



MySQL für Oracle DBAs

SOUG Day 2016, Baden

Cédric Bruderer

MySQL Support Engineer, FromDual GmbH

`cedric.bruderer@fromdual.com`



Über FromDual GmbH

Support



Beratung



remote-DBA



Schulung



Über mich

- **Cédric Bruderer**
- **Ausbildung**
 - 2010 – 2014: Lehre zum Informatiker**
 - Teilnahme an den Schweizer
Berufsmesterschaften**
- **Junior Engineer**
 - In einem international tätigen Unternehmen.**
- **MySQL Support Engineer bei FromDual**
 - seit Oktober 2015**

Inhalt

MySQL für Oracle DBAs

- › **Branches und Forks**
- › **Installation und Konfiguration**
- › **Tablespaces**
- › **Connections, Threads und Connectors**
- › **Tools**
- › **User und Schema**
- › **Storage Engine**
- › **Backup, Restore, Recovery**
- › **Performance Tuning**
- › **Hochverfügbarkeit**

Wie wichtig ist MySQL heute?

- Welche DB habt Ihr heute im Einsatz?

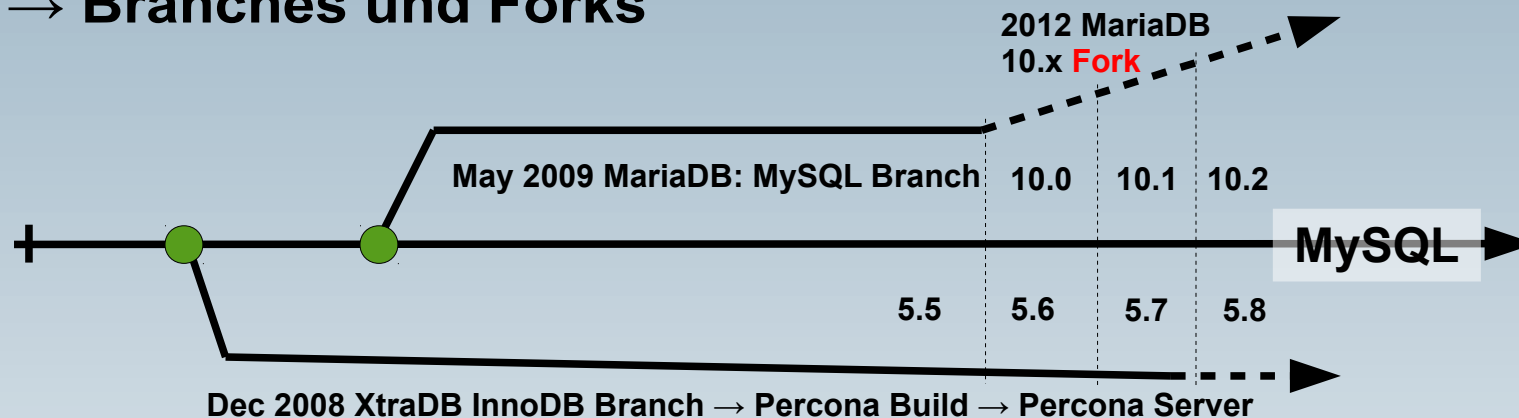
305 systems in ranking, June 2016

Rank			DBMS	Database Model	Score		
Jun 2016	May 2016	Jun 2015			Jun 2016	May 2016	Jun 2015
1.	1.	1.	Oracle	Relational DBMS	1449.25	-12.78	-17.11
2.	2.	2.	MySQL +	Relational DBMS	1370.13	-1.69	+91.78
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1165.81	+22.99	+47.76
4.	4.	↑ 5.	MongoDB +	Document store	314.62	-5.60	+35.57
5.	5.	↓ 4.	PostgreSQL	Relational DBMS	306.60	-1.01	+25.70
6.	6.	6.	DB2	Relational DBMS	188.57	+2.61	-10.12
7.	7.	↑ 8.	Cassandra +	Wide column store	131.12	-3.38	+22.21
8.	8.	↓ 7.	Microsoft Access	Relational DBMS	126.22	-5.35	-20.27
9.	↑ 10.	9.	SQLite	Relational DBMS	106.78	-0.48	-1.19
10.	↓ 9.	10.	Redis +	Key-value store	104.49	-3.74	+9.00

<http://db-engines.com/en/ranking>

Branches und Forks

- MySQL ist Open Source (GPL)
 - d. h. jede/r darf selber ...
 - → Branches und Forks



- Wer setzt MySQL Branch/Fork ein?
- Babylonische Sprachverwirrung steht uns bevor :-(
 - Sybase ASE vs. MS SQL Server
 - Ingres vs. PostgreSQL
- Bsp: Replication, Protocol X, MariaDB Column Store, Window Functions, ...

