

# MySQL in großen Umgebungen

07.03.2012 | CeBIT

 **NETWAYS**®

Bernd Erk  
[www.netways.de](http://www.netways.de)



**KURZVORSTELLUNG**

**MySQL STATUS QUO**

**STORAGE  
ENGINES**

**MONITORING & MANAGEMENT**

**ENTERPRISE FEATURES**

**REPLIKATION**

**FRAGEN UND ANTWORTEN**



# KURZVORSTELLUNG

# NETWAYS Expertise

## OPEN SOURCE SYSTEMS MANAGEMENT

- Monitoring & Reporting
- Configuration Management
- Service Management
- Knowledge Management
- Backup & Recovery

## OPEN SOURCE DATA CENTER

- High Availability & Clustering
- Cloud Computing
- Load Balancing
- Virtualization
- Database Management

MANAGED SERVICES

MONITORING HARDWARE

KONFERENZEN

# Community Aktiviäten



**[www.netways.org](http://www.netways.org)**

- NETWAYS Addons
- NETWAYS Plugins



**[www.icinga.org](http://www.icinga.org)**

- Development
- Hosting



**[www.monitoringexchange.org](http://www.monitoringexchange.org)**

- Repository of Icinga / Nagios addons and plugins
- ~2000 projects

# NETWAYS Konferenzen



**OSDC.de**  
OPEN SOURCE DATA  
CENTER CONFERENCE

## Open Source Data Center Conference 25 – 26 April 2012

- 100 Teilnehmer (2011)
- "Agile Infrastructures"
  - ▶ Devops & methods
  - ▶ Databases
  - ▶ Scalability & infrastructure



OPEN SOURCE  
**MONITORING**  
CONFERENCE

on Nagios

## Open Source Monitoring Conference 17 – 18 Oktober 2012

- 260 Teilnehmer (2011)
- Icinga / Nagios case studies & best practices
- Latest monitoring technologies & addons

Sonstiges



[www.netways.de/jobs](http://www.netways.de/jobs)

[www.netways.de](http://www.netways.de)



# MYSQL – STATUS QUO



# MySQL Status Quo

- MySQL 5.1
  - ▶ Partitioning
  - ▶ Row-based replication
  - ▶ Better XML-Support
- MySQL 5.5
  - ▶ InnoDB Standard Engine
  - ▶ Performanceoptimierung
  - ▶ Semisynchronous-Replikation
- MySQL 5.6
  - ▶ Time-Delayed-Replication
  - ▶ Partitioning improvements
  - ▶ NoSQL Interface



# STORAGE ENGINES

# InnoDB vs. MyISAM

## ■ InnoDB


- ▶ Row-Level-Locking
- ▶ Foreign-Key Constraints
- ▶ Transaktionssicherheit
- ▶ „Langsamer“
- ▶ Hoher Speicher und Platzbedarf

## ■ MyISAM

- ▶ Table-Level-Locking
- ▶ Unterstützung Volltextindexierung
- ▶ Support für Geodaten

# Die richtige Wahl

- Prüfung der verfügbaren Features der Version
- Prüfung auf Bedarf von Transaktionssicherheit
- Prüfung des Datenmodells
- Prüfung spezieller Anforderungen
- DIE richtige Engine gibt es nicht!





# MONITORING & MANAGEMENT

# Monitoring & Management


- Überwachung der vorhandenen DB-Services
- Erkennung von Auslastungsveränderungen
- Langfristige Identifizierung von Trends
- Replikationsüberwachung
- Unterstützung im Performancetuning
- Sicherheitsüberwachung von DB-Accounts

# Monitoring

- MySQL mit allen gängigen Monitoringlösungen überwachbar
- Gute Plugins und Trendanalyse für Nagios & Icinga



