



# InnoDB (default SE $\geq$ 5.5.)

- Transaktionen (ACID)
- Isolation Level (repeatable-read vs read-committed)
- Tablespaces
  - System TS = `ibdata1`
  - Table TS (`innodb_files_per_table = 1`)
- InnoDB: PK geclusterte Tabellen → IOT
- InnoDB Buffer Pool (16k Block Buffer)
- REDO Logs: `ib_logfile?` (5M default)



# Logging

- **Error Log (~ alert\_<SID>.log)**
  - /var/log/mysql\*
- **General Query Log**
  - `general_log = 1`
- **Slow Query Log**
  - `slow_query_log = 1`
- **Binary Log (~ archive.log)**
  - `log_bin = 1 / binary-log`
  - DML + DDL aller SE
- **(Transaktions-Log, ib\_logfile?, SE abhängig) (~ redo.log)**
  - DML InnoDB

# Logisches Backup

Oracle: `exp/imp` (vor Datapump)

- MySQL: `mysqldump/mysql`
- Jede Row wird angelangt:
  - Langsam
  - Restore SEHR langsam!

```
mysql> mysqldump --all-databases --single-transaction --master-data=1 > full_dump.sql
```

```
mysql> mysql < full_dump.sql
```

# Physisches Backup

Oracle: rman, ALTER TABLESPACE ... BEGIN|END BACKUP;

- **MySQL:**

- Snapshot mit LVM, BtreeFS oder ZFS

- `mylvmbbackup`

- Xtrabackup, MySQL Enterprise Backup (MEB)

```
shell> innobackupex /data/backups
```

```
shell> innobackupex --apply-log /data/backups/2012-11-13/
```

```
shell> innobackupex --copy-back /data/backups/2012-11-13/
```

```
shell> chown -R mysql:mysql /var/lib/mysql
```

