



# **Need for Speed: MySQL Indexing**

**Percona Live 2013,  
November 11 – 12, London**

**Oli Sennhauser**

**Senior MySQL Consultant at FromDual GmbH**

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



# About FromDual GmbH

- **FromDual provides neutral and independent:**
  - **Consulting for MySQL, Percona Server, MariaDB**
  - **Support for all MySQL and Galera Cluster**
  - **Remote-DBA Services**
  - **MySQL Training**
- **Oracle Silver Partner (OPN)**



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

# Our customers



# MySQL and Indexing

- MySQL documentation says:

*The best way to improve the performance of `SELECT` operations is to create indexes on one or more of the columns that are tested in the query.*

- Great! But:

*Unnecessary indexes waste space and waste time to determine which indexes to use. You must find the right balance to achieve fast queries using the optimal set of indexes.*

- ... hmm so we have to think a bit... :-)



# What is an Index technically?

Aa-Bv	Abächerli	Ptr
A-E Bw-Bz	Bucher	Ptr
Ca-Ez	Cathomas	Ptr
Fa-Gm	Fuchs	Ptr
F-M Gn-Jf	Haas	Ptr
Jg-Mu	Kübler	Ptr
Mv-Se	Sennhauser	233
N Sf-Wt	Vögeli	Ptr
Wc-Zz	Zürcher	Ptr

1, Cathomas, Marco, Bachweg  
 3, 1234, Lausanne; 2,  
 Abächerli, Hans,  
 Hauptstrasse 13, 8001,  
 Zürich; 3, Haas, Daniel,  
 Rainstrasse 34, 8600,  
 Dübendorf; 4, Fuchs, Hans-  
 Peter, Rebenweg 3, 8610,  
 Uster; 6, Zürcher, Barbara,  
 Zürichstrasse 1a, 8000,  
 Turbental; 8, Sennhauser,  
 Oli, Rebenweg 6, 8610 Uster;  
 9, Kübler Köbi, Tschutiplatz  
 1, 8000, Zürich; 10, Vögeli  
 Fridolin, Peterstrasse 3,  
 8610, Uster; 11, Bucher,  
 Walti, Peterstrasse 1, 8610,  
 Uster

Aa-Bv	Abächerli	PK
A-E Bw-Bz	Bucher	PK
Ca-Ez	Cathomas	PK
Fa-Gm	Fuchs	PK
F-M Gn-Jf	Haas	PK
Jg-Mu	Kübler	PK
Mv-Se	Sennhauser	8
N Sf-Wt	Vögeli	PK
Wc-Zz	Zürcher	PK

1  
 1-3 2  
 3  
 4  
 4-6 5  
 6  
 7  
 7-8 8  
 9

1, Cathomas, Marco, Bachweg  
 3, 1234, Lausanne; 2,  
 Abächerli, Hans,  
 Hauptstrasse 13, 8001,  
 Zürich; 3, Haas, Daniel,  
 Rainstrasse 34, 8600,  
 Dübendorf; 4, Fuchs, Hans-  
 Peter, Rebenweg 3, 8610,  
 Uster; 6, Zürcher, Barbara,  
 Zürichstrasse 1a, 8000,  
 Turbental; 8, Sennhauser,  
 Oli, Rebenweg 6, 8610 Uster;  
 9, Kübler Köbi, Tschutiplatz  
 1, 8000, Zürich; 10, Vögeli  
 Fridolin, Peterstrasse 3,  
 8610, Uster; 11, Bucher,  
 Walti, Peterstrasse 1, 8610,  
 Uster

# MySQL uses indexes:

- To enforce uniqueness (**PRIMARY KEY, UNIQUE KEY**)
- To fast access and filter rows (**WHERE**)
- To perform joins fast (**JOIN**)
- To find **MIN ( )** and **MAX ( )** values
- For sorting and grouping (**ORDER BY, GROUP BY**)
- To avoid joins by using covering indexes
- To enforce **FOREIGN KEY Constraints (FOREIGN KEY)**