<a href="#">net-mysql bufferpool-hitrate</a>		OK	27-02-2011 21:21:21	130d 7h 5m 33s	1/5	OK - Innodb Buffer Pool Hitrate at 99.84%
<a href="#">net-mysql keycache-hitrate</a>		OK	27-02-2011 21:21:24	115d 9h 7m 20s	1/5	OK - MyISAM Key Cache Hitrate at 92.36%
<a href="#">net-mysql slave-lag</a>		CRITICAL	27-02-2011 21:24:06	0d 0h 15m 19s	5/5	CRITICAL - Slave lag NULL
<a href="#">net-mysql threadcache-hitrate</a>		OK	27-02-2011 21:21:25	25d 23h 15m 53s	1/5	OK - Thread Cache Hitrate at 99.99%
<a href="#">net-mysql threads</a>		OK	27-02-2011 21:21:14	130d 7h 5m 36s	1/5	OK - Threads_connected: 3
<a href="#">net-mysql tmp-disk-tables</a>		OK	27-02-2011 21:21:25	130d 7h 4m 38s	1/5	OK - 24.93% of 1542748 temp tables were created on disk



# Monitoring & Management

- MySQL Enterprise Monitor
- Grid Control
- WebYog
- Kontrollbase





# Kontrollbase

- Webbasierte Managementplattform
- Open Source
- Zentrale oder dezentrale Agenten
- <http://kontrollsoft.com/software-kontrollbase>

The screenshot shows the 'Overview' page of the Kontrollbase management platform. The 'Summary Data' section displays a table with columns for server hostname, load, mem%, CPU, schema, tables, dbconn, total size, #innodb, #myisam, myisam\_total, and innodb\_total. Below this, the 'Critical Environment Alerts' section shows a list of alerts with columns for id, server, state, and name.

server	hostname	load	mem%	CPU	schema	tables	dbconn	total size	#innodb	#myisam	myisam_total	innodb_total
demo01.kontrollbase.com		0.01	32	0.56	3	49	0	47 MB	11	25	538.3 KB	46.5 MB
demo02.kontrollbase.com		0.01	32	0.56	3	49	0	47 MB	11	25	538.3 KB	46.5 MB
mysql.bumpaddler.com		0.02	32	0.99	2	44	0	3.4 MB	0	28	3.4 MB	0 bytes
mysql.relationnews.com		0.01	32	12.54	2	24	0	219.6 MB	2	6	202.5 MB	17 MB
total values				14.65	10	166	16	317 MB	24	84	207 MB	110 MB

id	server	state	name
1825	mysql.bumpaddler.com	crit	max_used_connections too high compared to max_connections (ratio over 85% utilization)
1828	mysql.bumpaddler.com	crit	memory usage over 85%
1830	mysql.bumpaddler.com	crit	tmp table to disk ratio too high
1856	mysql.bumpaddler.com	crit	binary logs not set to auto-purge

The screenshot shows the 'Alerts' page of the Kontrollbase management platform. It displays a table with columns for Hostname, Level, Name, Ack, and Ign. The table lists various alerts such as 'sync\_binlog not set', 'query cache size too large', 'innodb buffer too large', etc.

Hostname	Level	Name	Ack	Ign
demo01.kontrollbase.com	warn	sync_binlog not set	II	⊘
demo01.kontrollbase.com	warn	query cache size too large	II	⊘
demo01.kontrollbase.com	warn	innodb buffer too large	II	⊘
demo01.kontrollbase.com	warn	key buffer too large	II	⊘
demo01.kontrollbase.com	warn	sort buffer too large	II	⊘
demo01.kontrollbase.com	warn	join buffer too small	II	⊘
demo01.kontrollbase.com	warn	thread cache too large	II	⊘
demo01.kontrollbase.com	info	innodb_lock_wait_timeout too high	II	⊘
demo01.kontrollbase.com	info	Uptime less than 48 hours	II	⊘
demo02.kontrollbase.com	warn	sync_binlog not set	II	⊘
demo02.kontrollbase.com	warn	query cache size too large	II	⊘
demo02.kontrollbase.com	warn	innodb buffer too large	II	⊘
demo02.kontrollbase.com	warn	key buffer too large	II	⊘
demo02.kontrollbase.com	warn	sort buffer too large	II	⊘
demo02.kontrollbase.com	warn	join buffer too small	II	⊘
demo02.kontrollbase.com	warn	thread cache too large	II	⊘
demo02.kontrollbase.com	info	innodb_lock_wait_timeout too high	II	⊘
demo03.kontrollbase.com	info	3 replicas, less than 48 hours	II	⊘