- **Links:**

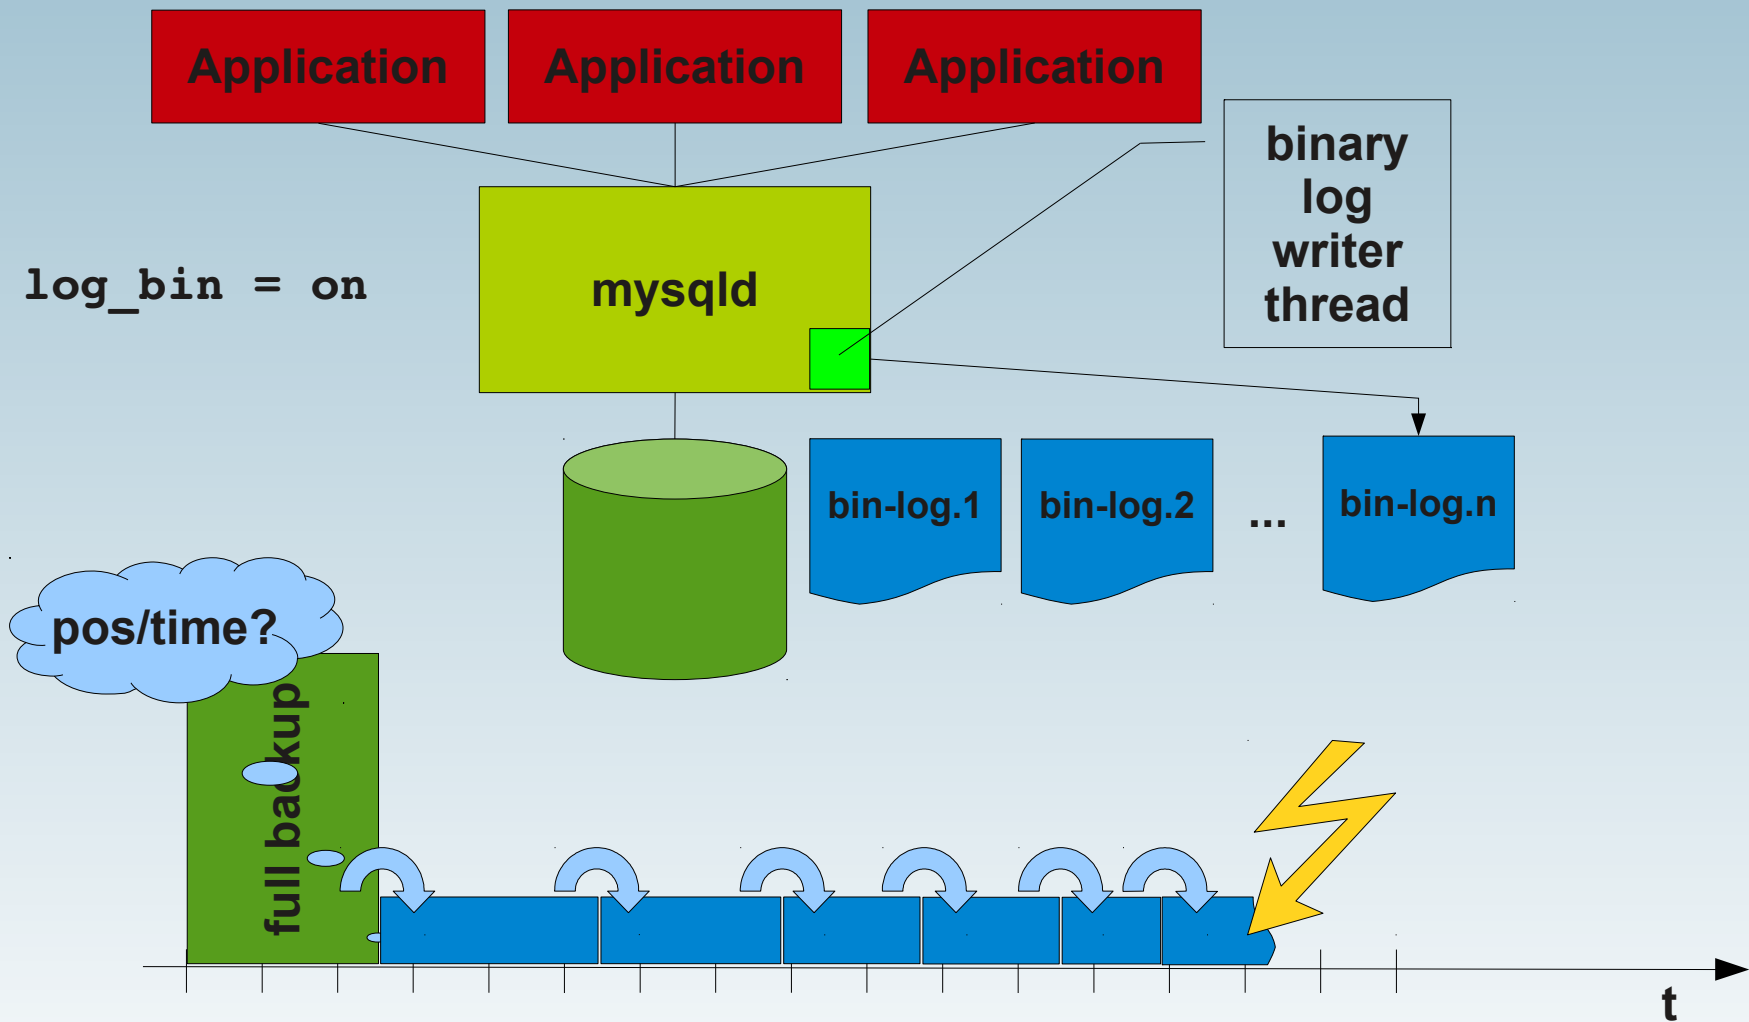
- <http://www.lenzg.net/mylvmbbackup/>

- <http://www.percona.com/doc/percona-xtrabackup>

# Binary Log

- ~ Oracle Archive Logs
- **MySQL Binary-Log:**
  - DDL + DML aller SE (nicht nur InnoDB)!
- **für:**
  - Replikation
  - Point-in-Time-Recovery (PiTR)
- **3 Varianten:**
  - Statement Based Replication (SBR)  $\leq$  5.0
  - Row Based Replication (RBR) ab 5.1
  - Transaction Based Replication (TBR) ab 5.6