# WHERE clause 1

```
SELECT *
  FROM customers
 WHERE name = 'No
```

```
SHOW CREATE TABLE customers\G
```

```
CREATE TABLE `customers` (
  `customer_id` smallint(5) unsigned
, `name` varchar(64) DEFAULT NULL
  PRIMARY KEY (`customer_id`))
```

**EXPLAIN**

```
SELECT *
  FROM customers
 WHERE name = 'No Clue of MySQL LLC';
```

table	type	possible_keys	key	rows	Extra
customers	ALL	NULL	NULL	31978	Using where



# How to create and Index?

**ALTER TABLE ...**

- **ADD PRIMARY KEY (id);**
- **ADD UNIQUE KEY (uuid);**
- **ADD FOREIGN KEY (customer\_id)  
REFERENCES customers (customer\_id);**
- **ADD INDEX (last\_name, first\_name);**
- **ADD INDEX pre\_ind (hash(8));**
- **ADD FULLTEXT INDEX (last\_name,  
first\_name);**

# WHERE clause 2

```
ALTER TABLE customers
  ADD INDEX (name);
```

**Gain: 20 ms → 5 ms**

```

s` (
nt(5) unsigned
DEFAULT NULL
er_id`)
, KEY `name` (`name`)
)
```

table	type	possible_keys	key	key_len	ref	rows
customers	ref	name	name	67	const	1

# JOIN clause

```
EXPLAIN SELECT *
  FROM customers AS c
  JOIN orders AS o ON c.customer_id = o.customer_id
 WHERE c.name = 'No Clue of MySQL LLC';
```

table	type	possible_keys	key	key_len	ref	rows
c	ref	PRIMARY	name	67	const	1
o	ALL	NULL				1045105

**Gain: 450 ms → 6 ms**

```
ALTER TABLE orders
  ADD INDEX (customer_id);
```

table	type	possible_keys	key	key_len	ref	rows
c	ref	PRIMARY	name	67	const	1
o	ref	customer_id	customer_id	3	c.customer_id	8

# For sorting/grouping tables

## ORDER BY, GROUP BY

```
EXPLAIN SELECT *
  FROM contacts AS c
 WHERE last_name = 'Sennhauser'
 ORDER BY last name, first name;
```