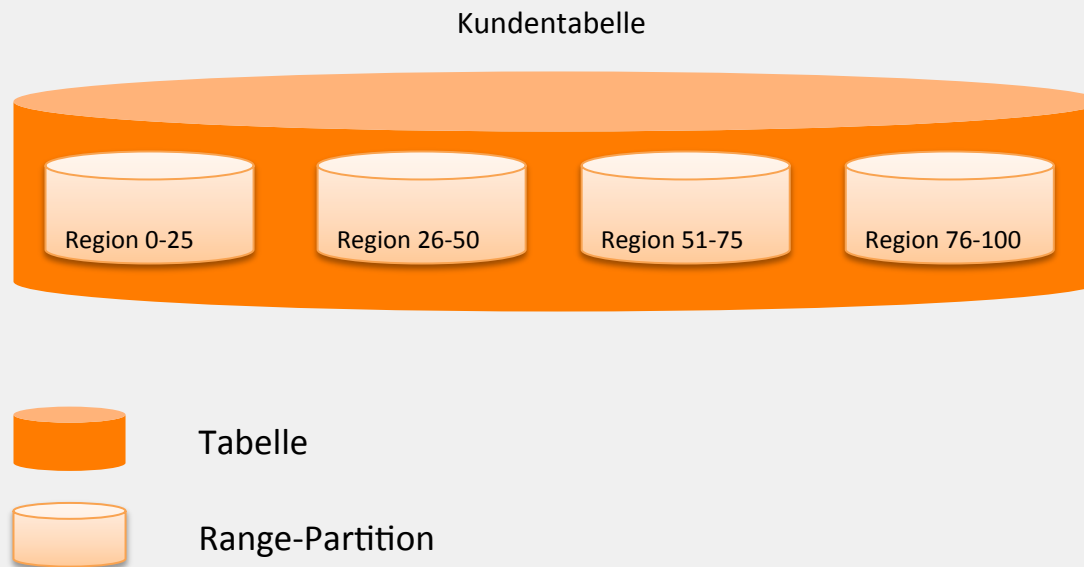


# ENTERPRISE FEATURES

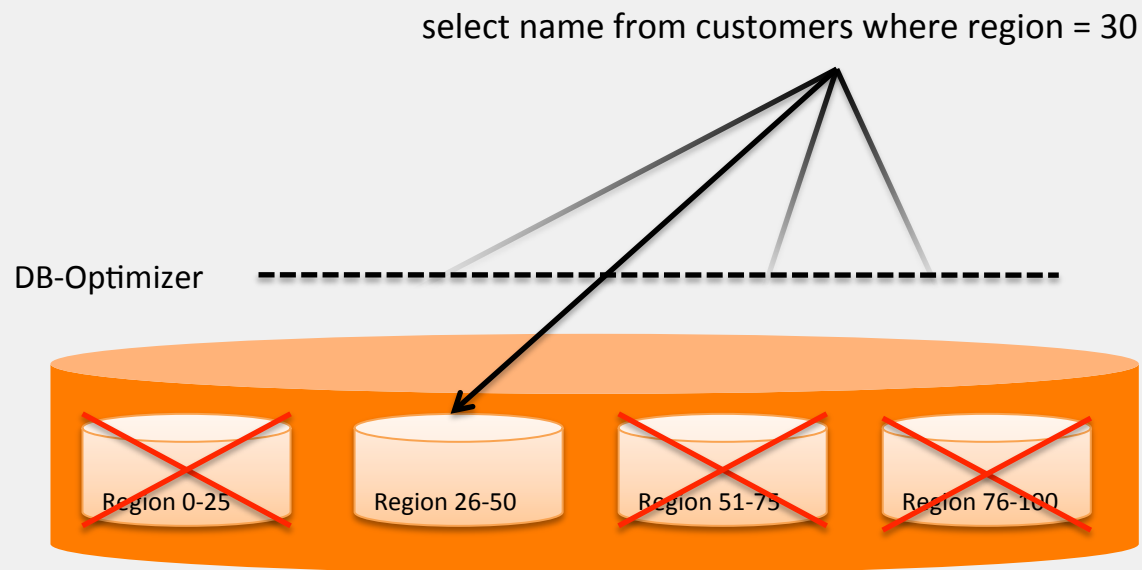
# Enterprise Features

- Partitionierung (5.1)
- mysqlslap (5.1)
- Performance Optimierung (5.1 und 5.5)
- Replikation
  - ▶ Row-Based-Replikation (5.1)
  - ▶ Mixed-Mode-Replikation (5.5)
  - ▶ Semisynchronous-Replikation (5.5)