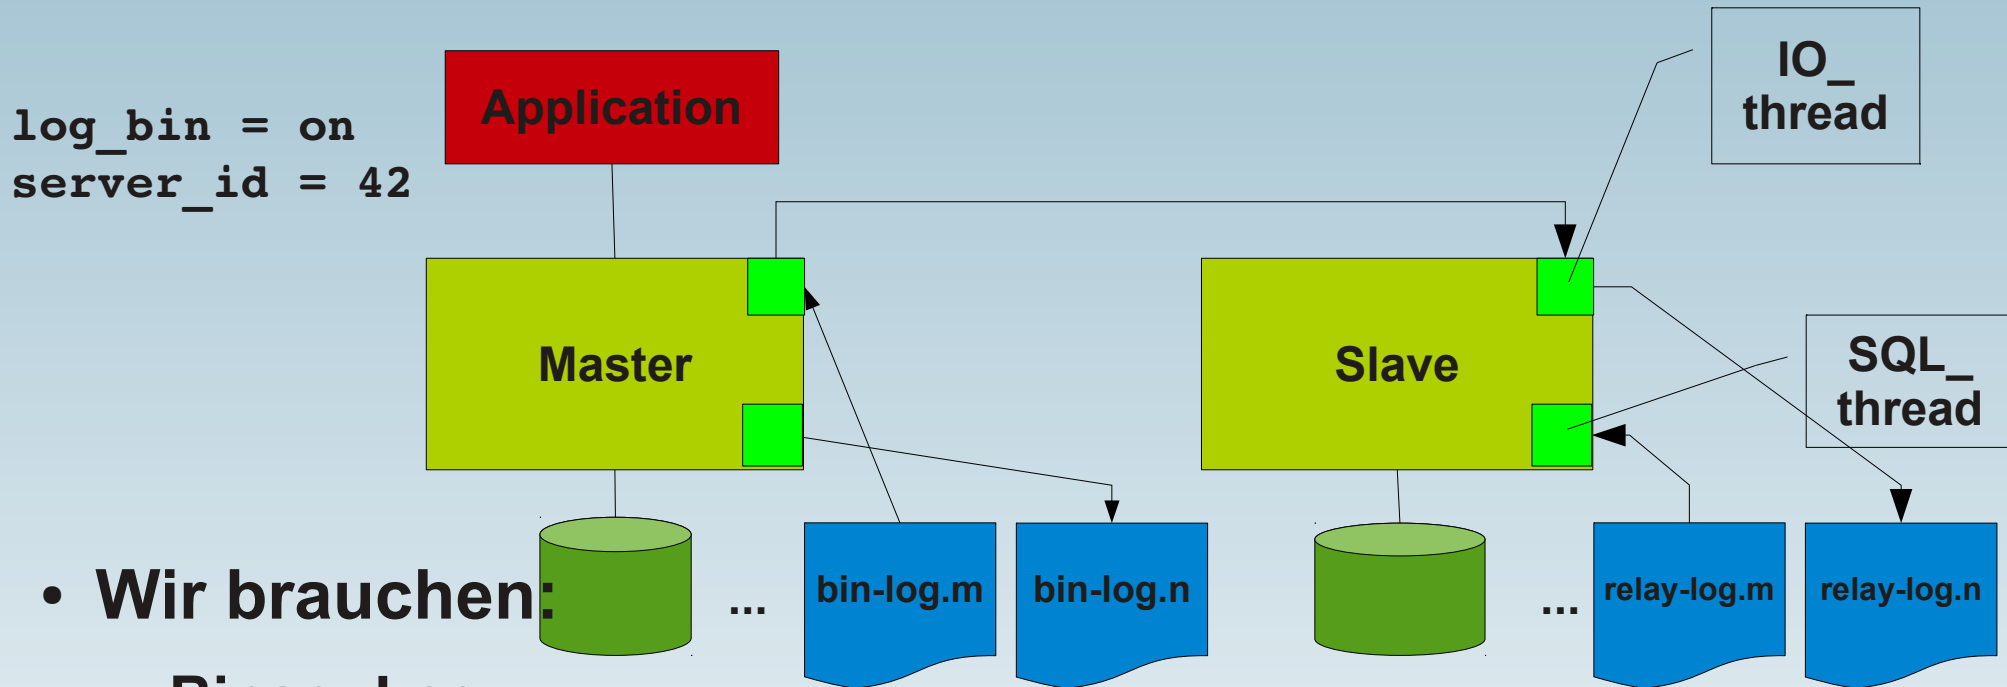
# Point-in-Time-Recovery (PITR)



# MySQL Replikation

- **MySQL Basis-Funktionalität**
- **Sehr einfach aufzusetzen (5 min)**
- **Streaming-Replication (kein Log Shipping)**
- **Basiert auf MySQL Binary Logs**
- **Braucht:**
  - **Unique server\_id (Master und Slave) (restart)**
  - **Binary Loggin auf Master einschalten (restart)**
  - **Replikations-User**
  - **Konsistentes Backup + Binary Log Position**

