MySQL Konfiguration - die wichtigsten Parameter

DOAG SIG MySQL – Performance
13. März 2012, Wiesbaden

Oli Sennhauser

Senior MySQL Consultant, FromDual GmbH

oli.sennhauser@fromdual.com

FromDual GmbH

- **FromDual bietet neutral und unabhängig:**
  - **Beratung für MySQL**
  - **Support für MySQL und Galera Cluster**
  - **Remote-DBA / MySQL Betrieb**
  - **Schulung für MySQL**
- **Oracle Silber Partner (OPN)**

www.fromdual.com

Kunden


HolidayCheck.com
linuxhotel
Schulungen und Universitäten machen Sinn.
Hostloco
mespas
OX
open-xchange
RYGOL
DÄMMSTOFFE
Boerse Stuttgart
sedo
TEAM NETSOL
careerJET.com
iNFOnline
verified by
swiss sign
GOLD
STRATO
fashion days
EXCLUSIVE SHOPPING
Flat
Planet
Phone
Company
billiger-mietwagen.de
ERICSSON
BOSCH
Invented for life
CoreTelecom
INNOVATIONS INC.
The Future of Telephony
nexellent
communications.
delivered.
Diogenes
integrics
Integrated Telecommunication Solutions

Inhalt

MySQL Konfiguration

- **FromDual Perfromance Waage**
- **The Big 9!**
- **InnoDB**
- **MyISAM**
- **MySQL**
- **Weitere Parameter**

FromDual Performance Waage

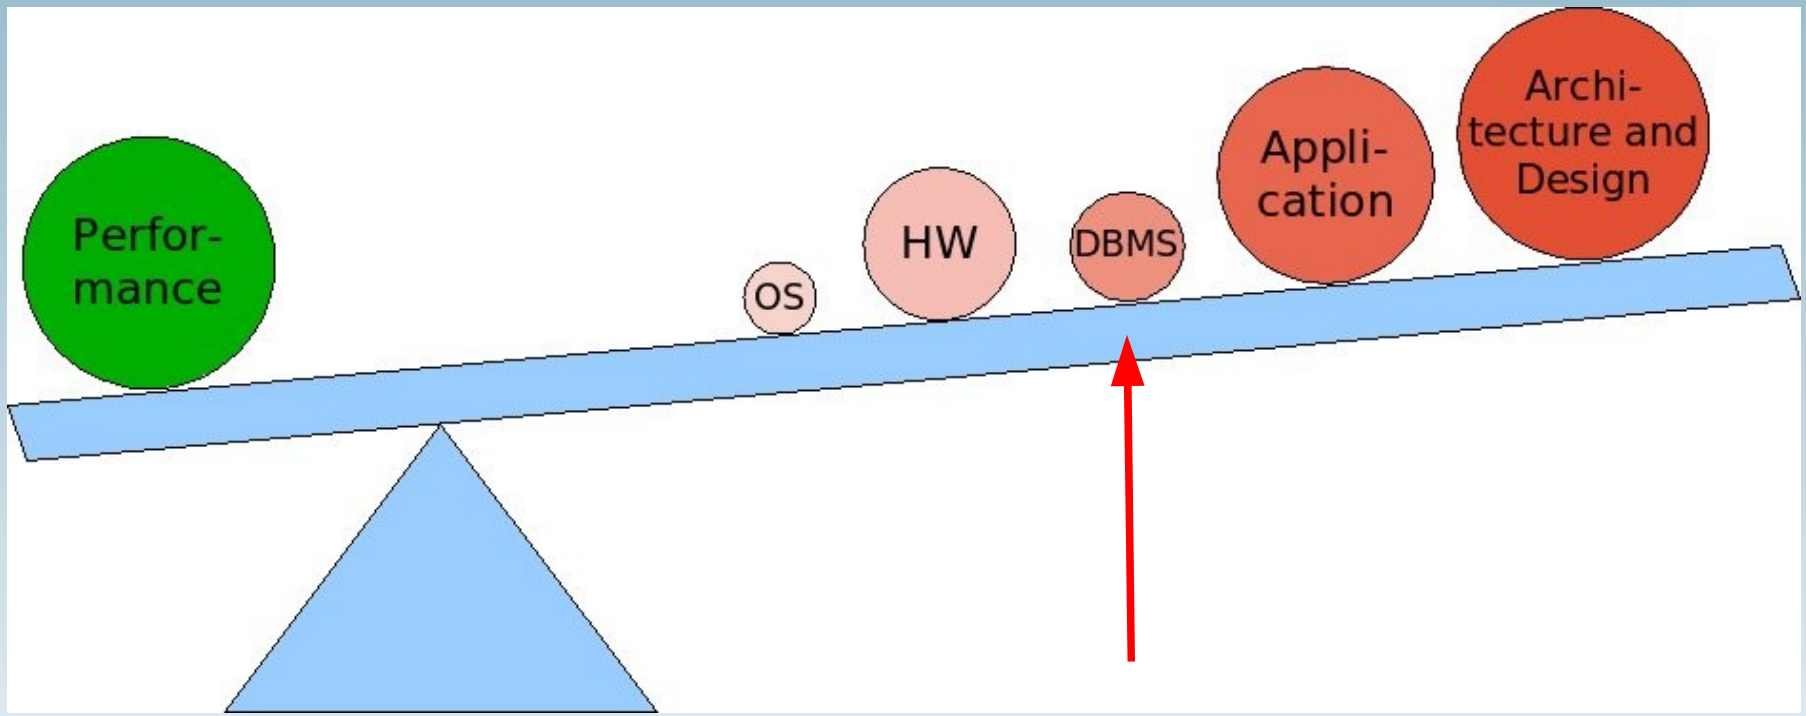

MySQL Tuning

- Welche Storage Engine verwendet Ihr zur Zeit?
- Welchen MySQL Release? (→ 5.1 und neuer)
- Zur Zeit: ca. 330 MySQL Parameter
  - → aber nur ca. 8 (9) davon sind signifikant!
  - → Grob-Tuning
- Alle anderen nur nach ausführlichem Benchmarken
  - → Fine-Tuning