# Partitionierung



# Partitionierung



# Partitionierung

- RANGE Partitioning
  - ▶ Basierend auf Ausdrücken wie z.B. LESS THAN (100)
- LIST Partitioning
  - ▶ Basierend auf Werten wie z.B. VALUES IN (1,10,20)
- HASH Partitioning
  - ▶ Basierend auf Wertehalt und Aufteilung
- KEY Partitioning
  - ▶ Basierend auf internen Schlüsseln
- Subpartitioning
  - ▶ Subpartition einer anderen Partition

# Partitionierung-Neuerungen

## ■ Manuelle Teilselektion

- ▶ `SELECT * FROM customer PARTITION (p0, p2);`
- ▶ `DELETE FROM customer PARTITION (p3, p4);`

## ■ Partitioning Import/Export

- ▶ `ALTER TABLE customer EXCHANGE PARTITION p0 WITH TABLE new_customer;`

# mysqlslap

- Kleines Werkzeug zur Simulation von Userload
- Durchführung kleiner Last- und Performanctests
- Test und Benchmark in drei Phasen
  - ▶ Erstellung der entsprechenden Tabellen in EINER Connection
  - ▶ Simulation von Last in einer konfigurierbaren Anzahl Connections
  - ▶ Löschung der Tabellen in EINER Connection



# mysqlslap - Beispiel

- `mysqlslap --delimiter=";" --create="CREATE TABLE cebit (halle int, stand int);INSERT INTO cebit VALUES (2, 144)" --query="SELECT * FROM cebit" --concurrency=50 --iterations=200`

Benchmark

```
Average number of seconds to run all queries: 0.019 seconds
Minimum number of seconds to run all queries: 0.000 seconds
Maximum number of seconds to run all queries: 0.039 seconds
Number of clients running queries: 50
Average number of queries per client: 1
```

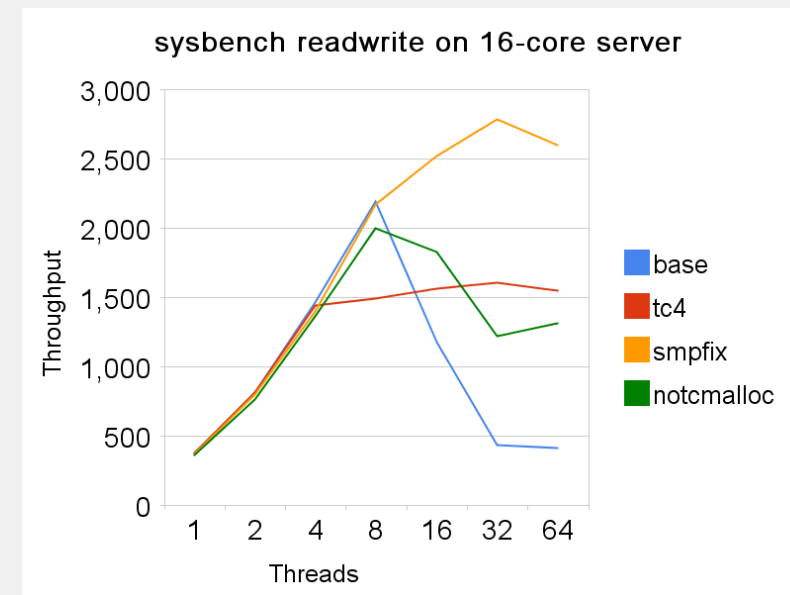

- `mysqlslap --delimiter=";" --create="CREATE TABLE cebit (halle int, stand int);INSERT INTO cebit VALUES (2, 144)" --query="SELECT * FROM cebit" --concurrency=100 --iterations=200`