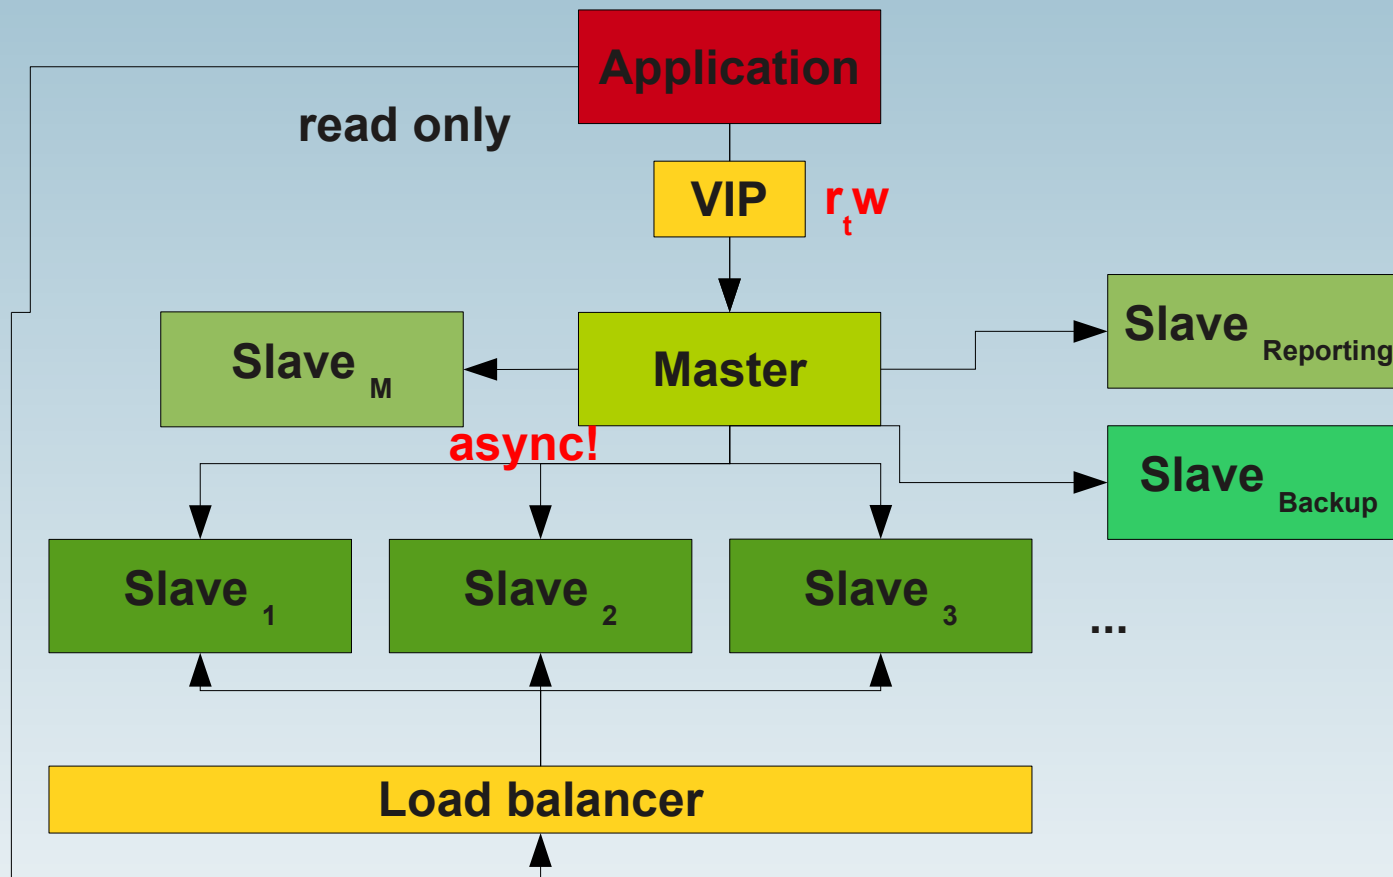
# Master – Slave Replikation



- **Wir brauchen:**

- Binary Log
- Server Id
- User für die Replikation (auf dem Master)
- Konsistentes Backup MIT Binary Log Position