MySQL Architektur

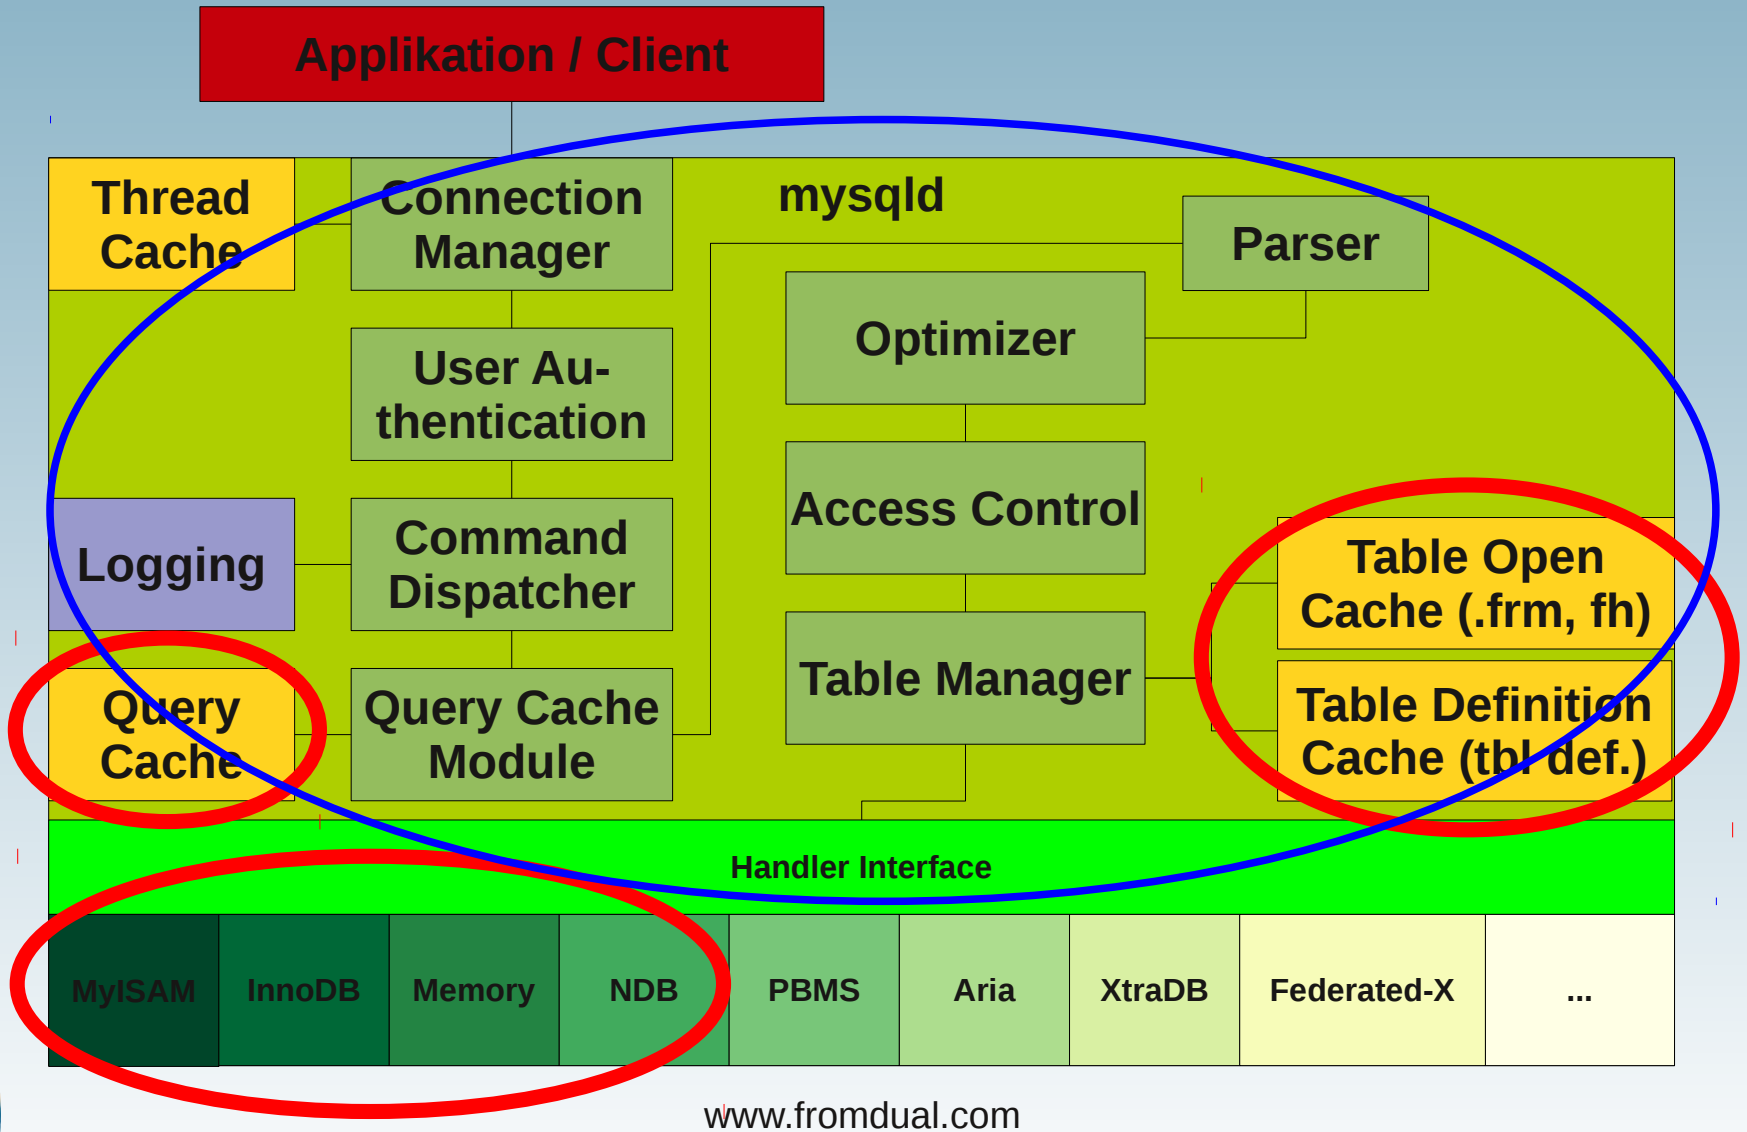

Bevor wir anfangen!

- **Kunde fragt mich: Kann mein System 30% mehr Last vertragen?**
- **Chemische Verfahrenstechnik:**

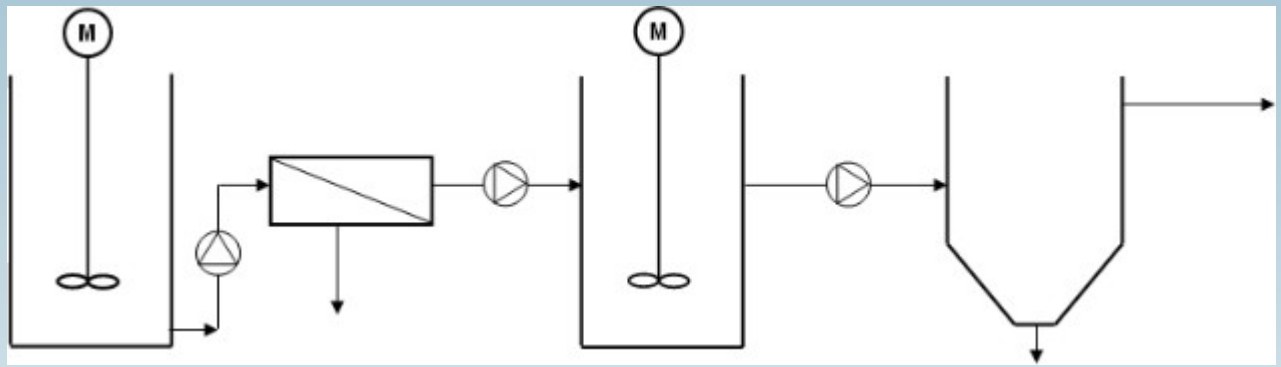

- **Gibt es Unterschiede zu eine DB basierten System?**
- **Was brauche ich um diese Frage zu beantworten?**

Messen, messen, messen...

