

Performance Tuning & Scale-Out mit MySQL

**Erfa-Gruppe Internet Briefing
2. März 2010**

Oli Sennhauser

Senior MySQL Consultant, FromDual

oli.sennhauser@fromdual.com






Inhalt

- **Allgemeines zu Performance Tuning**
- **MySQL Performance Tuning**
- **Scale-Out Architekturen**
- **Und wie weiter?**
- **Fragen & Antworten**




Szenarien

- 3 Entwicklungsszenarien:



Die FromDual Performance-Wippe

Ursachen für schlechte Performance?



Wo ist der Flaschenhals?

- **Wir haben 4 Möglichkeiten:**
 - **CPU** (Taktrate, Single-Threading, SQL-Queries)
 - **Memory** (RAM, 32/64-bit OS)
 - **Disk** (RAID, SAN, I/O Cache)
 - **Network** (selten und wenn, dann Konfiguration)

Virtualisierung?



Wie finde ich den Flaschenhals?

- **shell> free**
- **shell> vmstat 1**
- **shell> iostat -xk 1**
- **shell> ifconfig**
- **Nagios, Cacti, MySQL Enterprise Monitor, etc.**
- **Windows: Performance Monitor**

Siehe man-Pages.





Wie messe ich MySQL?

- `mysql> SHOW PROCESSLIST;`
- `mysql> SHOW GLOBAL STATUS;`
- `mysql> SHOW ENGINE INNODB STATUS;`

- **MySQL Enterprise Montior** [1]

- **MyTop** [2], **InnoTop** [3]

- **MySQL Activity Report** [4]



oder: **MySQL Monitoring Solutions** [5]



Die wichtigsten MySQL Datenbank Parameter ^[6]

- `key_buffer_size` (MyISAM: 25 – 33% RAM)
- `innodb_buffer_pool_size` (InnoDB: 80% RAM)
- `innodb_log_file_size` (InnoDB: 32 – 256 Mbyte)
- `innodb_flush_log_at_trx_commit` (InnoDB: 2)
- `sync_binlog` (InnoDB: 0 oder >> 1)
- `query_cache_size` (MySQL: 16 – 128 Mbyte)
- Database Health Check for MySQL ^[7]

