



PERCONA
LIVEONLINE

Jose Luis Martinez
Voovio Technologies



Boosting MySQL Performance

October 21
1:00 AM
New York

October 21
1:00 PM
Singapore

October 21
6:00 PM
London

This is not a DB optimization talk

It's a talk about how doing things in a specific way leads to getting way better results by default

"We should forget about small efficiencies, say about 97% of the time: premature optimization is the root of all evil"

-- Donald Knuth



Preoptimization



Don't denormalize



Maintenance prevails
over performance



Use InnoDB

Preoptimization



Don't denormalize
by default




Maintenance prevails
over performance
by default



Use InnoDB
by default

Preoptimization





Tuning MySQL Parameters

MySQL Parameter Tuning



Will get you out of SOME trouble

MySQL Parameter Tuning



Will get you out of SOME trouble



But not a good default strategy,
specially if the base is flawed

MySQL Parameter Tuning



Will get you out of SOME trouble



But not a good default solution,
specially if the base is flawed



Please do tune
MySQLs defaults

Not the theme
for this talk



Start with
your schema

Start with your Schema

- ▶ Your schema is probably the root cause of your “My DB doesn’t scale” problems
 - ▶ The solution is not “have a loose/no schema”
- ▶ How to fake a DB Design (Curtis Ovid Poe)
 - ▶ <https://www.youtube.com/watch?v=y1tcbhWLiUM>

Medlin



Care for
your
datatypes

Data types: how (not) to bloat your DB



SELECTING DATA TYPES WITH A BIT OF CARE IS VERY PRODUCTIVE



IT MAKES MORE DATA FIT IN LESS SPACE



OPTIMIZES USE OF INNODB BUFFER POOL, MYISAM KEY BUFFER, JOIN BUFFER, SORT BUFFER, SMALLER INDEXES



▶ Integers

Type	Bytes	Unsigned	Signed
INT	4	4000M	-2000M to 2000M

INT(1) == INT(10) == 4 bytes

Same range!




Type	Bytes	Unsigned	Signed
INT	4	4000M	-2000M to 2000M

The Integer family

Type	Bytes	Unsigned	Signed
TINYINT	1	255	-128 to 127
SMALLINT	2	65K	-32K to 32K
MEDIUMINT	3	16M	-8M to 8M
INT	4	4000M	-2000M to 2000M
BIGINT	8	1.8×10^{19}	-9×10^{18} to 9×10^{18}

<https://dev.mysql.com/doc/refman/8.0/en/integer-types.html>

Strings



i am in ur
dictionaries

verbing
ur nounz

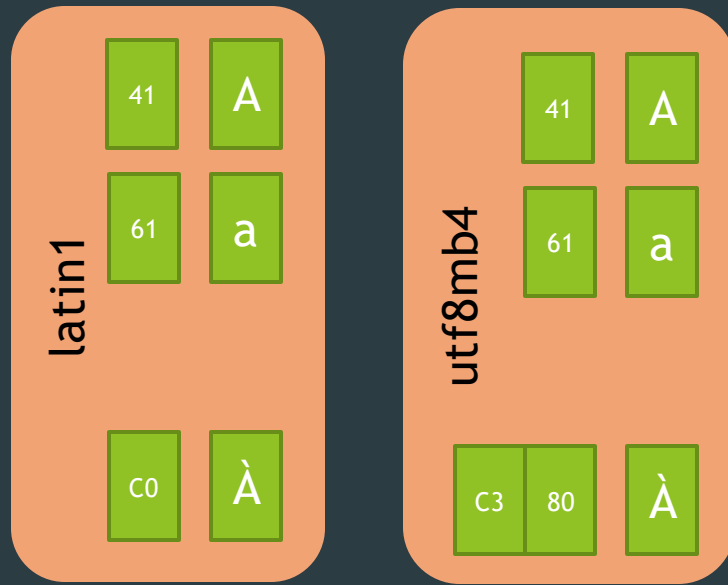
Strings are sequences of bytes with a charset and a collation