- **Messen!**
- **Idealfall: Nur ein Parameter aufs mal ändern!**

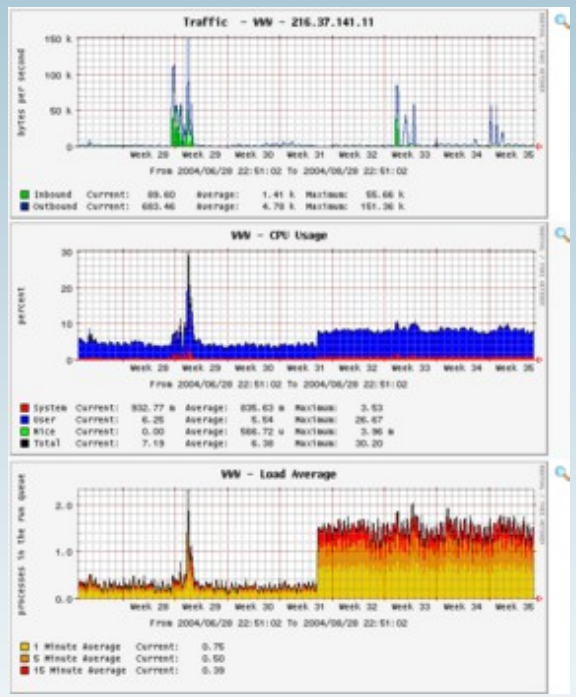

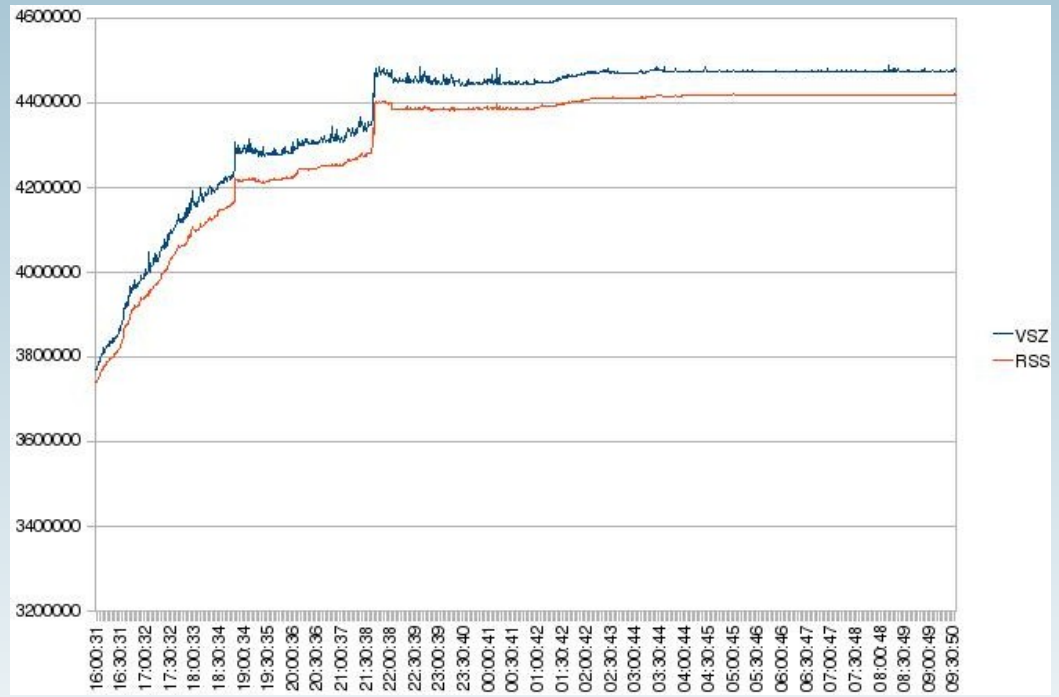

The big 9!

- InnoDB: 5
  - InnoDB Buffer Pool: 2
  - InnoDB Log File: 3
- MyISAM: 1
  - Key Buffer: 1
- MySQL: 3
  - Query Cache: 1
  - Table_*_cache: 2

InnoDB Buffer Pool

- InnoDB cached: Daten UND Indizes
- Grösse: `innodb_buffer_pool_size`
  - in Byte (M, G)
  - Pages à 16k (5.6: 4k, 8k, 16k)
- Repräsentation der Platte im Speicher:
- Ca. 80% vom RAM auf dedizierter InnoDB Maschine

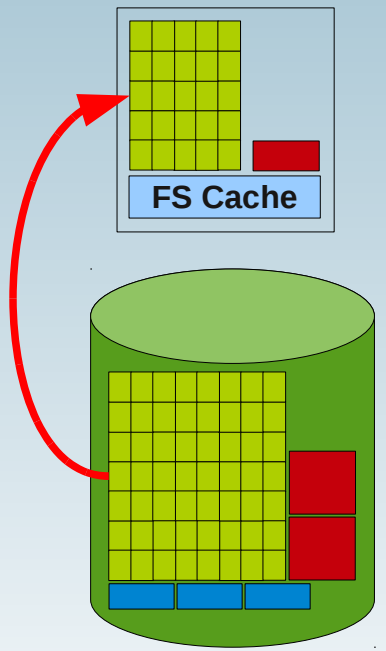

InnoDB Buffer Pool Informationen

- **Messen:**

```
SHOW GLOBAL STATUS LIKE 'innodb_buffer_pool_pages%';

+---------------------------------+----------+
| Variable_name                   | Value    |
+---------------------------------+----------+
| Innodb_buffer_pool_pages_data   | 9175     |
| Innodb_buffer_pool_pages_misc   | 40       | +
| Innodb_buffer_pool_pages_free   | 1024     | +
| Innodb_buffer_pool_pages_total  | 10239    | =

| Innodb_buffer_pool_pages_dirty  | 289      | 3.1%
| Innodb_buffer_pool_pages_data   | 9175     |
+---------------------------------+----------+
```

InnoDB Buffer Pool

- **Buffer Pool Hit Ratio:**

  read_requests / (read_requests + reads) * 100
  = 99.9%

```
SHOW GLOBAL STATUS LIKE 'innodb_buffer_pool%';

+---------------------------------------+-------------+
| Variable_name                         | Value       |
+---------------------------------------+-------------+
| Innodb_buffer_pool_read_requests      | 1507235721  |
| Innodb_buffer_pool_reads              | 1193121     |

| Innodb_buffer_pool_wait_free          | 5           |
+---------------------------------------+-------------+
```

InnoDB Buffer Pool von InnoDB Status

```
SHOW ENGINE INNODB STATUS\G
----------------------
BUFFER POOL AND MEMORY
----------------------
Total memory allocated 171704320
Dictionary memory allocated 2473598
Buffer pool size   10239
Free buffers       1024
Database pages     9185
Old database pages 3370
Modified db pages  812
Pending reads 0
Pending writes: LRU 0, flush list 0 single page 0
Pages made young 1768431, not young 0
1.00 youngs/s, 0.00 non-youngs/s
Pages read 1197328, created 215334, written 22307309
1.00 reads/s, 1.00 creates/s, 0.00 writes/s
Buffer pool hit rate 1000 / 1000, young-making rate 0 / 1000 not 0 / 1000
Pages read ahead 0.00/s, evicted without access 0.00/s,
Random read ahead 0.00/s
LRU len: 9185, unzip_LRU len: 0
I/O sum[3125]:cur[2], unzip sum[0]:cur[0]
```