Benchmark

```
Average number of seconds to run all queries: 0.022 seconds
Minimum number of seconds to run all queries: 0.000 seconds
Maximum number of seconds to run all queries: 0.083 seconds
Number of clients running queries: 100
Average number of queries per client: 1
```

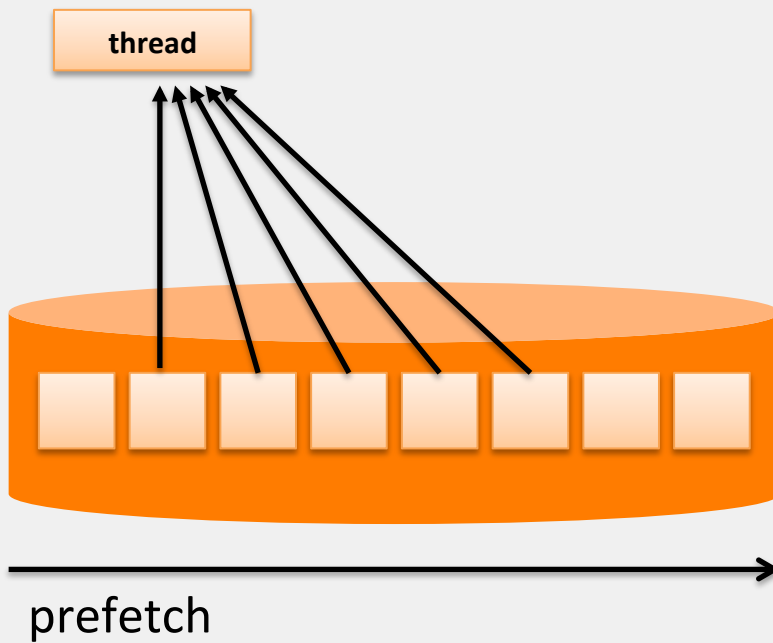
# Optimierung - SMP (Google Patches)

- Ersatz von Mutex durch atomare CPU-Instruktionen
- Speicherallokation der InnoDB Engine