Strings are sequences of bytes with a charset and a collation

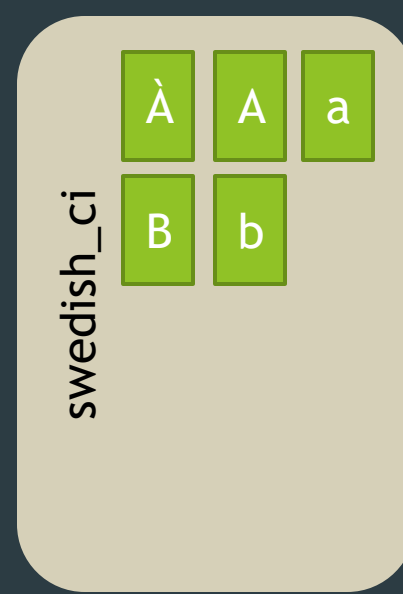
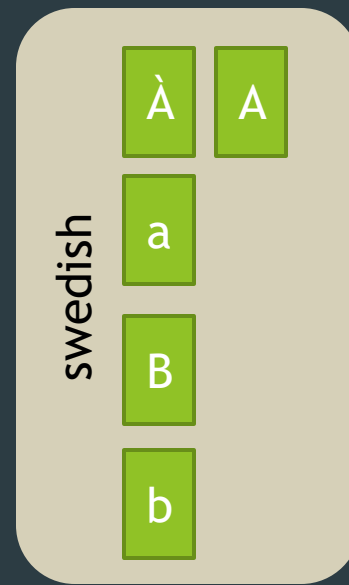
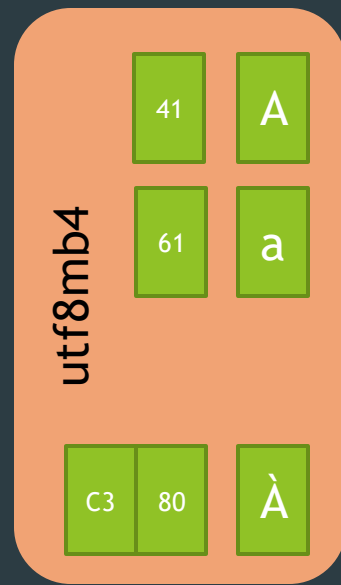

Strings are sequences of bytes with a charset and a collation



BINARY and VARBINARY



Strings are sequences of bytes with a charset and a collation



Strings are sequences of bytes with a charset and a collation

CHAR and VARCHAR

Strings are sequences of bytes with a charset and a collation

BINARY and VARBINARY

utf8mb4_swedish_ci

CHAR and VARCHAR

Strings are sequences of bytes with a charset and a collation

BINARY and VARBINARY

BINARY and CHAR: Fixed length

BINARY(10)

P	E	R	C	O	N	A			
---	---	---	---	---	---	---	--	--	--

CHAR(10) LATIN1

???

BINARY and CHAR: Fixed length

BINARY(10)

P E R C O N A

CHAR(10) LATIN1

O N L I N E

BINARY and CHAR: Fixed length

BINARY(10)

P	E	R	C	O	N	A			
---	---	---	---	---	---	---	--	--	--

CHAR(10) LATIN1

O	N	L	I	N	E				
---	---	---	---	---	---	--	--	--	--

CHAR(10) utf8 + InnoDB

2	0	2	0		☆	☆	☆	☆	☆
---	---	---	---	--	---	---	---	---	---

2	0	2	0						
---	---	---	---	--	--	--	--	--	--



Use `utf8mb4` by default

Don't use `utf8` (alias for `utf8mb3`)

CHAR and VARCHAR

Texts are sequences of bytes with a charset and a collation

BINARY and VARBINARY

✓
CHAR and VARCHAR

Texts are sequences of bytes with a charset and a collation

✓
BINARY and VARBINARY

VARBINARY and VARCHAR

VARBINARY(10)

7 P E R C O N A

VARCHAR(10) LATIN1

6 O N L I N E

VARCHAR(10) utf8mb4

$\frac{2}{0}$ 2 0 2 0 ☆ ☆ ☆ ☆ ☆

4 2 0 2 0

VARBINARY and VARCHAR

VARBINARY(10)

7 P E R C O N A

1 to 11
bytes

VARCHAR(10) L

6 O N L I N E

VARCHAR(10) utf8mb4

$\frac{2}{0}$ 2 0 2 0 ☆ ☆ ☆ ☆

1 to 41
bytes

4 2 0 2 0

VARBINARY(65535) and VARCHAR(65535)