InnoDB Buffer Pool Monitoring

- **Es ist wünschenswert, aber meist illusorisch, alle Daten im RAM zu halten!**

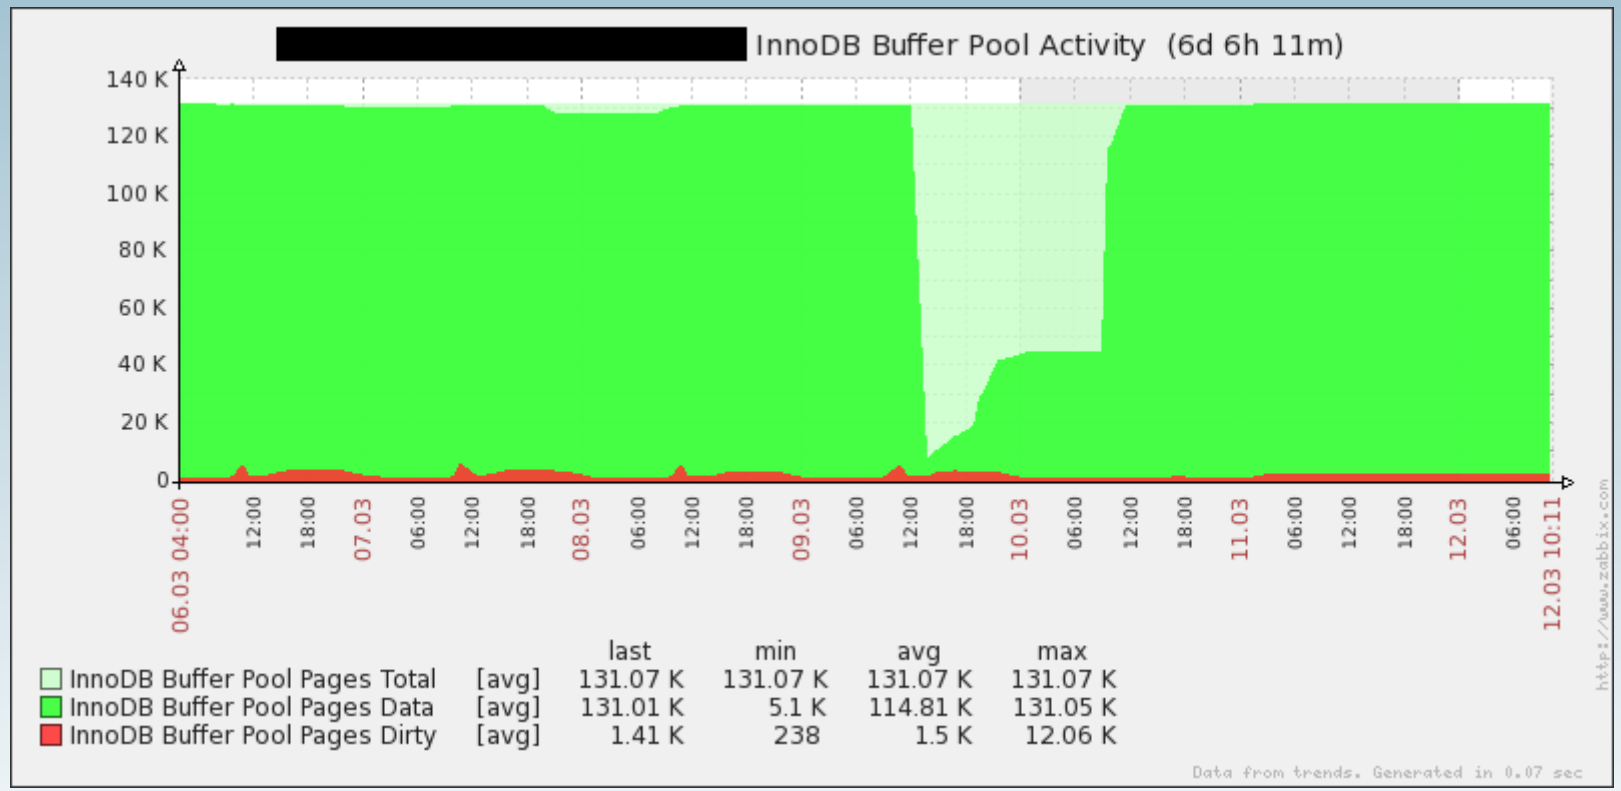

InnoDB Buffer Pool Monitoring

- **Einfluss von Partitionierung auf den InnoDB Buffer Pool:**

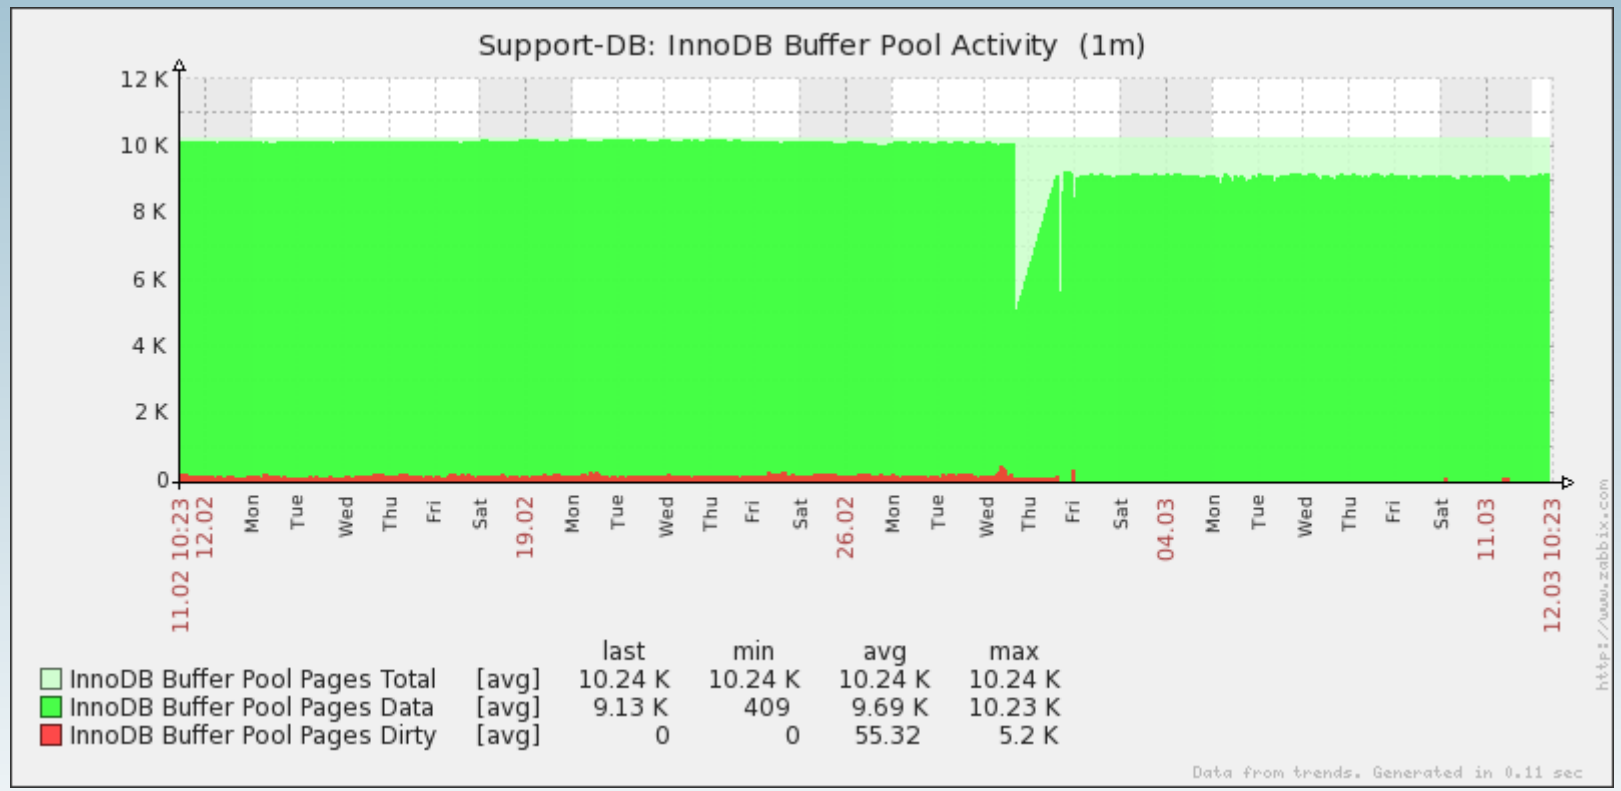

InnoDB Buffer Pool Instanzen

- **InnoDB Buffer Pool Instanzen!**
  - **Neu mit MySQL 5.5**
  - **Bei sehr viel RAM**
  - **Bei hoher Concurrency**
    - **→ Locks auf Buffer Pool Verwaltungsstruktur**
    - **→ Eigene Free List, Flust List, LRU List, Buffer Pool Mutex, etc.**
  - **Hash-Verteilung (random?)**
- **`innodb_buffer_pool_instances`**
  - **Sinnvoll ab 2 Gb**
  - **min. 1 Gb pro Instanz**
  - **max. #cores**