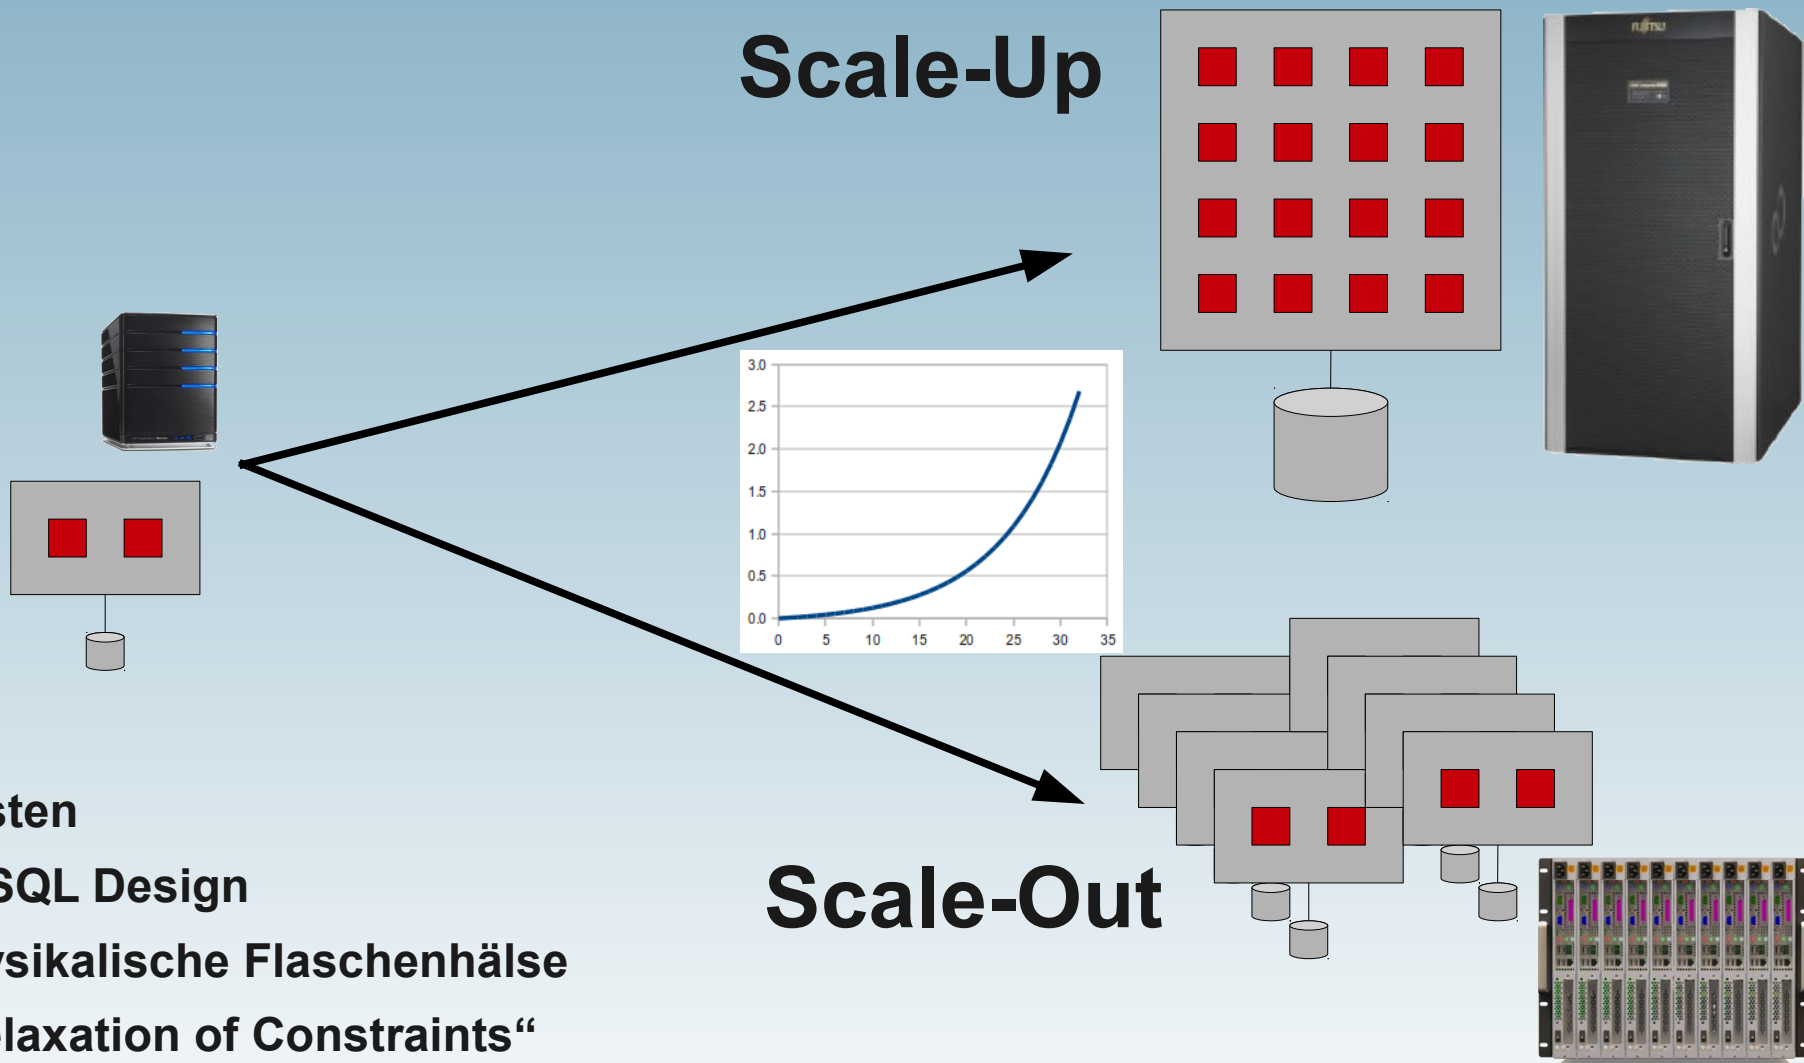


SQL Query Tuning

- **Slow Query Log** (`slow_query_log`, `slow_query_log_file`, `long_query_time`)
- **`log_queries_not_using_indexes`** (deprecated!)
- **`mysqldumpslow -s t slow.log`**
 - **Profile**
- **`EXPLAIN SELECT ... FROM ...`**
 - **Execution Plan der Abfrage** (nur `SELECT`)