table	type	key	rows	Extra
c	ref	last_name	1561	Using filesort

**Gain: 20 ms → 7 ms**

```
ALTER TABLE contacts
  ADD INDEX (last_name, first_name);
```

table	type	key	rows	Extra
contacts	ref	last_name_2	1561	Using where; Using index

# Covering Indexes

```
EXPLAIN
SELECT customer_id, amount
FROM orders AS o
WHERE customer_id = 59349;
```

table	type	key	rows	Extra
o	ref		15	Using index

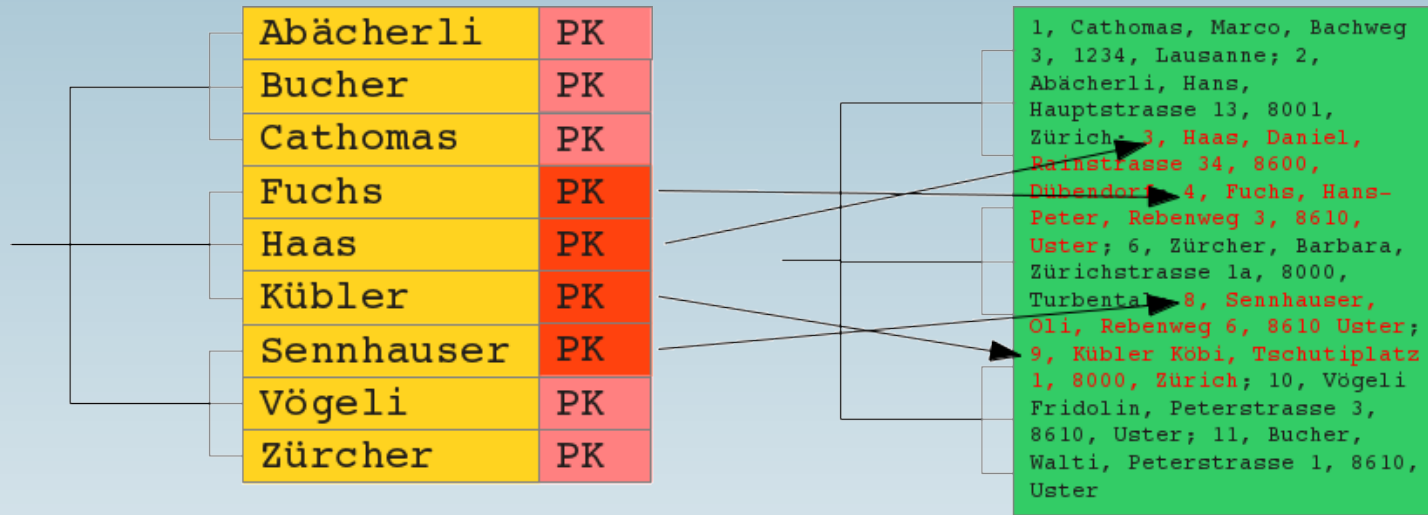
So what?

```
ALTER TABLE orders
ADD INDEX (customer_id, amount);
```

table	type	key	rows	Extra
o	ref	customer_id_2	15	Using index

# Benefit of Covering Indexes

- Why are Covering Indexes beneficial



Abächerli	Hans	PK
Bucher	Walti	PK
Cathomas	Marco	PK
Fuchs	Hans-Peter	PK
Haas	Daniel	PK
Kübler	Köbi	PK
Sennhauser	Oli	PK
Vögeli	Fridolin	PK
Zürcher	Barbara	PK

# How-to find missing indexes?

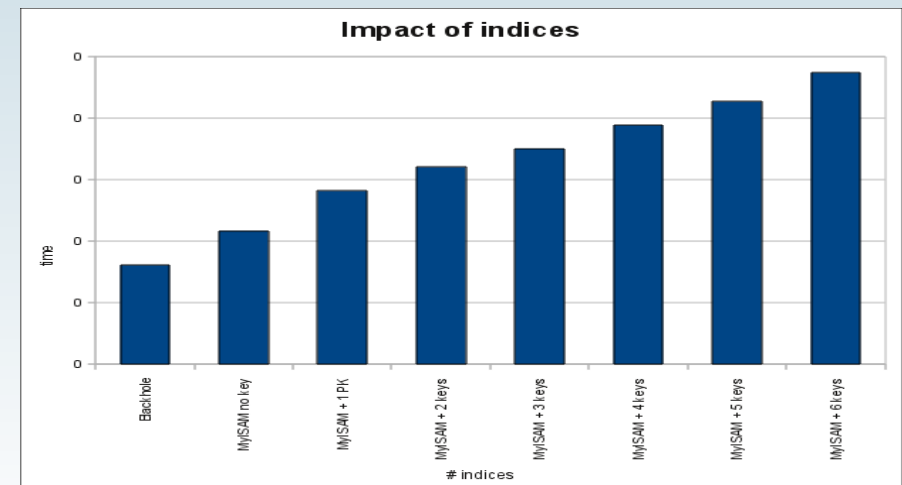
- ER Diagram? :-(
  - Most of them rely to you business logic...
- How-to FIND?
- Slow Query Log
- MySQL Variables:
- Since v5.1 on-line!

Variable_name	Value
log_queries_not_using_indexes	ON
long_query_time	0.250000
min_examined_row_limit	100
slow_query_log	ON
slow_query_log_file	slow.log

# Indexes are not only good

- Indexes use space (Disk, hot data in RAM!)
- Indexes use time to maintain (CPU, RAM, I/O)
- Optimizer needs time to determine which indexes to use.
- Sometimes optimizer is completely confused and does wrong decisions if too many (similar) indexes are there.