InnoDB Log File

- **innodb_flush_log_at_trx_commit**
  - 0, 2 für Performance, 1 für Sicherheit
    - 0: 1/s + fsync
    - 1: COMMIT + fsync
    - 2: COMMIT + 1/s fsync
- **innodb_log_file_size**
  - Grösser = schneller, aber längere Recovery Zeiten → 2 x 256 M
- **sync_binlog**
  - != 0 → langsam(er)

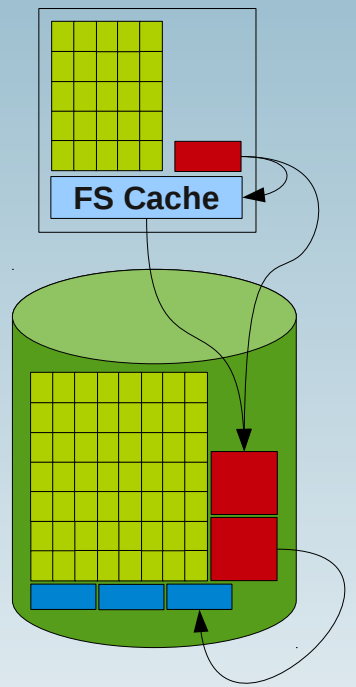

InnoDB Log File Information

```
SHOW ENGINE INNODB STATUS\G
---
LOG
---

Log sequence number 404010398185 -
Log flushed up to   404010198753 = 199432 byte log buffer

Log flushed up to   404010198753 -
Last checkpoint at  404010110312 =  88441 byte log file

0 pending log writes, 0 pending chkp writes
152844 log i/o's done, 0.50 log i/o's/second

SHOW GLOBAL STATUS LIKE 'innodb_os_log_%';
+-------------------------------+-------------+
| Variable_name                 | Value       |
+-------------------------------+-------------+
| Innodb_os_log_fsyncs          | 153643      |
| Innodb_os_log_pending_fsyncs  | 0           |
| Innodb_os_log_pending_writes  | 0           |
+-------------------------------+-------------+
```

InnoDB Log File Monitoring

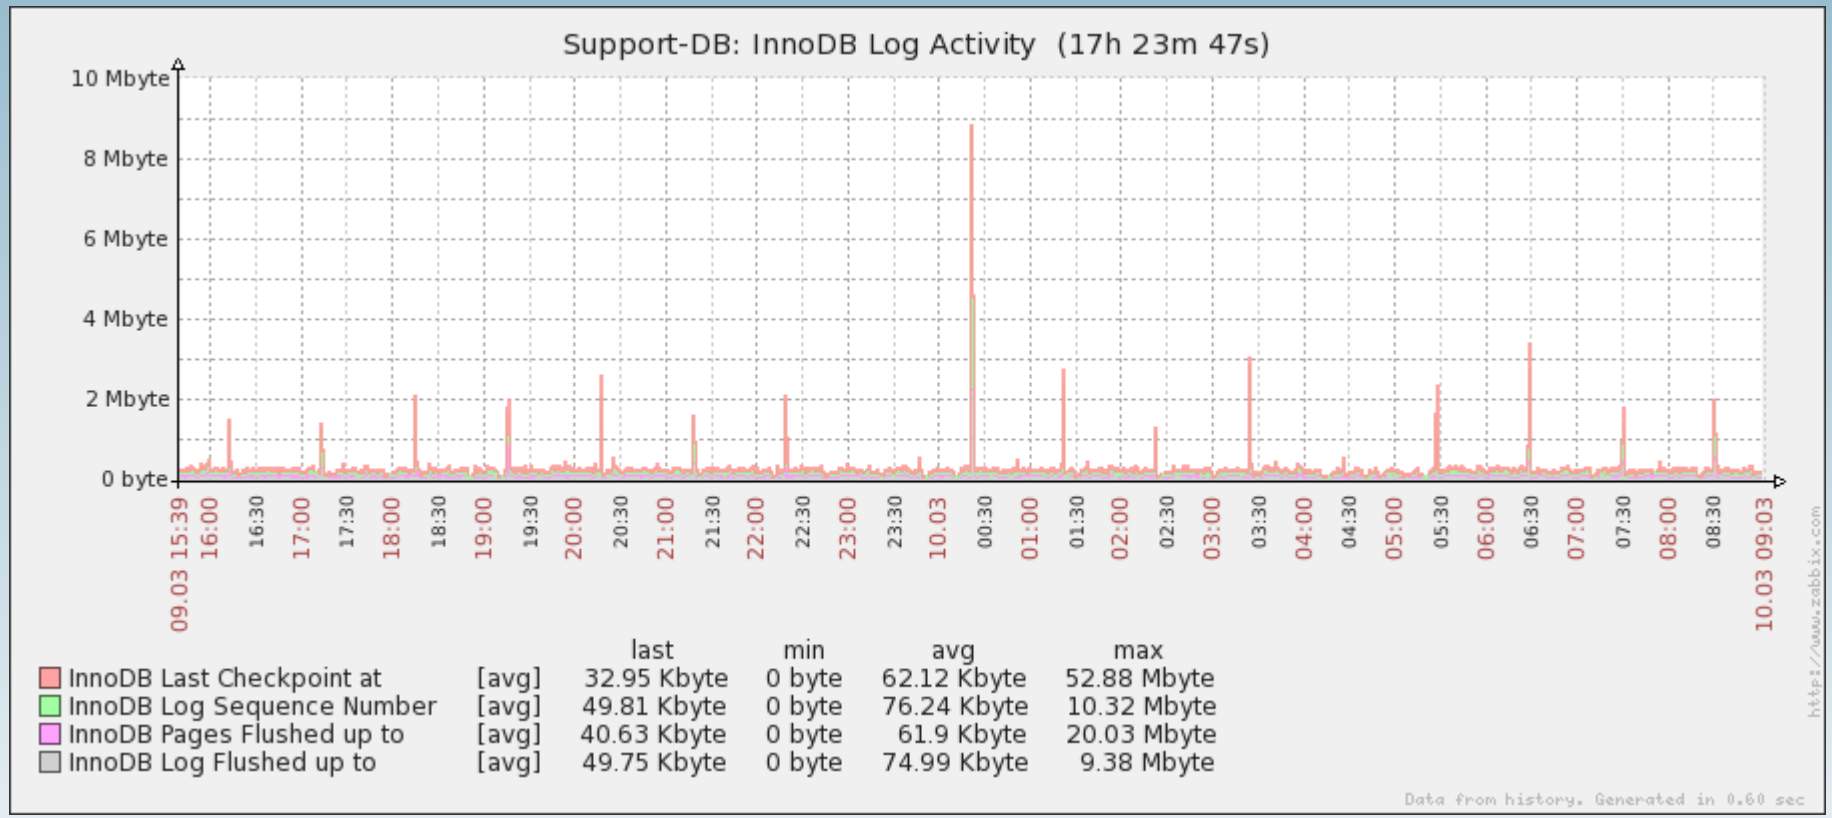

MyISAM Key Buffer

- MyISAM cached nur Indizes!
  - → Key Buffer
  - → Daten: File System Cache
- `key_buffer_size`
  - ca. 25 – 33% vom RAM auf dedizierter Maschine
  - ca. 67 – 75% vom RAM für File System Cache
- Üblicherweise 1k pages

MyISAM Key Buffer Information

- **Key Buffer Hit Ratio:**

```
read_requests / (read_requests + reads) x 100
```