Scale-Up vs. Scale-Out



- **Kosten**
- **MySQL Design**
- **Physikalische Flaschenhalse**
- **„Relaxation of Constraints“**




Was für Komponenten nehme ich idealerweise?

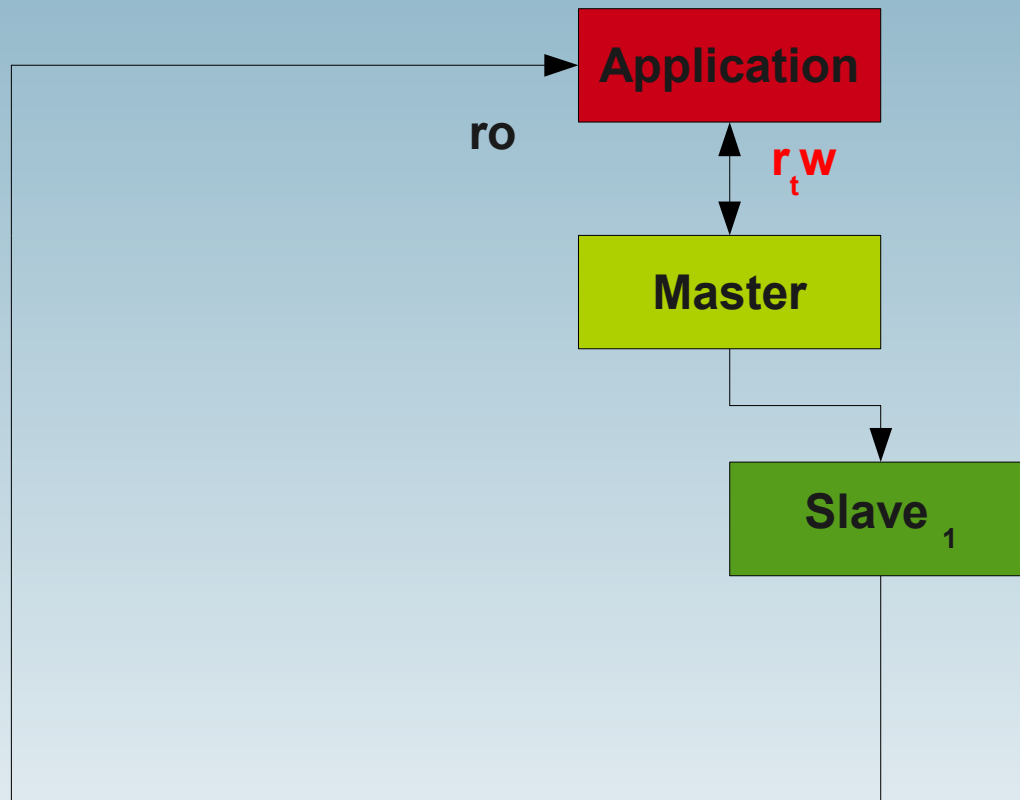
- **LAMP³ evtl. Java?**
- **„Mainstream“ → Linux (Solaris, Windows)**
- **64-bit OS und Hardware (Intel, AMD, KEIN Sparc!) und viel RAM!**
- **Dediziertes I/O System (KEIN SAN, KEIN RAID-5!), schnelle Disks, Write Cache.**
- **>= 1 Gbit Netzwerk**
- **KEINE Virtualisierung!**