# High-Availability mit Replikation

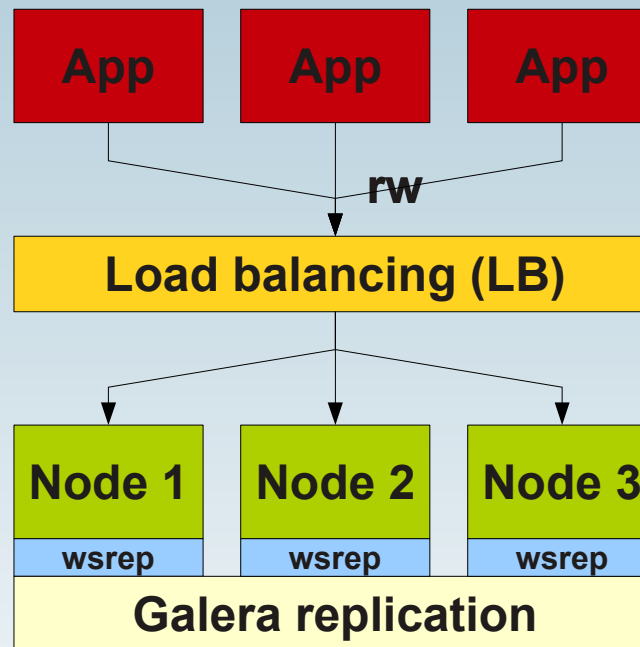




# RAC: Galera Cluster

## Oracle Real Application Cluster (RAC)

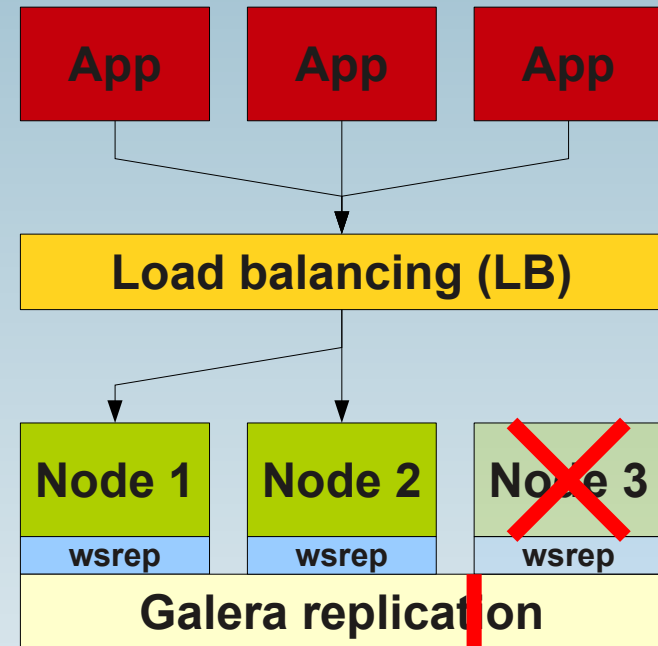
- **MySQL: Galera Cluster**



- **Shared-Nothing Architektur**

# Galera Cluster für MySQL

- Hardware-Ausfall
- Wartungsarbeiten
  - HW/OS/DB Upgrade
- 5x9 HA: 99.999%



# Monitoring

- OEM/DBC/Grid-Control/Cloud-Control
- MySQL Enterprise Monitor (MEM) €€€
- MySQL Performance Monitor (mpm)
- `mysql> SHOW GLOBAL STATUS;`
- Nagios / Icinga
- Links:



<http://www.mysql.com/products/enterprise/monitor.html>

<http://www.fromdual.com/mysql-performance-monitor>

<http://www.fromdual.com/download#nagios>

# Performance Tuning

- `mysql> SHOW GLOBAL STATUS;`
- `PERFORMANCE_SCHEMA`
- **Slow Query Log**
  - `slow_query_log = 1`
  - `long_query_time = 0.5`
  - `shell> mysqldumpslow -s t slow.log > profile`
- **Query Execution Plan:**
  - `mysql> EXPLAIN SELECT * FROM test;`

# Stored Programs

**Oracle: PL/SQL, Java**

- **MySQL:**

- **Stored Procedures**
- **Stored Functions**
- **User Defined Functions (UDF): C/C++**
- **Plugin: C/C++**

# Volltext-Suche

## Oracle: Kostenpflichtiges Modul?

- **MySQL: Standardmässig dabei!**

```
ALTER TABLE test ADD FULLTEXT INDEX (data);  
SELECT * FROM test WHERE MATCH data AGAINST('DBA');  
+-----+-----+  
| id | data |  
+-----+-----+  
| 1 | Wir suchen zur Zeit einen Support DBA! |  
+-----+-----+
```

# Q & A



[www.fromdual.com](http://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](http://www.fromdual.com/presentations)**