- **Key_blocks_used: High Water Mark!**

```
SHOW GLOBAL STATUS LIKE 'key%';
+------------------------+--------+
| Variable_name          | Value  |
+------------------------+--------+
| Key_blocks_not_flushed | 0      |
| Key_blocks_unused      | 1674   |
| Key_blocks_used        | 9      |
| Key_read_requests      | 130434 |
| Key_reads              | 0      |
+------------------------+--------+

shell> free
             total       used       free     shared    buffers     cached
Mem:      16431960    6190408   10241552          0     408832    1719944
```

MyISAM Monitoring

- **Key Buffer mehr als genügend gross!**

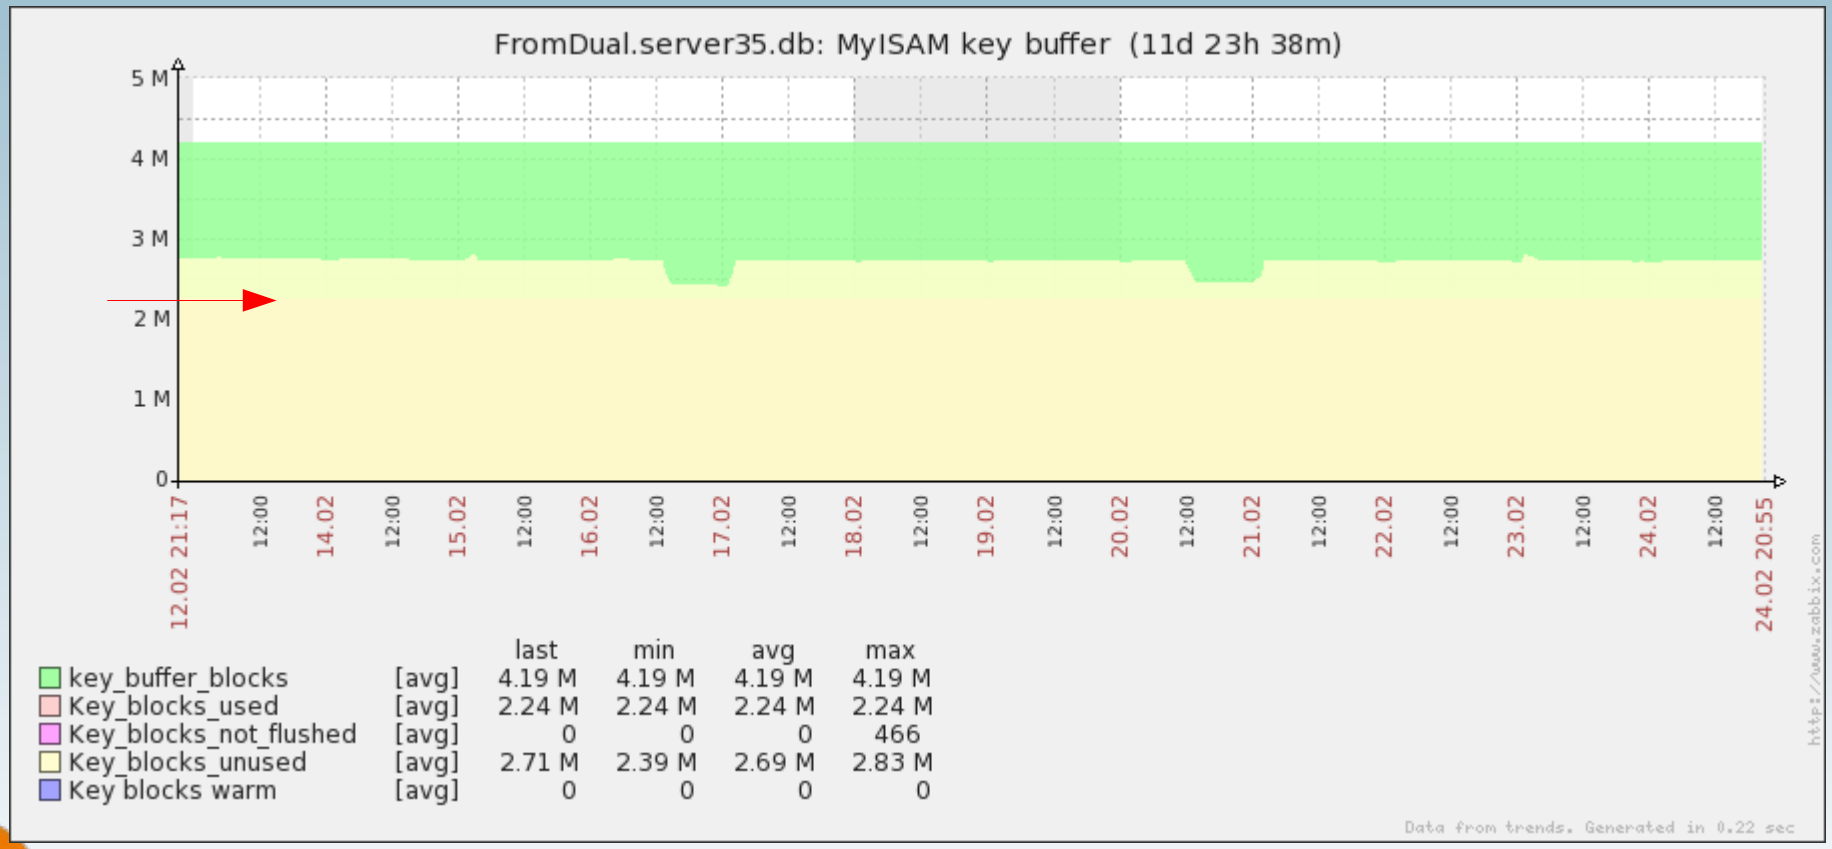

Query Cache

- Cached `SELECT` Queries
  - Hash, Tabellen, Resultat
- Performance Booster bei:
  - Lesen >> Schreiben
  - Geringer Concurrency
- `query_cache_size / query_cache_type`
  - Nicht zu gross machen (<= 128 M)
- Bei sehr hoher Concurrency schädlich!
  - Global Query Cache Mutex

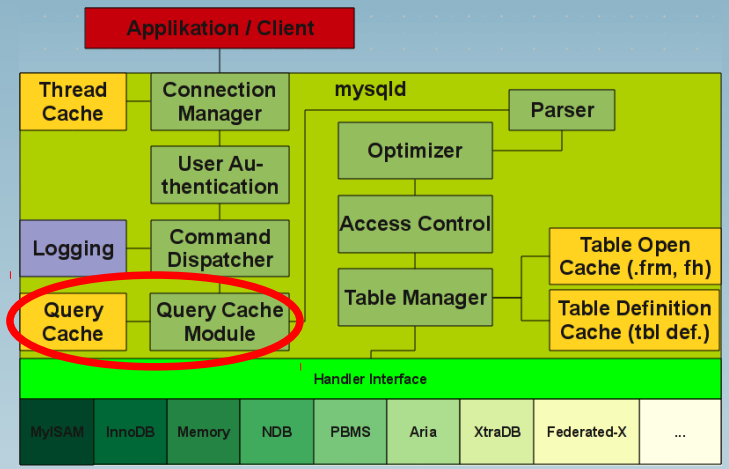

Query Cache Informationen

- Schlechtes Hits / Inserts Verhältnis:
  - Ist: 0.26, Soll: >= 1.2
  - Insert: 10 – 20 % Overhead
- Query Cache Hit Ratio: 16.9%
  - `Qcache_hits / (Com_select + Qcache_hits) x 100`

```
SHOW GLOBAL STATUS LIKE 'qcache%';
+-------------------------+---------+
| Variable_name           | Value   |
+-------------------------+---------+
| Com_select              | 5049720 |
| Qcache_free_memory      | 1995376 |
| Qcache_hits             | 1026597 |
| Qcache_inserts          | 3911380 |
| Qcache_lowmem_prunes    | 2529208 |
| Qcache_not_cached       | 1138305 |
| Qcache_queries_in_cache | 1357    |
+-------------------------+---------+
```