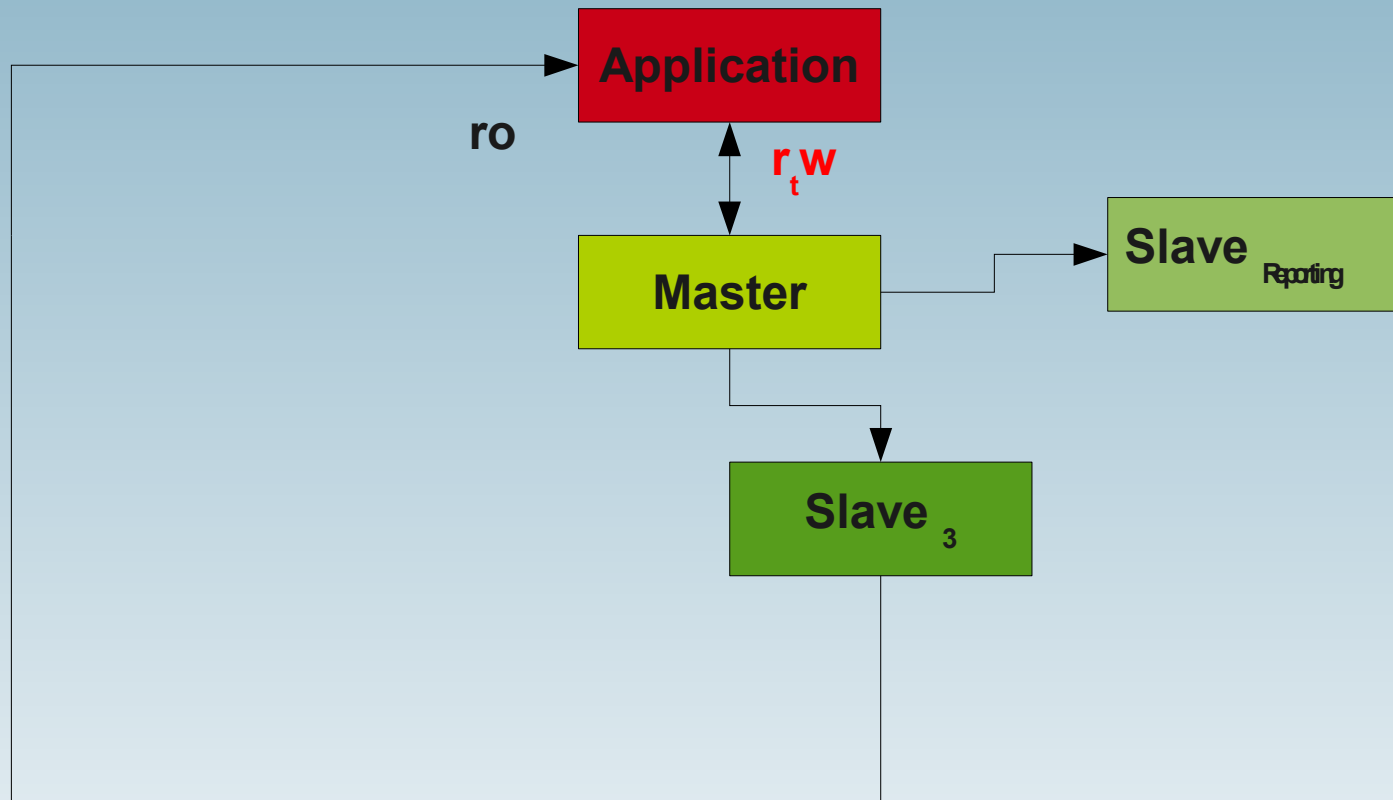
Der MySQL Scale-Out Ansatz



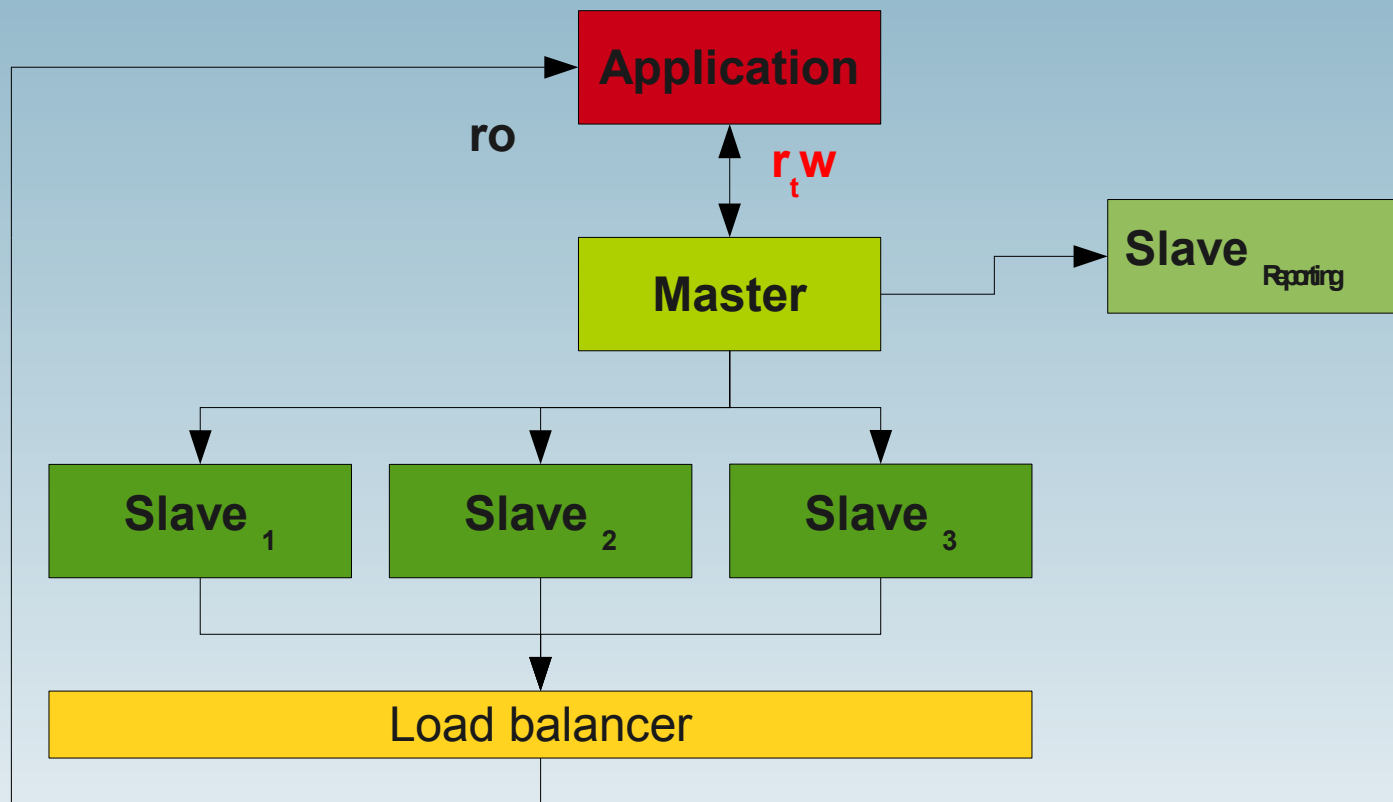
Der MySQL Scale-Out Ansatz



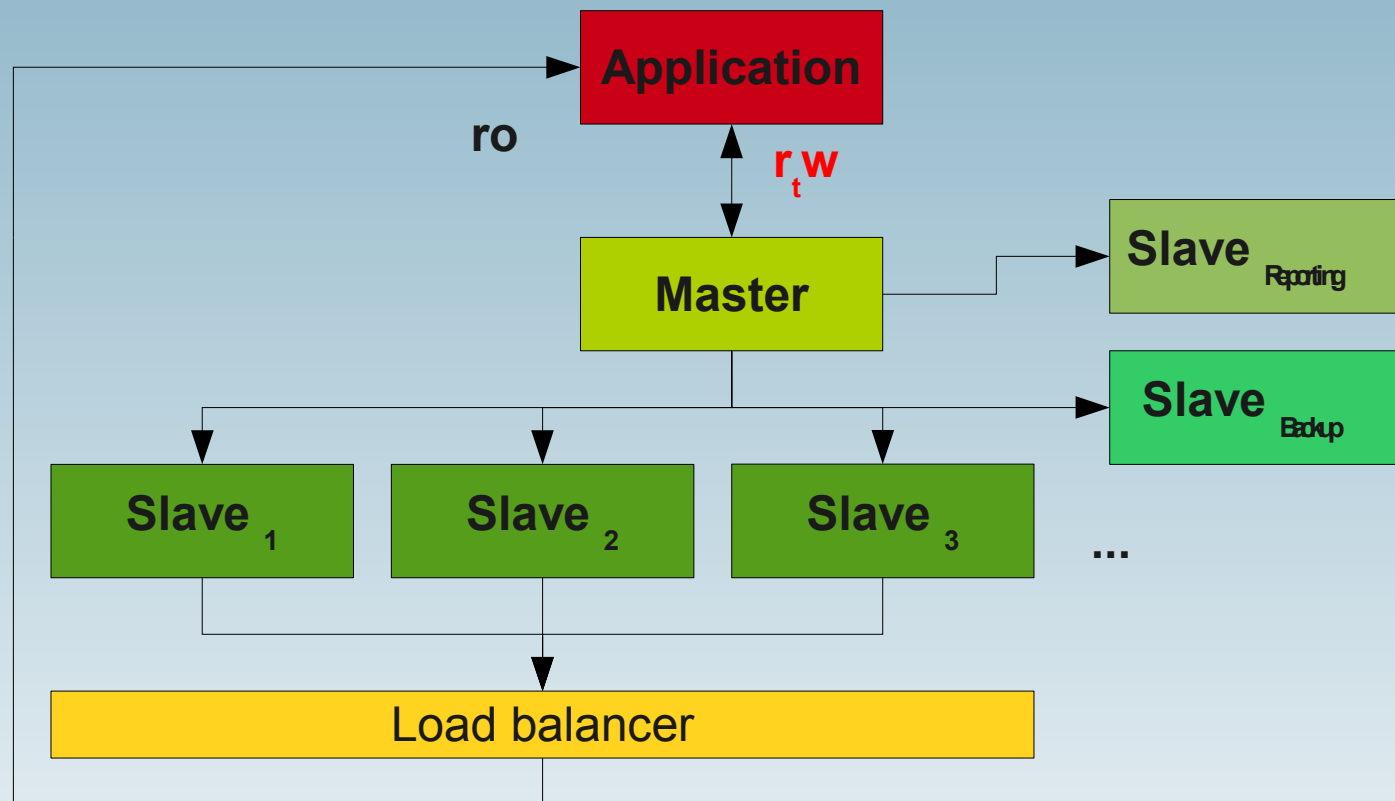
Der MySQL Scale-Out Ansatz



Der MySQL Scale-Out Ansatz



Der MySQL Scale-Out Ansatz



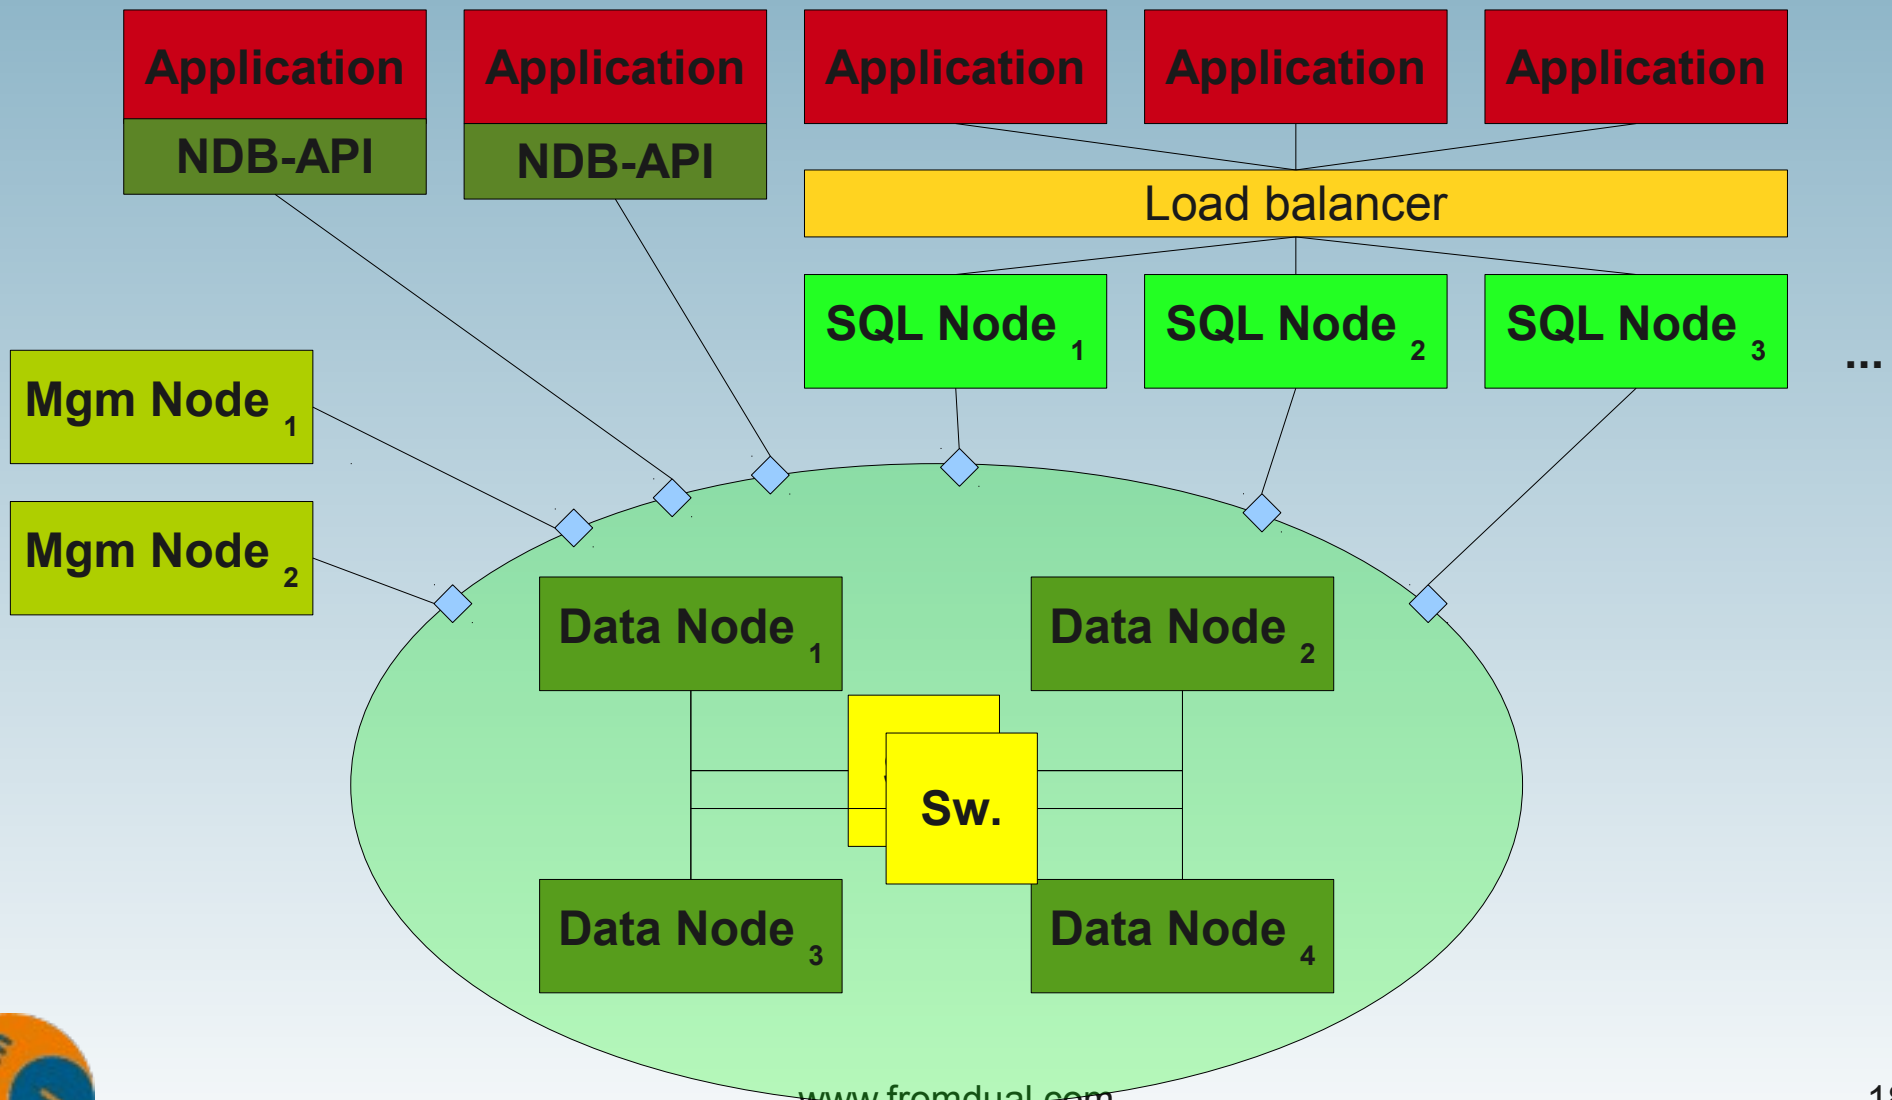
Weitere Möglichkeiten

Irgendwann stehen wir beim Schreib-Durchsatz an...

- Stärkeres I/O system
- Solid State Drive/Disk (SSD)
Bis > 1 Tbyte, €€€, ca. 10-20 x schneller (Marketing: 100x!)
- MySQL Cluster