VARCHAR(560) latin1



MySQL Max row size: 65K





I'll just go for VARCHAR(255) on all text columns



I'll just go for VARCHAR(255) on all text columns

“After all... I'll just consume the number of bytes + 1



All your VARCHAR(255) are now CHAR(255) in
memory temp tables

That's 255 bytes binary/latin-1. Or 1K
utf8mb4 O_o

Big texts



The Blob family

Text	Blob	Bytes Max
TEXT	BLOB	65K

<https://dev.mysql.com/doc/refman/8.0/en/blob.html>

The Blob family

Text	Blob	Bytes Max
TINYTEXT	TINYBLOB	255
TEXT	BLOB	65K
MEDIUMTEXT	MEDIUMBLOB	16M
LONGTEXT	LOB	4GB

<https://dev.mysql.com/doc/refman/8.0/en/blob.html>

The Blob family

Text	Blob	Bytes Max
TINYTEXT	TINYBLOB	255
TEXT	BLOB	65K
MEDIUMTEXT	MEDIUMBLOB	16M
LONGTEXT	LOB	4GB

<https://dev.mysql.com/doc/refman/8.0/en/blob.html>

The Blob family

Text	Blob	Bytes Max
TINYTEXT	TINYBLOB	255
TEXT	BLOB	65K
MEDIUMTEXT	MEDIUMBLOB	16M
LONGTEXT	LOB	4GB



Don't
use



FS /
Object
Store

<https://dev.mysql.com/doc/refman/8.0/en/blob.html>

IF a `SELECT` references a `BLOB/TEXT` column

Temporary tables go **DIRECTLY** to **DISK**



Small sets



ENUM

One value out of a set of possibilities ('big', 'small')

1 or 2 bytes

Looks like a string to the client

SET

Choose a set of non-exclusive possible values ('pool', 'terrace', 'fence', 'guard')

SETS are NOT good for finding stuff

FIND_IN_SET is a function. No indexes

Model by default to a separate table + relation

Date and Time



The Date and Time Family

Text	Bytes
YEAR	1
DATE	3
TIME	3 + ...
DATETIME	8 + ...
TIMESTAMP	4 + ...

Careful with ranges. Consult the manual frequently

<https://dev.mysql.com/doc/refman/8.0/en/date-and-time-types.html>

The Date and Time Family

Text	Bytes
YEAR	1
DATE	3
TIME	3 + ...
DATETIME	8 + ...
TIMESTAMP	4 + ...




An epoch!
Things
that
happen
"now"

Careful with ranges. Consult the manual frequently.

<https://dev.mysql.com/doc/refman/8.0/en/date-and-time-types.html>

NULL vs NOT NULL



Gives the DB hints

Set NULL where ever it makes sense

Masked Data

**By day,
she was
Hello
Kitty**

**By night,
she
dropped
the "o"**

Choose type for

An IP Address

234.34.123.92

Choose type for

An IP Address

234.34.123.92

CHAR(15)? VARCHAR(15)?

An IPv4 address (dotted-decimal notation)

172 . 16 . 254 . 1



10101100 . 00010000 . 11111110 . 00000001



Thirty-two bits (4 x 8), or 4 bytes

INT



```
INSERT INTO table (col) VALUES (INET_ATON('10.10.10.10'));
```

```
SELECT INET_NTOA(col) FROM table;
```


Choose type for

MD5("The quick brown fox jumps over the lazy cat")

"71bd588d5ad9b6abe87b831b45f8fa95"

Choose type for

MD5("The quick brown fox jumps over the lazy cat")

"71bd588d5ad9b6abe87b831b45f8fa95"

CHAR(32)?

MD5

From Wikipedia, the free encyclopedia

The **MD5 message-digest algorithm** is a widely used [cryptographic hash function](#) producing a [128-bit \(16-byte\) hash value](#), typically expressed in text format as a 32 digit [hexadecimal](#) number. MD5 has been utilized in a wide variety of cryptographic applications, and is also commonly used to verify [data integrity](#).