MySQL Subskriptionen

- **Nicht Lizenz sondern eine Subskription (Abo) für Support + Tools**
- **MySQL Community Edition (CE)**
 - **Open Source MySQL Server**
 - **+ hilf dir selber**
- **MySQL Standard Edition (SE)**
 - **Open Source oder Commercial MySQL Server**
 - **+ MySQL Support**
- **MySQL Enterprise Edition (EE)**
 - **Standard Edition (SE)**
 - **+ Enterprise Tools (Monitor, Backup, Audit, ...)**
- **MySQL Commercial – OEM, ISV, VAR (Lizenz)**

Installation

Oracle: OUI (Oracle Universal Installer)

- **Windows: MySQL installer**
`C:\Program files\mysql\mysql-server-5.7`
- **MySQL auf Linux:**
 - **Pakete: *.rpm, *.deb**
`/usr`
 - **Binary Tar-Ball:**
`mysql-5.7.12-linux-x86_64.tar.gz`
- **Über Linux-Distro (oft alt)**
- **Vom Hersteller (MySQL/Oracle, ganz frisch!)**

Erstellen der Datenbank

- **Paket der Distribution legt Datenbank an:**

- `/var/lib/mysql`

- `C:\Program files\mysql\mysql-server-5.7\data`

- **Von Hand:**

- ```
mysql_install_db --user=mysql --datadir=...
```

- **Mit MySQL 5.7:**

- ```
mysqld --initialize --datadir=...
```

Konfiguration

- **MySQL Konfigurationsdatei:**
 - `my.cnf` oder `my.ini` (Win)
 - Liegt unter: `/etc/` oder `/etc/mysql`
- **MySQL: kein SPFILE**
- **Änderung in `my.cnf` → Restart oder**
- **`SET GLOBAL variable = wert;`**
- **→ von Hand in `my.cnf` schreiben!**
- **Gutes Template:**
 - <http://fromdual.com/mysql-configuration-file-sample>

Tablespaces (TS)

- **Bis MySQL 5.6: Alles oder Nichts**
 - `innodb_file_per_table = { 0 | 1 }`
 - 0 → alle Tabellen im System TS (`ibdata1`)
 - 1 → jede Tabelle in eigenem TS (`*.ibd`)
- **Ab MySQL 5.7:**
 - **Flexible/general Tablespaces:**

```
CREATE TABLESPACE customer0001
  ADD DATAFILE 'customer0001.ibd' ENGINE = InnoDB;
```

```
ALTER TABLE customer0001.invoices
  TABLESPACE = customer0001;
```

Starten / stoppen von MySQL www.fromdual.com

Oracle: `sqlplus / as sysdba STARTUP`

- MySQL automatisch: in Boot-Prozess eingebunden
- von Hand:

```
service mysql [start | stop]
systemctl start mysql.service
```

- Prüfen:

```
ps -ef | grep mysqld
```

- Killen von Hand (nett, dann hart):

```
kill -15 <pid>; kill -9 <pid>
```

- **Achtung: bei NICHT crash-safen Storage Engines (MyISAM)!**

Prozess vs. Thread

- MySQL: Threads
- Oracle: Prozess (+ neu auch Threads)
 - RECO, PMON, SMON, DBW0, LGWR, ARC0, ...

```
shell> ps -eLf | grep mysqld
UID          PID    PPID    LWP   CMD
mysql       6411     1    6411  mysqld_safe
mysql       7081   6411    7081  mysqld
mysql       7081   6411    7083  mysqld
mysql       7081   6411    7084  mysqld
mysql       7081   6411    7085  mysqld
```

- Welche?

```
SELECT * FROM performance_schema.threads;
```

MySQL Thread Architektur

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

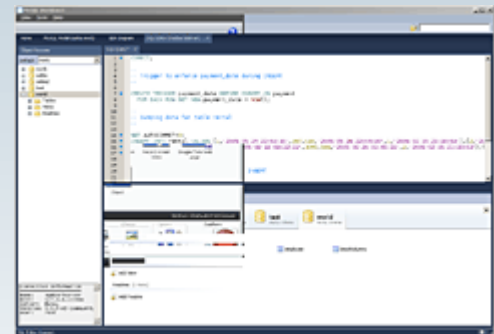
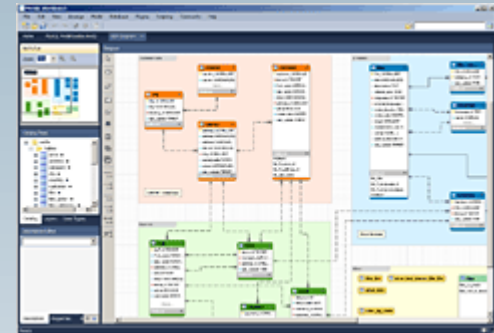
thread_id	thread_name	type	user
1	thread/sql/main	BACKGROUND	NULL
2	thread/innodb/io_ibuf_thread	BACKGROUND	NULL
3	thread/innodb/io_log_thread	BACKGROUND	NULL
4	thread/innodb/io_read_thread	BACKGROUND	NULL
11	thread/innodb/io_write_thread	BACKGROUND	NULL
14	thread/innodb/srv_error_monitor_thread	BACKGROUND	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