Query Cache Monitoring

- **Um xx:25 macht „jemand“ den Query Cache „kaputt“:**

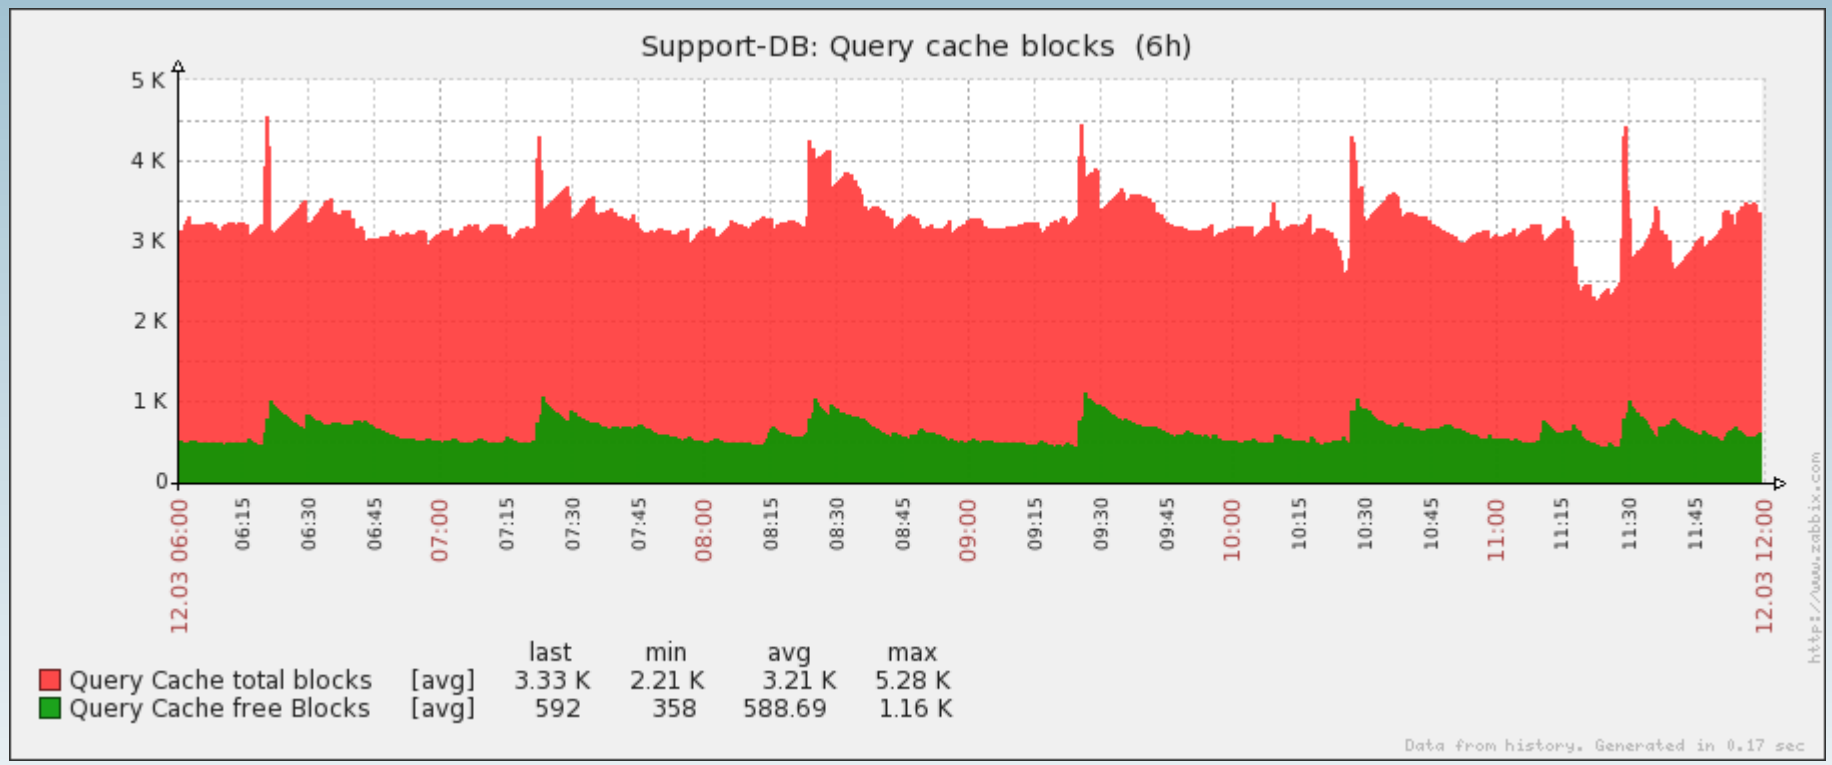

MySQL Parameter

- `table_open_cache`
  - Cached File Handles
  - Laufende Connections x benutze Tabellen
    → 2 – 4k ist nicht ungewöhnlich!
  - Wenn zu klein → massive Performance-Probleme!
  - Hat Einfluss auf open-files-limit!
- `table_definition_cache`
  - Cached `.frm` File (Tabellenstruktur)
  - 512 – 4096 ist nicht ungewöhnlich!

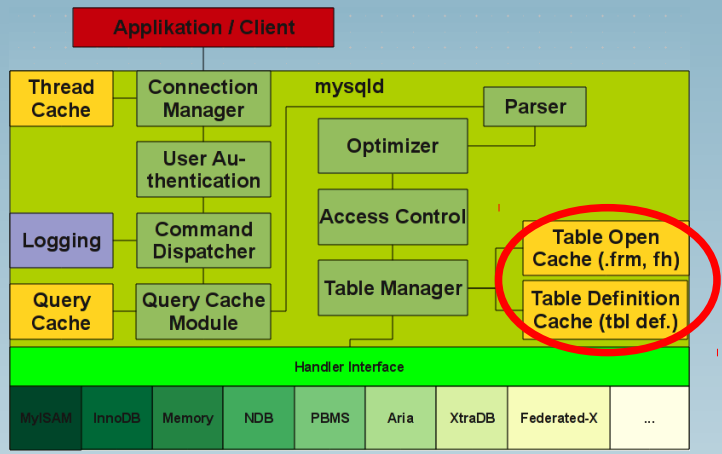

MySQL Parameter Informationen

```
SHOW GLOBAL STATUS LIKE 'open%table%s';
+--------------------------+-------+
| Variable_name            | Value |
+--------------------------+-------+
| Open_table_definitions   | 100   |
| Open_tables              | 102   |
| Opened_table_definitions | 100   |
| Opened_tables            | 109   |
+--------------------------+-------+

SHOW GLOBAL STATUS LIKE 'open%table%s';
+--------------------------+----------+
| Variable_name            | Value    |
+--------------------------+----------+
| Open_table_definitions   | 67321    |
| Open_tables              | 131023   |
| Opened_table_definitions | 6177337  |
| Opened_tables            | 12489056 |
+--------------------------+----------+

SHOW OPEN TABLES;
+--------------------+--------------------------+--------+-------------+
| Database           | Table                    | In_use | Name_locked |
+--------------------+--------------------------+--------+-------------+
| crm                | saved_search             |      0 |           0 |
| zabbix             | users                    |      0 |           0 |
...
```