→ *You must find the right balance to achieve fast queries using the optimal set of indexes.*

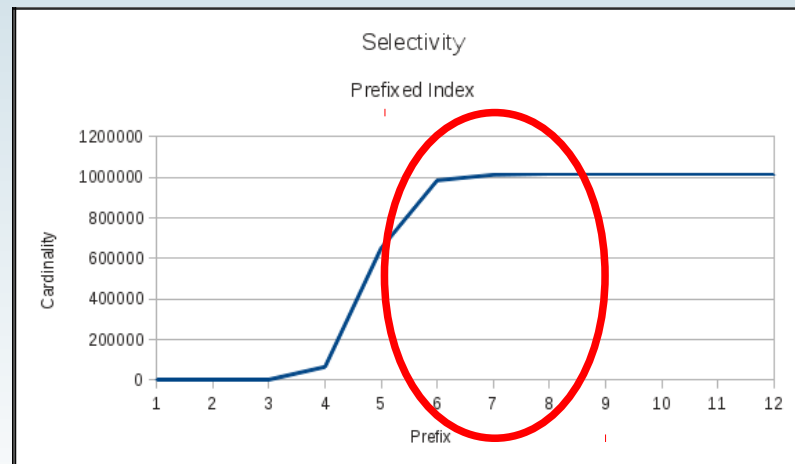




# Smaller indexes faster queries

- Better fit into memory (less I/O)
- Higher data density (rows/block)
- Less CPU cycles (to crawl through)
- Prefixed indexes:

```
ADD INDEX pre_ind (hash(8));
```



# Avoid indexes

- **Avoid redundant (and thus unnecessary ) indexes**
- **How does it happen?**
  - **Developer 1: Creates a Foreign Key constraint → done**
  - **Developer 2: Ouu! Query is slow → Oli told me to create an index! → done**
  - **Developer 3: Ouu! Query is slow → Developer 2 is stupid! → Create an index → done**
- **Frameworks vs. Developer**
- **Upgrade process vs. Developer**
- **Avoid indexes which are not used / needed**

# How to find such indexes?

```
SHOW CREATE TABLE ... \G
```

```
mysqldump --no-data > structure_dump.sql
```

- **Since MySQL 5.6: PERFORMANCE\_SCHEMA**
  - Percona Server / MariaDB: Userstats
  - <http://fromdual.com/mysql-performance-schema-hints>

```
SELECT object_schema, object_name, index_name
FROM performance_schema.table_io_waits_summary_by_index_usage
WHERE index_name IS NOT NULL
AND count_star = 0
ORDER BY object_schema, object_name;
```

# Avoid partial redundant indexes

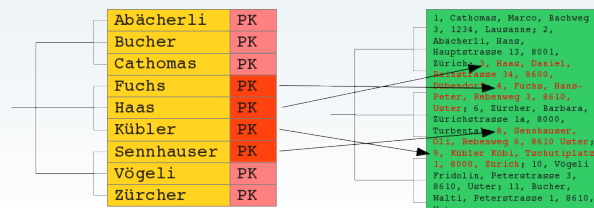
- INDEX (city, last\_name, first\_name)
- ~~INDEX (city, last\_name)~~
- ~~INDEX (city)~~
- INDEX (last\_name ↔ city) ???
- INDEX (first\_name, last\_name) !!!