Connections / Connectors

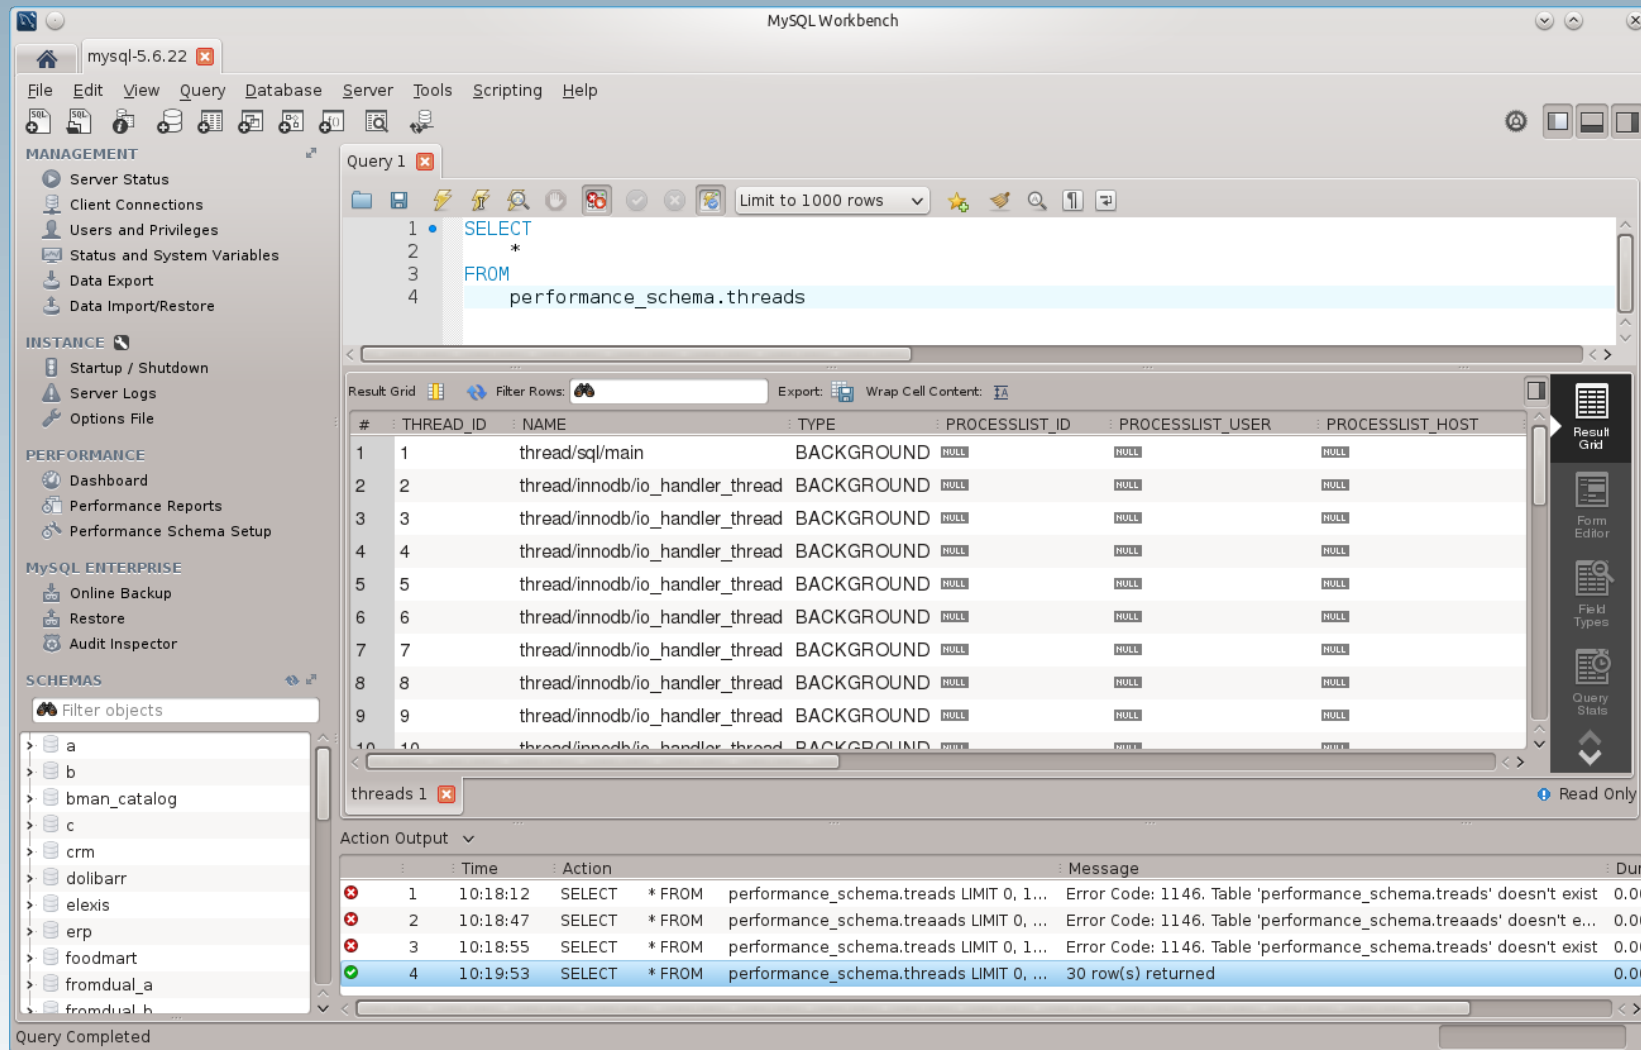
- **Verbindung**
 - In MySQL billig: oft KEIN Connection-Pooling
 - 1 Verbindung = 1 Thread → 1 Query → 1 Core
 - MySQL skaliert heute (5.7) mit der Hardware (144 cores)!
 - Thread Pool (1000e von Verbindungen)
- **MySQL Port: 3306**
- **Connectors:**
 - JDBC/ODBC
 - PHP, Perl, Python, Ruby, .NET