# Performanceoptimierung - InnoDB I/O

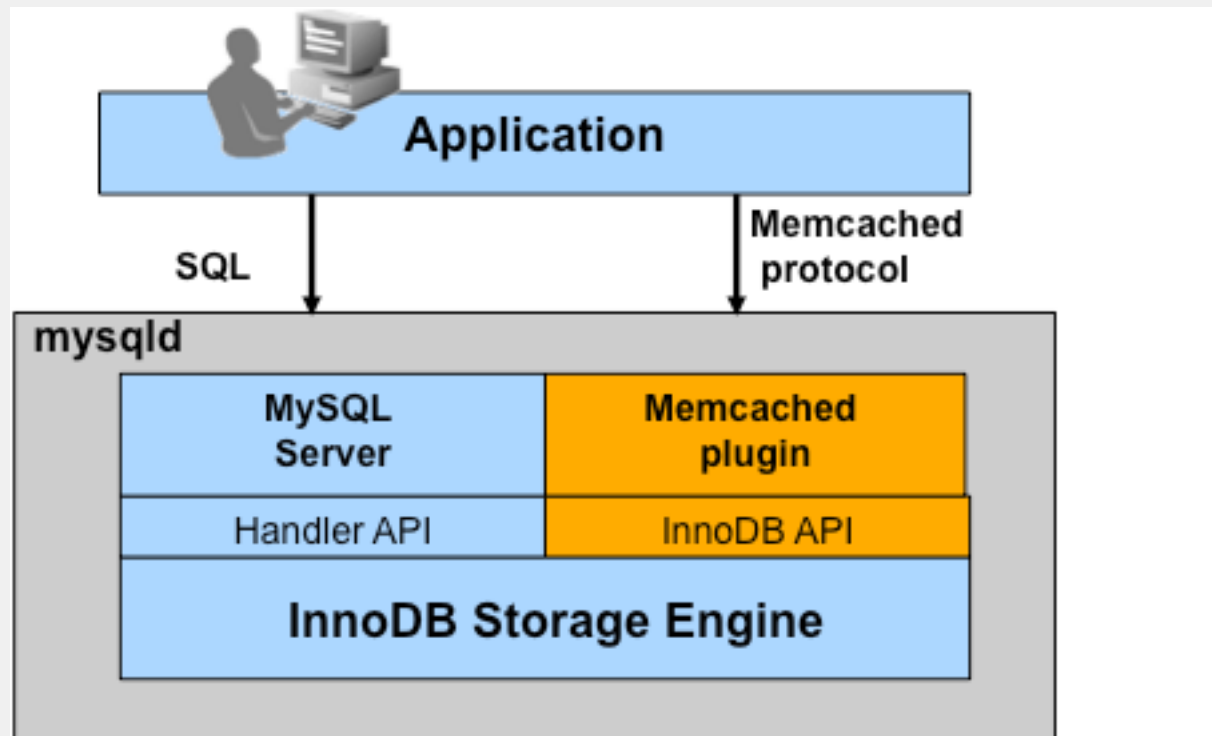
innodb\_read\_io\_threads



innodb\_write\_io\_threads



# Memcached - Plugin






# DATENBANKREPLIKATION

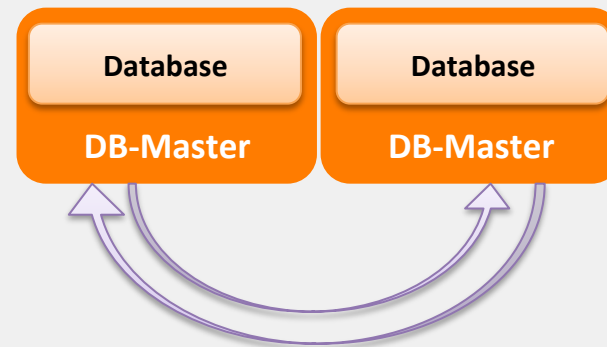
# Datenbankreplikation

- Erhöht die Stabilität und bietet Ausfallschutz
- Erlaubt die kostengünstige Kapazitätssteigerung bei horizontaler Skalierung
- Möglichkeit eines entkoppelten Datenbankbackups ohne Störung des Hauptservers