MySQL Parameter Monitoring

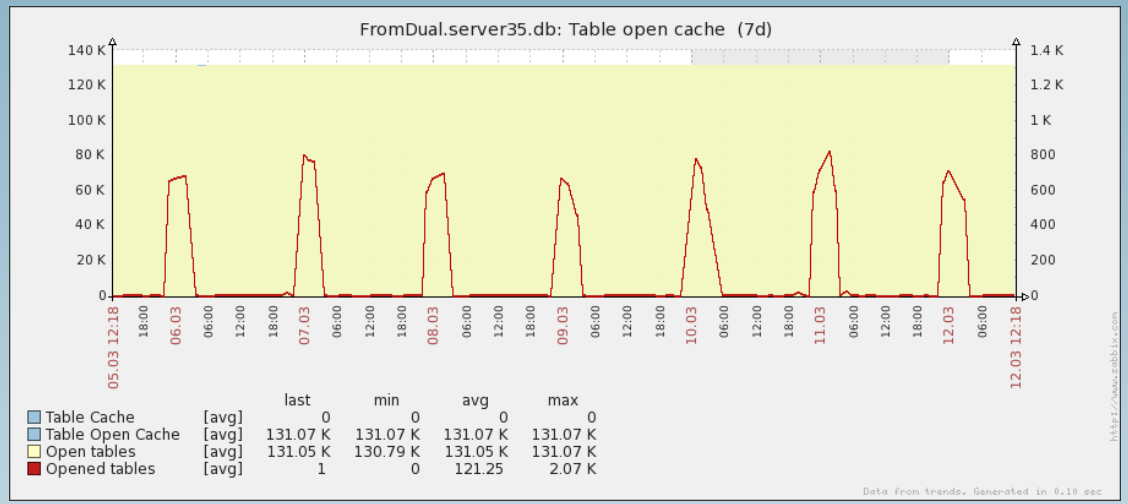

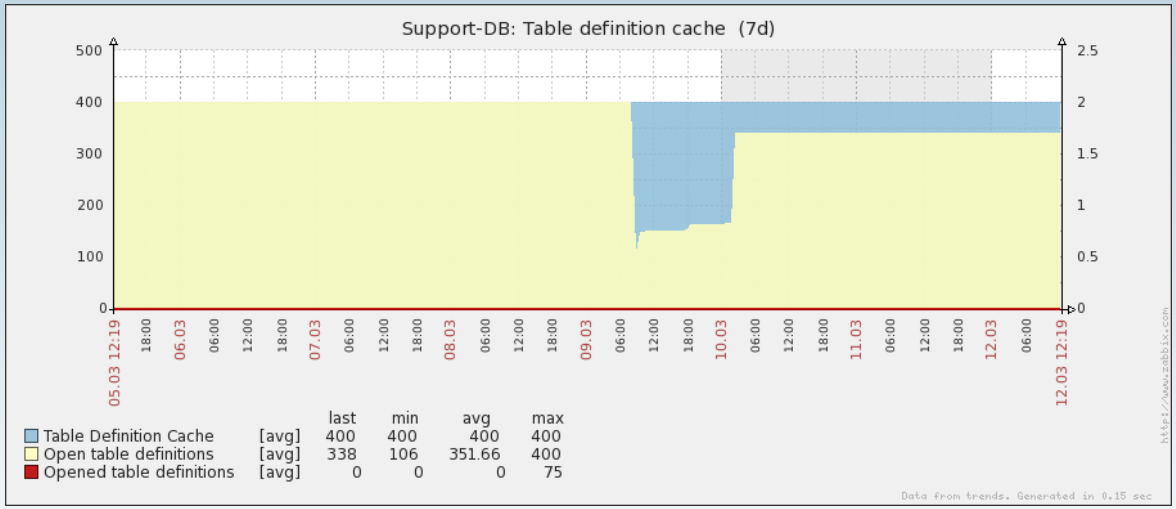

Weitere InnoDB Parameter

- Viele Neuerungen mit MySQL 5.5 und 5.6!
- InnoDB Log File / Buffer:
- `innodb_log_group_home_dir = ./`
  - Default datadir
  - Trennen von sequentiellem I/O und random I/O
- `innodb_log_buffer_size = 8M`
  - Für grosse Transaktionen
  - `Innodb_log_waits > 0`

Weitere InnoDB Parameter

- `innodb_file_per_table = 1`
  - 1 grosses Sytem TS File vs.
  - 1 kleines System TS File + 1 File pro Tabelle
  - Vorteil: Diskplatz wird wieder freigegeben
- `innodb_data_home_dir`
  - Default datadir
  - Splitten von random I/O auf Daten und sequential I/O auf Log Files
- `innodb_data_file_path = ibdata1:10M:autoextend`
- `innodb_autoextend_increment = 8`
  - Achtung: 8 Mbyte!!!
  - Gilt nur für System TS File
- `innodb_file_format = Antelope`
  - Default (= alt)
  - Barracuda ermöglicht File compression

Weitere InnoDB Parameter

- **innodb_flush_method =**
  - **Ausprobieren!**
  -

	data files		log files	
	open	**flush**	**open**	**flush**
default	normal	fsync	normal	fsync
O_DSYNC	O_SYNC	fsync	O_SYNC	O_SYNC
O_DIRECT	O_DIRECT	fsync	normal	fsync

- **innodb_io_capacity = 200**
  - **Raid-1/10 #Disks / 2 x 200**
- **innodb_read_io_threads = 4**
- **innodb_write_io_threads = 4**
  - **Pending reads > innodb_read_io_threads x 64**
  - **Pending writes**
- **innodb_use_native_aio = 1**
  - **Linux und Windows**
  - **Pending reads / writes**

Weitere MySQL Parameter

- thread_cache_size = 8
- tmp_table_size = 16M
- max_heap_table_size = 16M
- max_connections = 151
- max_user_connections = 0
- open_files_limit = 1024

Weitere MySQL Parameter

- join_buffer_size = 128k --> 8M
- read_buffer_size = 128k --> 2M
- read_rnd_buffer_size = 256k --> 16M
- sort_buffer_size = 256k --> 8M

- binlog_format = ROW
- binlog_cache_size = 32k --> 1M
- binlog_stmt_cache_size = 32k --> 1M
- binlog_row_image = FULL

Weitere MySQL Parameter

- `datadir = /var/lib/mysql/`
- `default_storage_engine = InnoDB`
- `default_tmp_storage_engine = InnoDB`

- `innodb_old_blocks_pct = 37`
- `innodb_old_blocks_time = 0 --> 500?`

- `key_cache_age_threshold = 300`
- `key_cache_division_limit = 100 --> 63?`

Weitere MySQL Parameter

- `long_query_time = 10.000000 --> 0.5?`
- `slow_query_log = OFF --> 1`
- `slow_query_log_file = /var/lib/mysql/slow.log`

- `optimizer_switch = index_merge=on, index_merge_union=on,index_merge_sort_union=on`
- `optimizer_trace = enabled=off,end_marker=off, one_line=off`
- `optimizer_trace_features = greedy_search=on, range_optimizer=on,dynamic_range=on,repeated`

- `performance_schema = ON`

Weitere Hilfe

- Wie messen?
  - `SHOW GLOBAL STATUS;`
  - `SHOW ENGINE INNODB STATUS\G`
- ca. 330 Variablen
- ca. 310 Status Informationen
- MySQL Database Health Check:
  - http://www.fromdual.com/mysql-database-health-check
- MySQL Doku, Server Status Variablen:

  http://dev.mysql.com/doc/refman/5.5/en/server-status-variables.html
- MySQL Performance Monitor:
  - http://www.fromdual.ch/mysql-performance-monitor
  - Auch als Monitoring as a Service (MaaS) Angebot!

Q & A

Fragen ?

Diskussion?

Wir haben noch Zeit für persönliche und indviduelle Beratung

und bieten

Support, Schulung und Betrieb für MySQL