Tools

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



Query Browser



The screenshot shows the MySQL Workbench interface. The left sidebar contains navigation menus for MANAGEMENT, INSTANCE, PERFORMANCE, MySQL ENTERPRISE, and SCHEMAS. The main window displays a query editor with the following SQL code:

```

1 SELECT
2 *
3 FROM
4 performance_schema.threads
  
```

Below the query editor, the 'Result Grid' shows a table of threads. The table has the following columns: #, THREAD_ID, NAME, TYPE, PROCESSLIST_ID, PROCESSLIST_USER, and PROCESSLIST_HOST. The data rows are as follows:

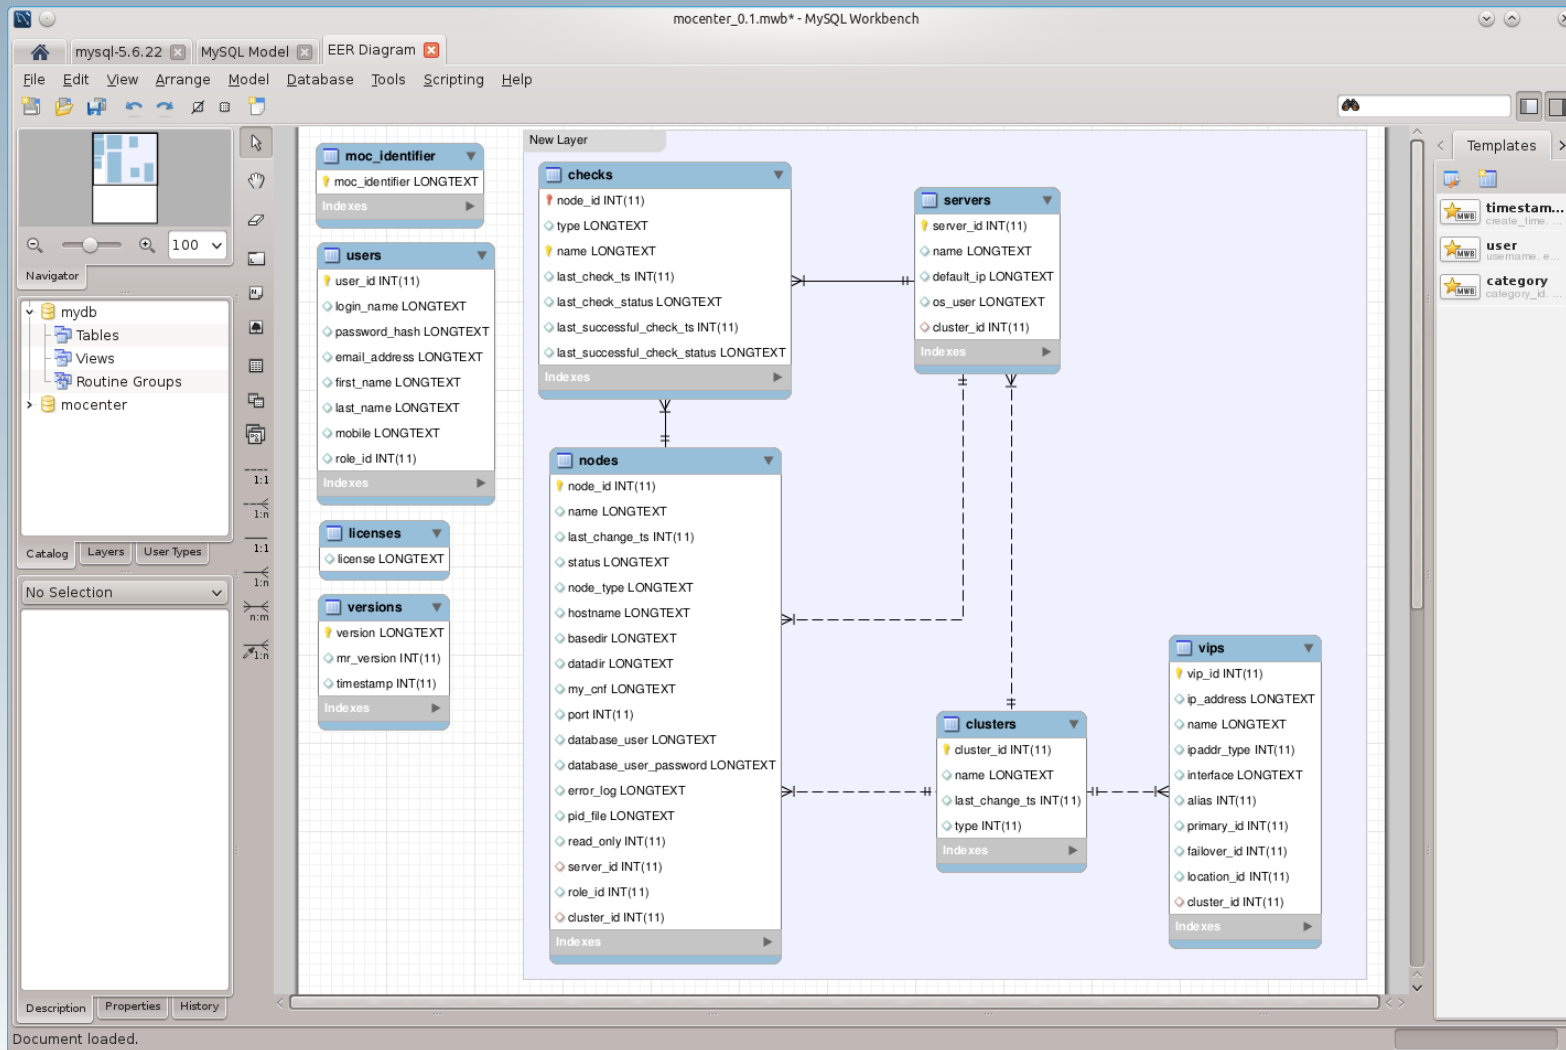
#	THREAD_ID	NAME	TYPE	PROCESSLIST_ID	PROCESSLIST_USER	PROCESSLIST_HOST
1	1	thread/sql/main	BACKGROUND	NULL	NULL	NULL
2	2	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
3	3	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
4	4	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
5	5	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
6	6	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
7	7	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
8	8	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
9	9	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL
10	10	thread/innodb/io_handler_thread	BACKGROUND	NULL	NULL	NULL