# Bad selectivity

- Remove indexes with bad selectivity (~= low cardinality)
- Candidates are:
  - status
  - gender
  - active
- How to find if field has bad selectivity?  
 Indexes (and Joins) are expensive!!!
- Break even between 15% and 66%
- Lets see if the MySQL optimizer knows about it... :-)

```
SELECT status, COUNT(*)
FROM orders
GROUP BY status;
```

status	cnt
0	393216
1	262144
2	12
3	36
4	24
5	4
6	8



# Optimizer is wrong!

```
SELECT * FROM orders WHERE status = 2;
+-----+-----+-----+-----+-----+
| table  | type  | possible_keys | key      | rows  |
+-----+-----+-----+-----+-----+
| orders | ref   | status        | status   | 12    |
+-----+-----+-----+-----+-----+
```

```
SELECT status, COUNT(*)
FROM orders
GROUP BY status;
```

```
+-----+-----+
| status | cnt    |
+-----+-----+
| 0      | 393216 |
| 1      | 262144 |
| 2      | 12     |
| 3      | 36     |
| 4      | 24     |
| 5      | 4      |
| 6      | 8      |
+-----+-----+
```

```
SELECT * FROM orders WHERE status = 0;
1.43 s
+-----+-----+-----+-----+-----+
| table  | type  | possible_keys | key      | rows  |
+-----+-----+-----+-----+-----+
| orders | ref   | status        | status   | 327469 |
+-----+-----+-----+-----+-----+
```

```
SELECT * FROM orders IGNORE INDEX (STATUS) WHERE status = 0;
0.44 s
```

```
+-----+-----+
| table  | type  | index |
+-----+-----+
| orders | ALL   | NULL  |
+-----+-----+
```

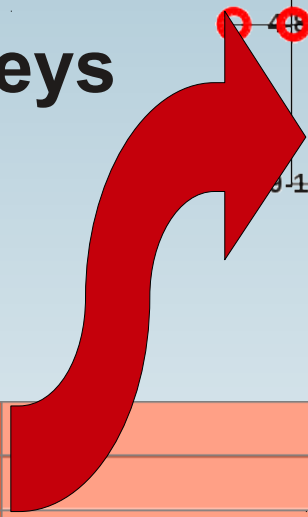
**5.6.12 (after analyze table)**

# InnoDB PK and SK

- InnoDB has
  - Primary Keys and
  - Secondary Keys

1-3	1	Cathomas, Marco, Bachweg 3, 1234, Lausanne
	2	Abächerli, Hans, Hauptstrasse 13, 8001, Zürich
	3	Haas, Daniel, Rainstrasse 34, 8600, Dübendorf
	4	Fuchs, Hans-Peter, Rebenweg 3, 8610, Uster
	6	Zürcher, Barbara, Zürichstrasse 1a, 8000, Turbental
	8	Sennhauser, Oli, Rebenweg 6, 8610 Uster
	9	Kübler Köbi, Tschutiplatz 1, 8000, Zürich
	10	Vögeli Fridolin, Peterstrasse 3, 8610, Uster
	11	Bucher, Walti, Peterstrasse 1, 8610, Uster

Aa-Bv	Abächerli	2	Long PK
A-E Bw-Bz	Bucher	11	Long PK
Ca-Ez	Cathomas	1	Long PK
Fa-Gm	Fuchs	4	Long PK
F-M Gn-Jf	Haas	3	Long PK
Jg-Mu	Kübler	9	Long PK
Mv-Se	Sennhauser	8	Long PK
N-O Sf-Wt	Vögeli	10	Long PK
Wc-Zz	Zürcher	6	Long PK