BINARY (16)



Others

UUIDs

HASH functions

Network masks



Your new best friend

- ▶ <https://dev.mysql.com/doc/refman/5.7/en/storage-requirements.html>
- ▶ <https://dev.mysql.com/doc/refman/8.0/en/storage-requirements.html>

Smaller is better



VS



Indexes



Datatype for keys



Default to an Integer for your PK

AUTONUMERIC UNSIGNED INT by default. Think if it can be smaller.
TINY, SMALL, MEDIUM. Always UNSIGNED!

UNIQUE INDEX your "natural keys"
This is called a surrogate key

InnoDB stores whole PK on every index.
The DB bloats with "big keys"

Do you really need an UNSIGNED BIGINT?

The magnitude comparison trick

Do you really need an UNSIGNED BIGINT?

The magnitude comparison trick

Epochs in Unix have 4 bytes (like an UNSIGNED INT)

Do you really need an UNSIGNED BIGINT?

The magnitude comparison trick

Epochs in Unix have 4 bytes (like an UNSIGNED INT)

The epoch has been counting every second since 1970

Do you really need an UNSIGNED BIGINT?

The magnitude comparison trick

Epochs in Unix have 4 bytes (like an UNSIGNED INT)

The epoch has been counting every second since 1970

Will run out in 2038

Do you really need an UNSIGNED BIGINT?

The magnitude comparison trick

Epochs in Unix have 4 bytes (like an UNSIGNED INT)

The epoch has been counting every second since 1970

Will run out in 2038

An INT can identify ONE THING HAPPENING EVERY SECOND for 68 YEARS!

Do you STILL need a BIGINT?