Below the result grid, the 'Action Output' pane shows the execution log:

#	Time	Action	Message	Dur
1	10:18:12	SELECT * FROM performance_schema.treads LIMIT 0, 1...	Error Code: 1146. Table 'performance_schema.treads' doesn't exist	0.00
2	10:18:47	SELECT * FROM performance_schema.treaads LIMIT 0, ...	Error Code: 1146. Table 'performance_schema.treaads' doesn't e...	0.00
3	10:18:55	SELECT * FROM performance_schema.treads LIMIT 0, 1...	Error Code: 1146. Table 'performance_schema.treads' doesn't exist	0.00
4	10:19:53	SELECT * FROM performance_schema.threads LIMIT 0, ...	30 row(s) returned	0.00

The 'Query Completed' status is shown at the bottom of the interface.

ER-Diagrammer



User und Schema

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


Wichtigste Befehle

- `SHOW DATABASES ;`
- `USE mysql ;`
- `SHOW TABLES ;`
- `SHOW [FULL] PROCESSLIST ;`
- `SELECT user, host FROM mysql.user ;`
- `SHOW GRANTS FOR user@host ;`
- `SHOW GLOBAL VARIABLES ;`
- `SHOW GLOBAL STATUS ;`
- `SHOW ENGINES ;`



Storage Engines

- MySQL Architektur ist speziell:



- Verschiedene Storage Engines (DB Engine):
InnoDB, MyISAM, MEMORY

```
SELECT * FROM information_schema.tables;
```

- Früher MyISAM (≤ 5.1)
- heute: InnoDB (≥ 5.5)

```
SHOW CREATE TABLE test\G
CREATE TABLE `test` (
  ...
) ENGINE=InnoDB;
```

Backup / Restore

- Logisch vs. physisch
- Logisches Backup mit `mysqldump` (`exp/imp`)
- Physisches Backup mit MySQL Enterprise Backup (MEB) (`rman`)
- Backup:
 - logisch: `mysqldump > full_dump.sql`
 - physisch: `mysqlbackup backup-and-apply-log`
- Restore:
 - logisch: `mysql < full_dump.sql`
 - physisch: `mysqlbackup copy-back`