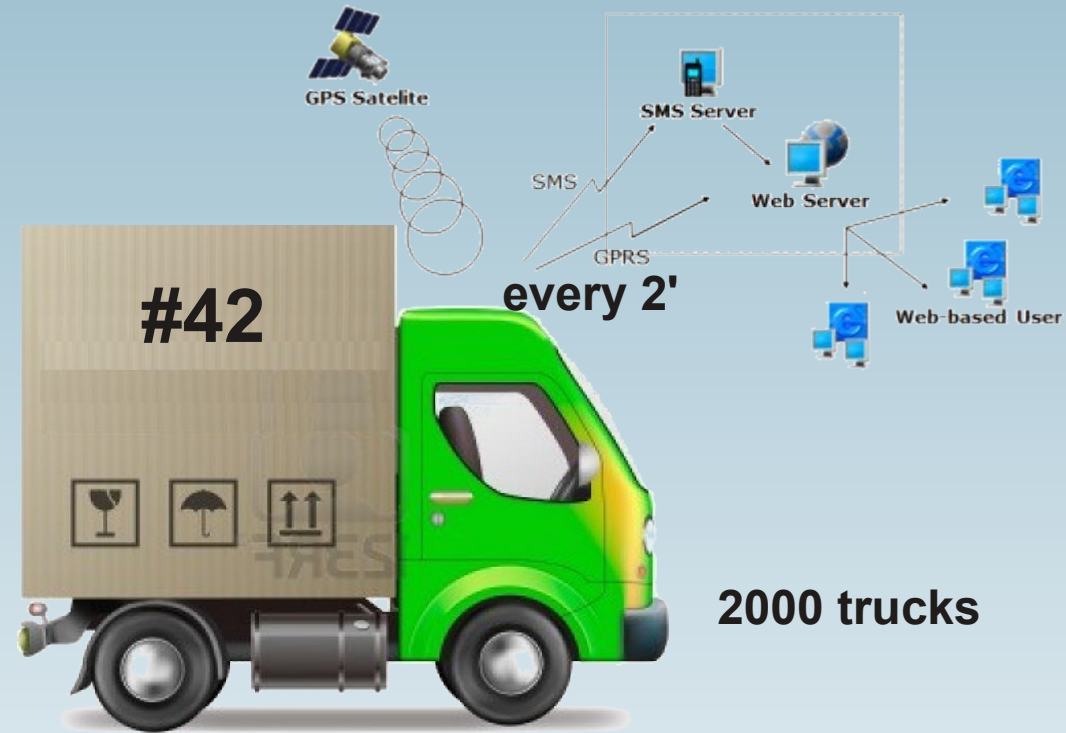
# Clustered Index

- **InnoDB: Data = Leaf of Primary Key**
  - We call this an **Index Clustered Table (IOT)**
    - Data are sorted like PK (key is sorted)!
    - PK influences **Locality of data** (physical location)
- **AUTO\_INCREMENT**  $\approx$  sorting by time!
- **Good for many things**
  - where hot data = recent data
- **Bad for time series**
  - Where hot data = per item data



# Example: InnoDB

<b>A_I</b>	<b>ts</b>	<b>v_id</b>	<b>xpos</b>	<b>ypos</b>	<b>...</b>
1	17:30	#42	x,	y,	...
2	17:30	#43	x,	y,	...
3	17:30	#44	x,	y,	...
...					
2001	17:32	#42	x,	y,	...
2002	17:32	#43	x,	y,	...
2003	17:32	#44	x,	y,	...



Q1:  $\Delta$  in rows? ~ 2000 rows

A1: 1 row ~ 100 byte

Q2:  $\Delta$  in bytes? ~ 200 kbyte

Q3: Default InnoDB block size? default: 16 kbyte

Q4: Avg. # of rows of car #42 in 1 InnoDB block? ~ 1

A2: 3 d and 720 pt/d  $\rightarrow$  ~2000 pt ~ 2000 rec ~ 2000 blk

Q5: How long will this take and why (32 Mbyte)?

~ 2000 IOPS ~ 10s random read!!!