# Replikationstypen

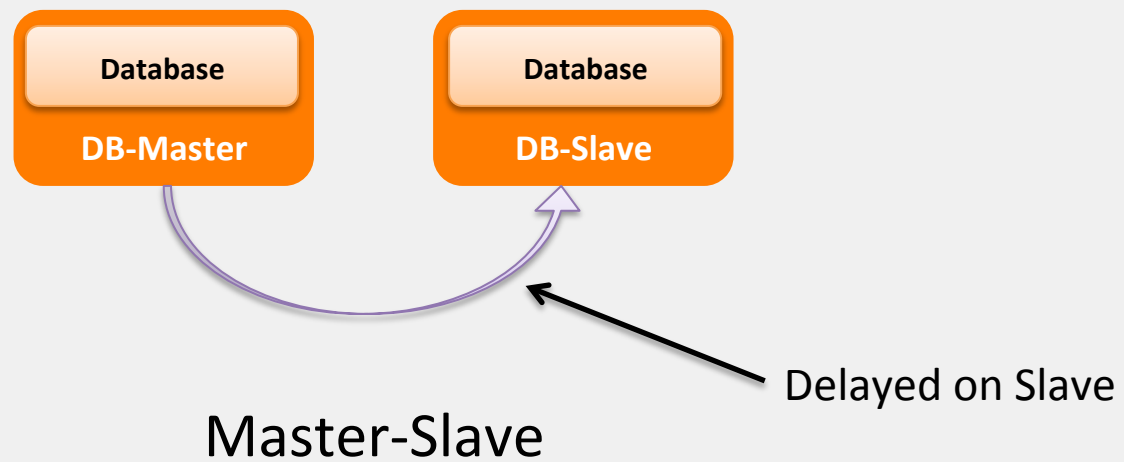


Master-Slave



Master-Master

# Time-Delayed-Replikation

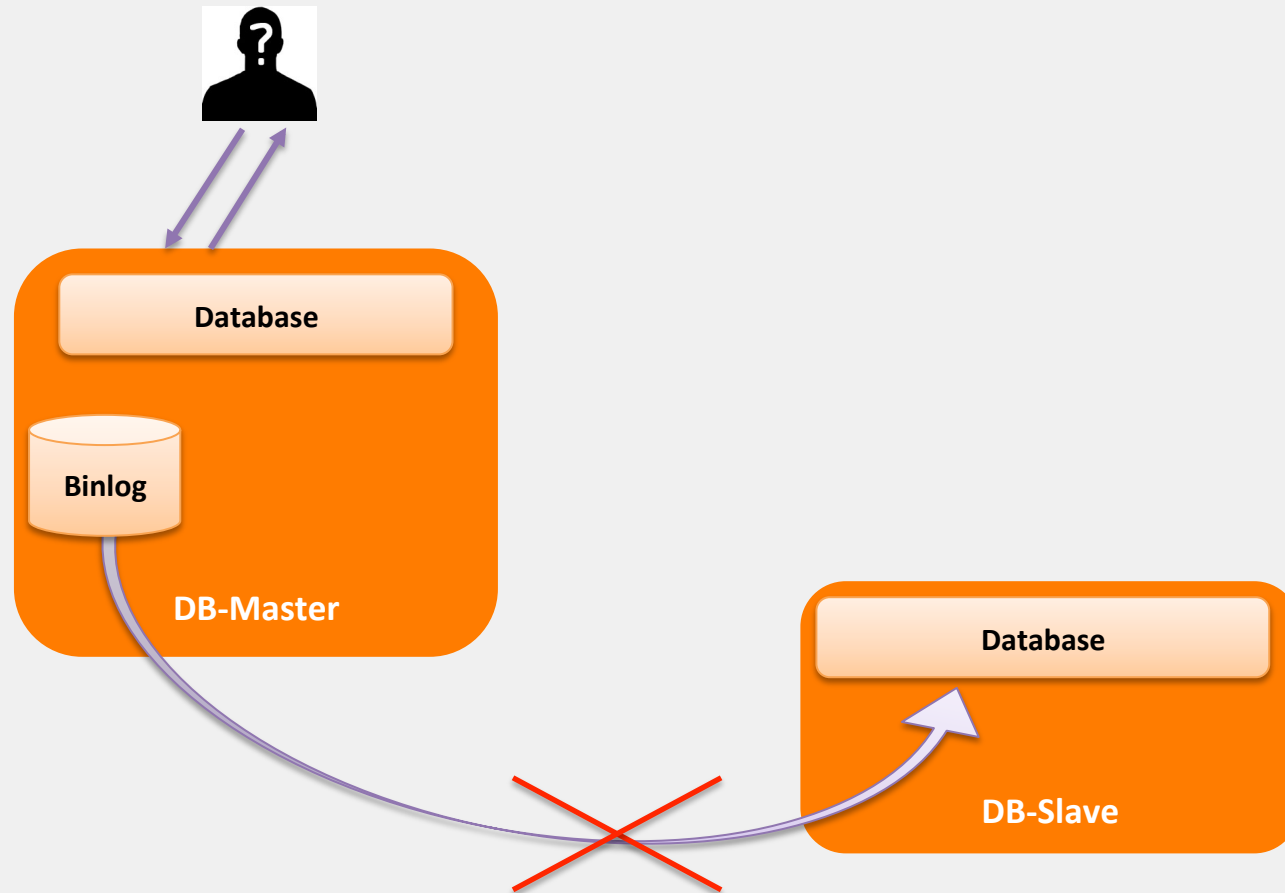




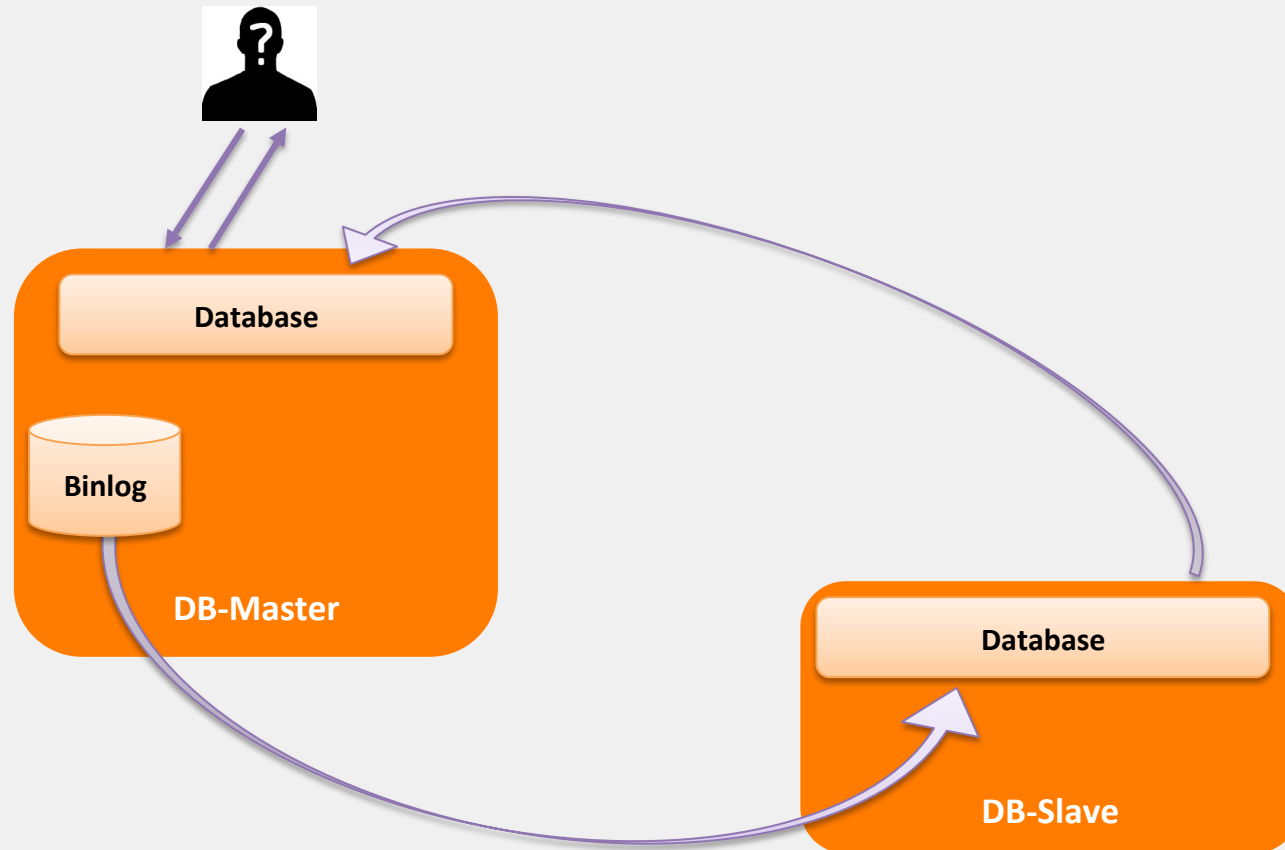
# Statement-Based vs. Row-Based

- Statement-Based-Replication
  - ▶ Standard-Replikationsmechanismus
  - ▶ Logging aller Statements erlaubt sicheren Audit-Trail
  - ▶ Keine Replikation von nicht deterministischer Operation
- Row-Based-Replikation
  - ▶ Alle Operation können repliziert werden, da ausschließlich die Veränderungen übertragen werden
  - ▶ DDL werden automatisch mittels SBR übertragen
  - ▶ Deutlich weniger Locks und sehr schnell bei Verarbeitung kleiner Datenmengen
- Mixed-Mode-Replikation

# Standard-Replikation



# Semisynchronous-Replikation





# FRAGEN UND ANTWORTEN



**Question  
& Answer**

**HALLE 2 – D58**

**NETWAYS GmbH**

**Deutschherrnstrasse 15-19  
90429 Nürnberg**

**Tel:** +49 911 92885-0

**Fax:** +49 911 92885-77

**Email:** [info@netways.de](mailto:info@netways.de)

**Twitter:** [twitter.com/netways](https://twitter.com/netways)

**Blog:** [blog.netways.de](http://blog.netways.de)