Point-in-Time-Recovery

- **Binary-Log für Point-in-Time-Recovery**
- **Binary-Log einschalten**
 - `log_bin = binary-log`
 - erfordert DB Neustart
 - ~ Oracle Archive Log (aber nur ähnlich!)
- **Bei Backup: Binary-Log und Position merken**
- **PiTR:**
 - `mysqlbinlog --start-position=45678 --stop-datetime=... binary-log.000042 ...`

Logging

- **Error Log**
 - `log_error = dateiname.log`
 - **Neustart erforderlich !**
 - **Fehler und Warnungen (`log_warnings = 2`)**
- **Slow Query Log**
 - **„langsame“ Abfragen**
- **General Query Log**
 - **Achtung: „alle“ Abfragen**
 - `general_query_log = ON`

MySQL Upgrade

- **2 Varianten:**
 - **Dump/Restore (logical upgrade) (5 TB DWH?)**
 - **Binary-Upgrade (in-place upgrade) (seit 5.5 supportet)**
- **Vorgehen (ca. 15 min):**
 - `SET GLOBAL innodb_fast_shutdown = 0;`
 - **DB stoppen**
 - **alte Pakete deinstallieren**
 - **neue Pakete installieren**
 - **DB starten**
 - `mysql_upgrade`
 - **DB nochmals stoppen und starten**
- **Major Releases NICHT überspringen (5.5 -X-> 5.7)**

Performance Tuning

- `mysql> SHOW GLOBAL VARIABLES ;`
- `mysql> SHOW GLOBAL STATUS ;`
- `mysql> SHOW ENGINE INNODB STATUS\G`
- `PERFORMANCE_SCHEMA` (seit 5.6)
- **Slow Query Log**
- **Query Execution Plan:**
`mysql> EXPLAIN SELECT * FROM test;`

Slow Query Log

- Systematischer Ansatz mit etwas Vorlauf:

```

+-----+-----+
| Variable_name          | Value      |
+-----+-----+
| slow_query_log         | OFF        |
| slow_query_log_file    | slow.log   |
| log_queries_not_using_indexes | OFF        |
| long_query_time        | 0.500000  |
+-----+-----+

```

- Kann dynamisch eingeschaltet werden:

- `SET GLOBAL slow_query_log = 1;`

- Profile vom Slow Query Log:

- `shell> mysqldump_slow -s t slow.log > slow.profile`

- `shell> pt_query_digest slow.log > slow.digest`

Optimiere das Query!

- Was machen mit den langsamen Abfragen?
- Query Execution Pläne (QEP) erstellen!
 - `mysql> EXPLAIN SELECT ...`
- Interpretieren von QEP:

```
EXPLAIN SELECT * FROM emp where name = 'Oli';
```

select_type	table	type	possible_keys	key	key_len	ref	rows
SIMPLE	emp	ALL	last	NULL	NULL	NULL	261369