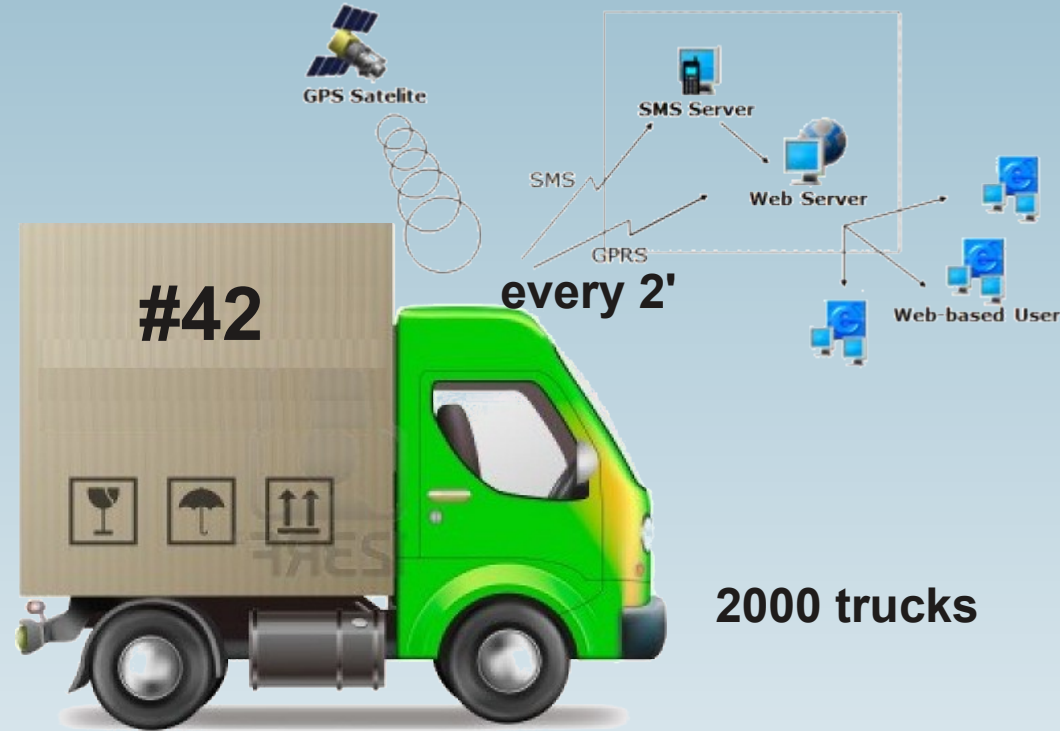
S: All in RAM or strong I/O system or ...?



over the last 3 days

# InnoDB PK safes the day!

ts	v_id	xpos	ypos	...
17:30	#42	x,	y,	...
17:32	#42	x,	y,	...
17:34	#42	x,	y,	...
...				
17:30	#43	x,	y,	...
17:32	#43	x,	y,	...
17:34	#43	x,	y,	...
...				
17:30	#44	x,	y,	...



Q1: Avg. # of rows of car #42 in 1 InnoDB block? ~ 120

A1: 3 d and 720 pt/d → ~2000 pt ~ 2000 rec ~ 20 blk

Q2: How long will this take and why (320 kbyte)?

~ 1-2 IOPS ~ 10-20 ms sequential read!

S: Wow f=50 faster! Any drawbacks?



over the last 3 days

# Index hints

- **MySQL optimizer is sometimes wrong!**
  - We have to help (= hint) him...
- **Index hints are:**
  - **USE INDEX (ind1, ind2)**
    - Only consider these indexes
  - **FORCE INDEX (ind3)**
    - Use this index without considering anything else
  - **IGNORE INDEX (ind1, ind3)**
    - Do NOT consider these indexes but everything else
- **Hints should be used only as a last resort**

# MySQL Variables

- MySQL variables influencing index use
  - MyISAM: `key_buffer_size`
  - InnoDB: `innodb_buffer_pool_size / innodb_buffer_pool_instances`
- InnoDB Change Buffer
  - `innodb_change_buffer_max_size`
  - `innodb_change_buffering`
- Adaptive Hash Index (AHI)
- MySQL 5.6.3 / 5.5.14 index length 767 → 3072 bytes
  - `innodb_large_prefix`

# Q & A



**Questions ?**

**Discussion?**

**We have time for some face-to-face talks...**

- **FromDual provides neutral and independent:**
  - **Consulting**
  - **Remote-DBA**
  - **Support for MySQL, Galera, Percona Server and MariaDB**
  - **Training**

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