In-Memory, Realtime, HP- und HA Datenbank



MySQL Cluster



Und wenn das noch nicht reicht? - Sharding! [9]

„Horizontales Partitionieren“ der Daten

- Separieren



The screenshot shows a Facebook news feed for user Oli Sennhauser. The interface includes a search bar at the top, navigation links for Home, Profile, and Account, and a left sidebar with navigation options like News Feed, Messages, Events, Photos, Friends, Applications, Games, Marketplace, and Zimride. The main content area displays a 'News Feed' with posts from Erlend Dahl Bøon, Andrew Hutchings, Bryan Alsdorf, Saskia Schweitzer, and Jad Khoury. A sponsored post for MySQL@Facebook is also visible. The right sidebar contains 'Suggestions' for Aksel Lund Svindal and Ulf Wendel, a 'Sponsored' section for creating an ad, and 'Events' for Andrew Hutchings's birthday. A 'Connect With Friends' section is also present. At the bottom, there is a chat notification for 19 messages.

- Splitten



Wie bereite ich meine Applikation darauf vor?

- ro / rw Datenbank-Verbindungen
- modularisieren / kapseln
- Asynchron und lose gekoppelt, Anforderungen lockern, Cachen
- Messen, messen, messen...
- Verschiedene Lasttest-Szenarios fahren.
- Ich muss meine Applikation kennen!



Literatur

- **Diese Präsentation:**
<http://www.fromdual.com/presentations>

[1] MySQL Enterprise Monitor: <http://dev.mysql.com/doc/refman/5.1/en/mem-introduction.html>

[2] MyTop: <http://jeremy.zawodny.com/mysql/mytop/>

[3] InnoTop: <http://www.xaprb.com/blog/2006/07/02/innotop-mysql-innodb-monitor/>

[4] MySQL Activity Report: <http://gert.sos.be/en/projects/mysqlar/>

[5] MySQL Monitoring Solutions: http://www.shinguz.ch/MySQL/mysql_monitoring.html

[6] MySQL Datenbank Parameter: <http://dev.mysql.com/doc/refman/5.1/en/server-system-variables.html>

[7] Database Health Check for MySQL: http://www.shinguz.ch/database_health_check/

[8] MySQL Cluster: <http://dev.mysql.com/doc/refman/5.1/en/mysql-cluster.html>

[9] Sharding: http://en.wikipedia.org/wiki/Shard_%28database_architecture%29



Fragen & Antworten

?

und sonst →

oli.sennhauser@fromdual.com