Operation

Genutzter
Index

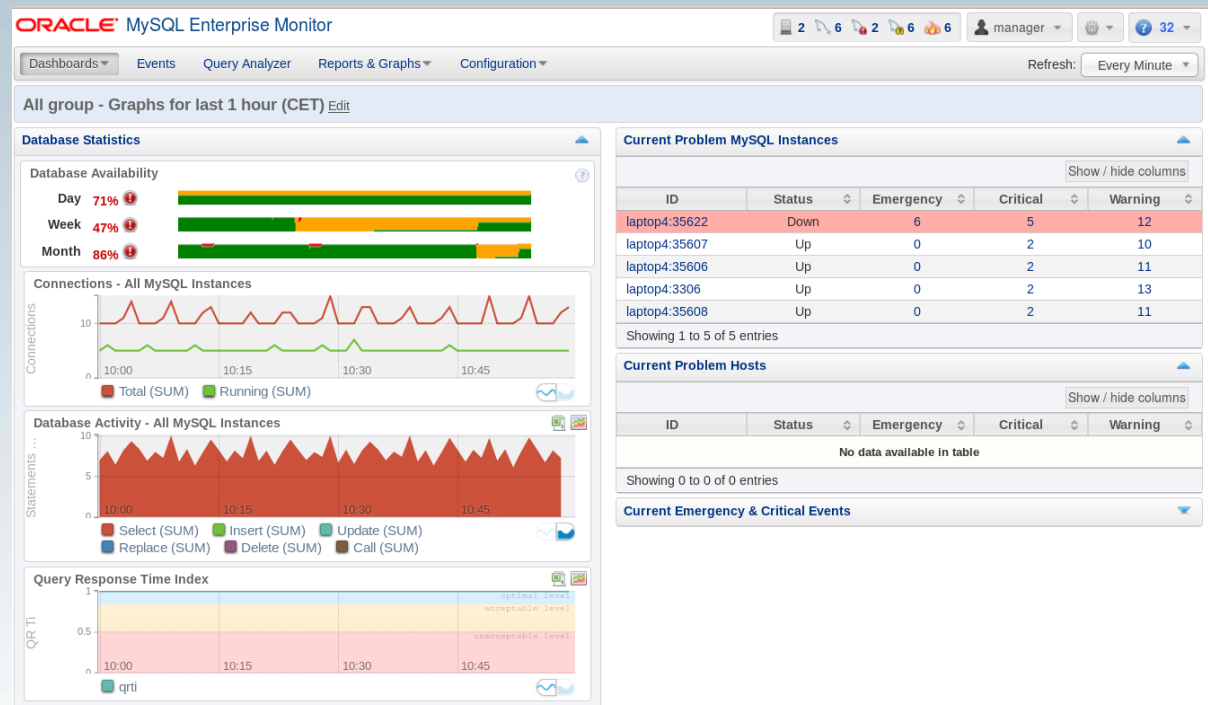
Angelangte
Zeilen

EXPLAIN Type Operationen

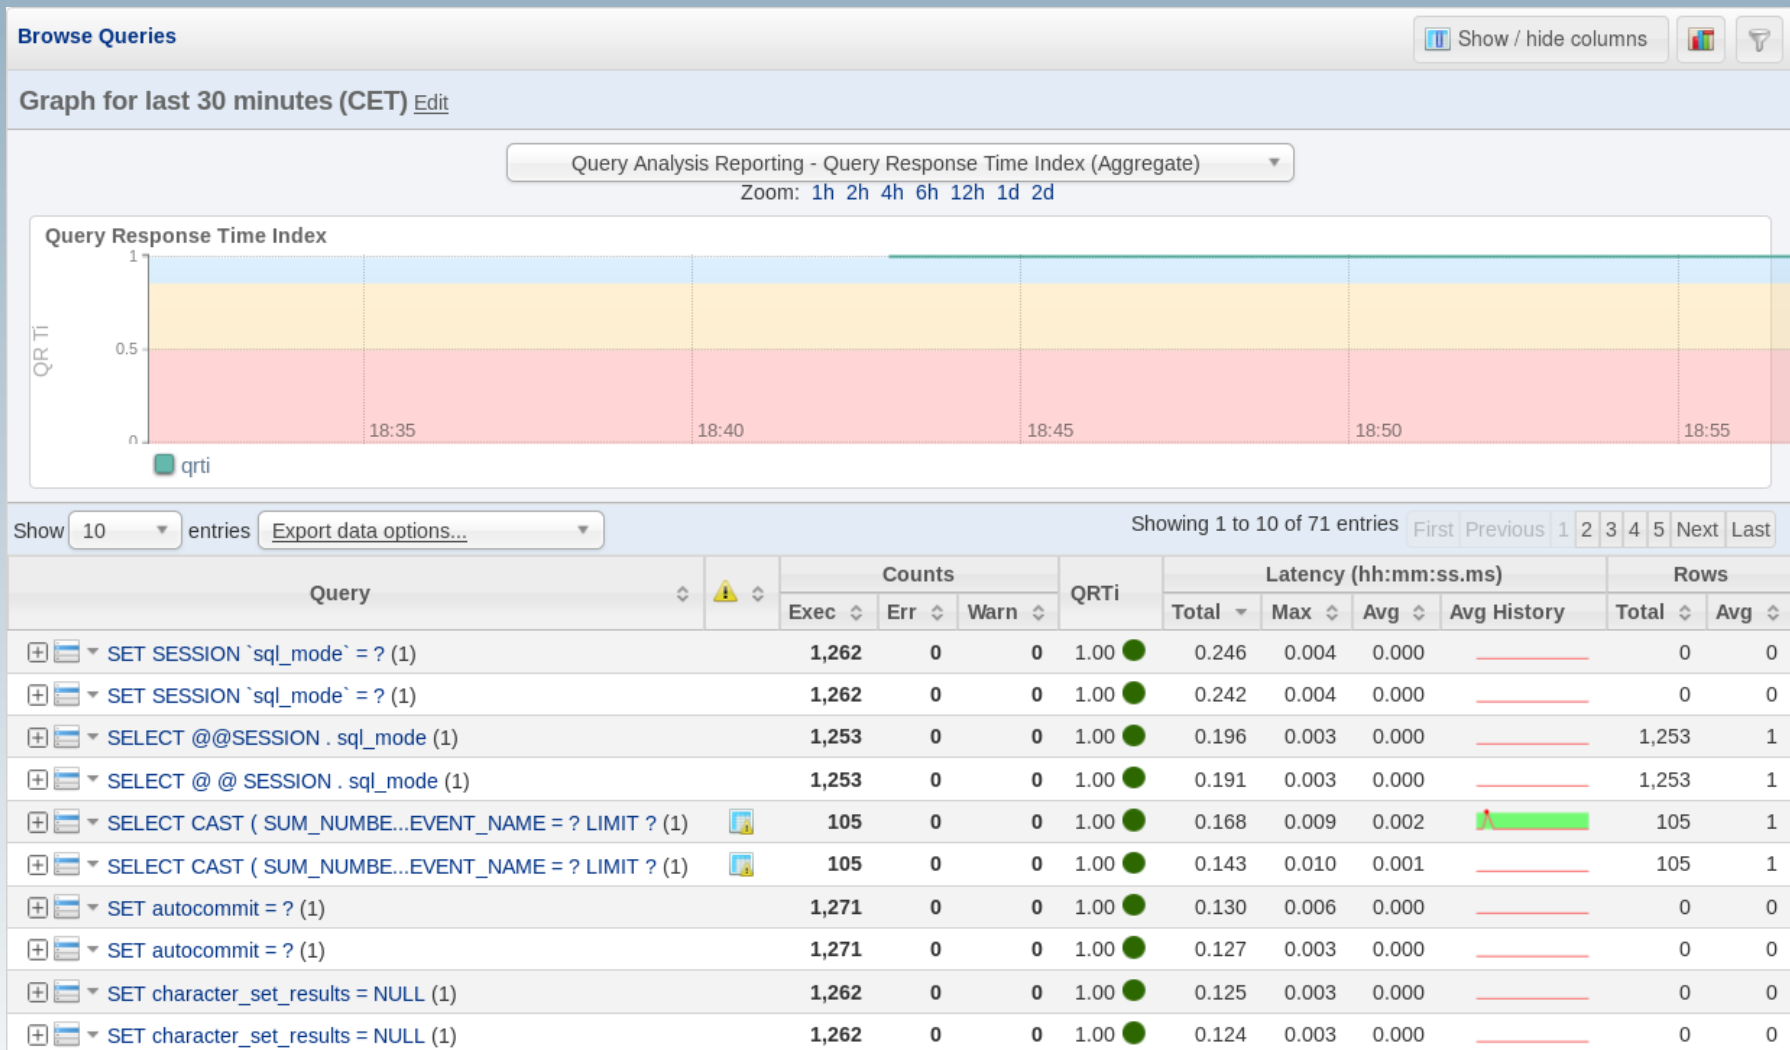
billig ↑ ↓ teuer	const	Höchstens eine passende Zeile, wird wie Konstante behandelt
	eq_ref	Ein Zeile pro Zeile aus vorheriger Tabelle (Primary Key- / Unique Key-Join)
	ref	Mehrere Zeilen pro Zeile aus vorheriger Tabelle (non-Unique Key-Join, Datenmenge nimmt zu!)
	fulltext	Der Join wird mittels FULLTEXT Index gelöst
	index_merge	Mehrere Index-Suchen werden gemerged
	xxx_subquery	Subqueries
	range	Index Range Scan
	index	Full Index Scan (IFFS)
	ALL	Full Table Scan

Monitoring von MySQL

- MySQL Enterprise Monitor (MEM)
- Oracle Cloud Contol (OEM/OMS) mit MySQL Plug-in
- 3rd Party Tools




Graphisch: Query Analyzer



Hochverfügbarkeit

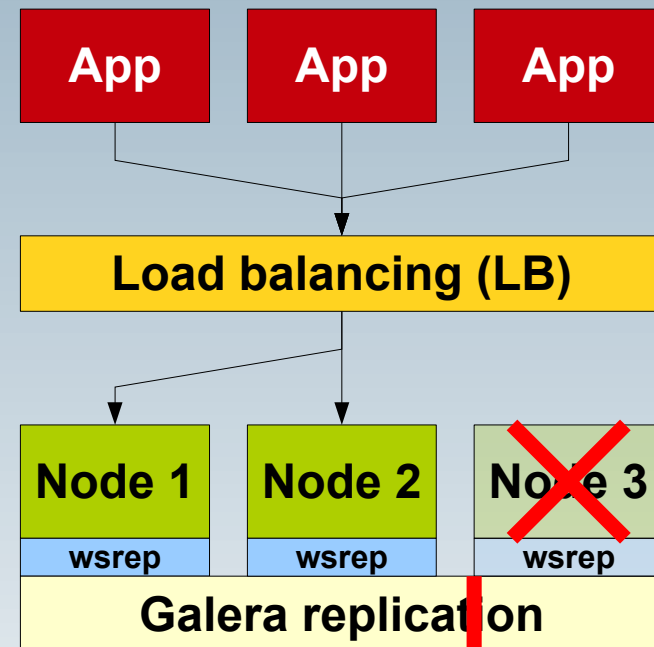
- **aktiv/passiv Failover-Cluster**
 - dito
- **Master/Slave Replikation**
 - Data-Guard
- **Galera Cluster für MySQL**
 - Oracle RAC

MySQL Replication



RAC: Galera Cluster

- Oracle Real Application Cluster (RAC)
- Hardware-Ausfall
- Wartungsarbeiten
 - HW/OS/DB Upgrade
- 5x9 HA: 99.999%



Q & A



Fragen ?

Diskussion?

Wir haben Zeit für ein persönliches Gespräch...

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

www.fromdual.com/presentations