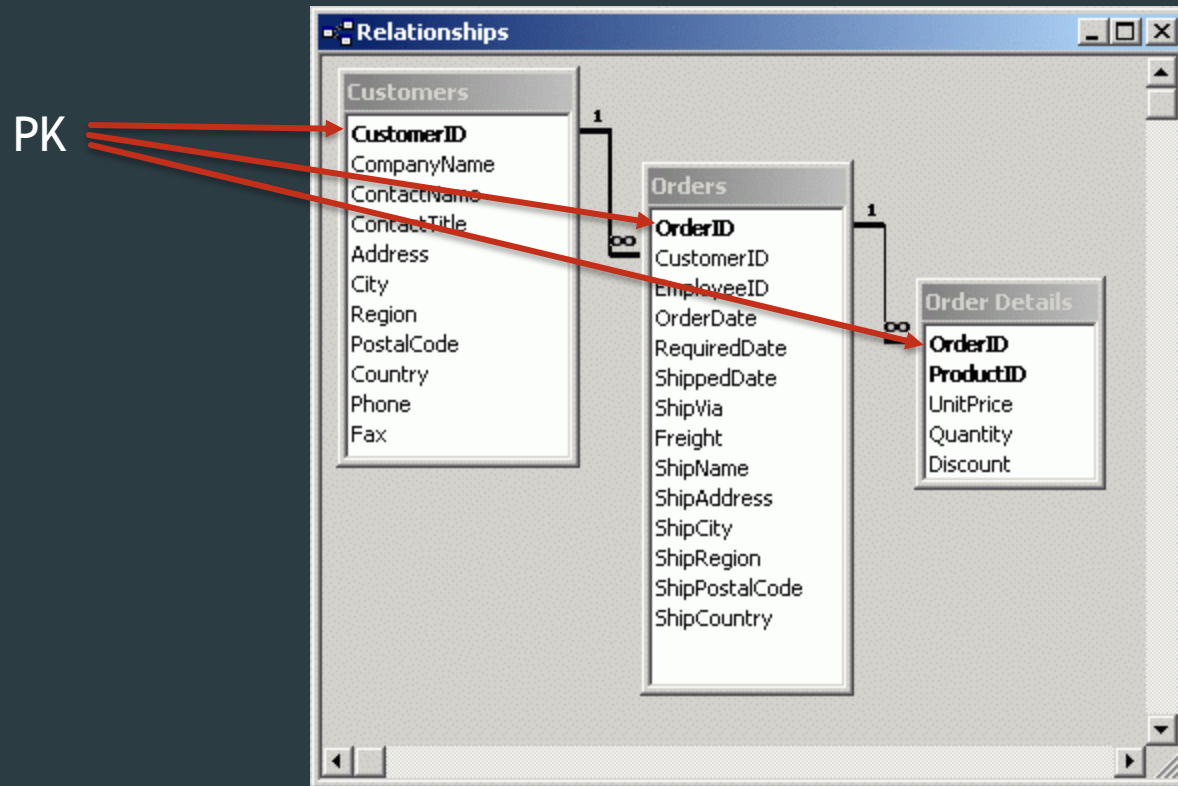
FIND THE CAT



FIND THE CAT

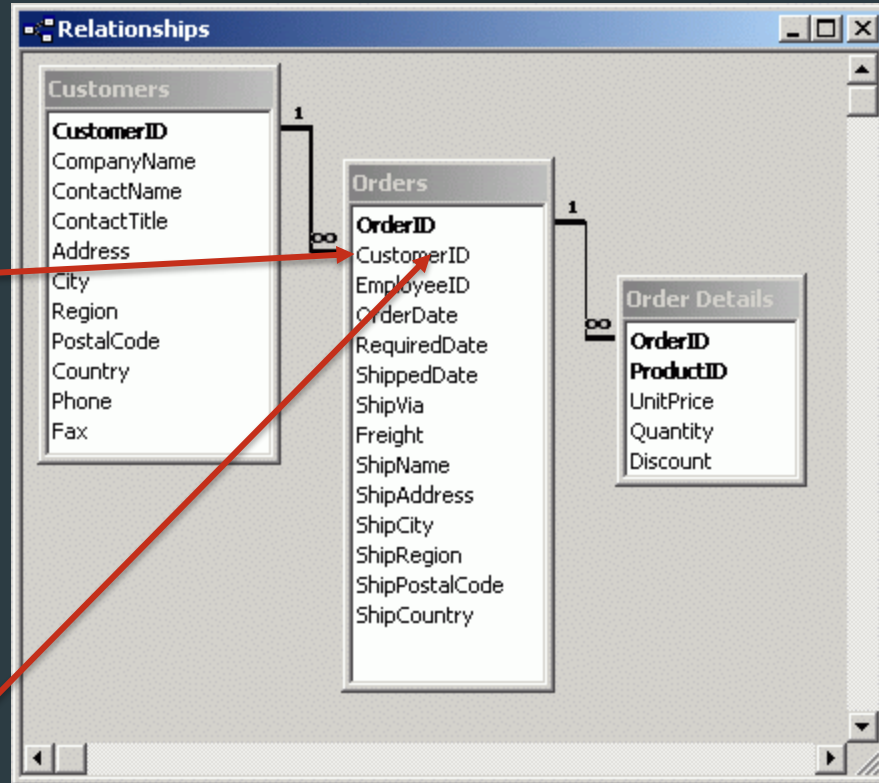


Index the two sides of relations



Index the two sides of relations

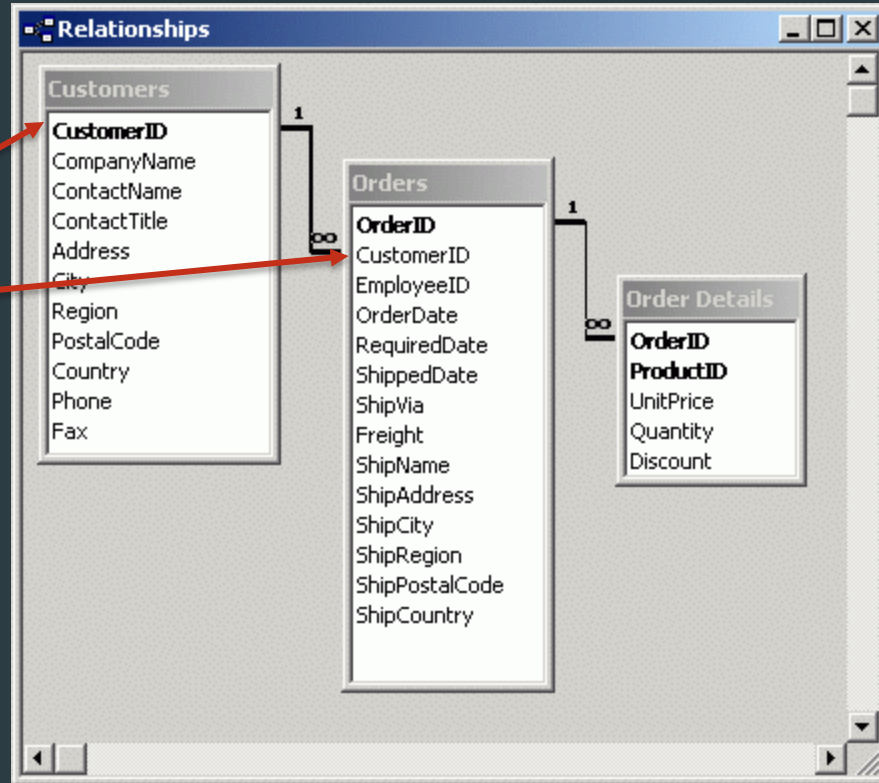
Non-Unique index



Use REFERENCES Customers.CustomerID

Index the two sides of relations

Same Data Type



FIND THE CAT



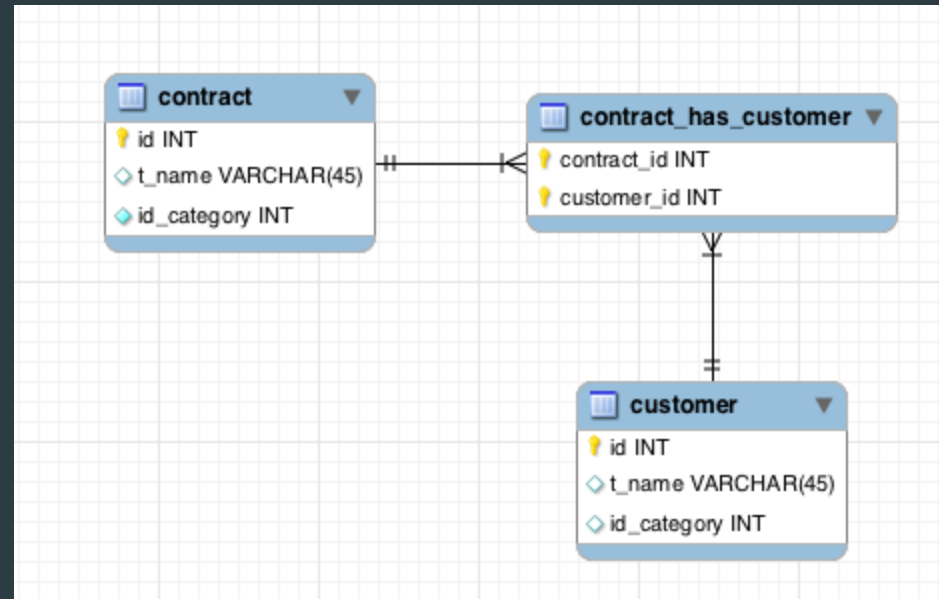
kappit.com



N-M relation: Junction tables

PK is (contract_id, customer_id)
(Implied uniqueness)

Index “both ways”:
(customer_id, contract_id)



InnoDB optimization: Don't index the full (customer_id, contract_id). The index ALREADY HAS customer_id in it's leafs. So just index (customer_id)



Find the cat.



Find the cat.

Don't operate on fields

- ▶ Because they can't use indexes
- ▶ WHERE column = 'x'
WHERE column > 2000
WHERE column LIKE 'prefix%'
- ▶ WHERE column + 2000 > 2013
WHERE FIND_IN_SET(column)
WHERE CONCAT(f1,f2) = "xxxx.com"
WHERE YEAR(date) = 2015
WHERE column LIKE '%.com'

Polish your maths: Algebra

- ▶ Doesn't use index

WHERE column + 2000 > 2013

WHERE FIND_IN_SET('pool',column)

WHERE CONCAT(f1,'.',f2) =
"xxxx.com"

WHERE YEAR(date) = 2015

WHERE column LIKE '%.com'

- ▶ Uses index

WHERE column > 13

?

WHERE f1 = 'xxxx' AND f2 = 'com'

WHERE date BETWEEN '01-01-2015'
and '31-12-2015'

?

The old switcheroo...

▶ Doesn't use index

WHERE column + 2000 > 2013

WHERE FIND_IN_SET('pool',column)

WHERE CONCAT(f1,'.',f2) =
"xxxx.com"

WHERE YEAR(date) = 2015

WHERE column LIKE '%.com'

▶ Uses index

WHERE has_pool = 1

WHERE column_rev LIKE 'moc.%'

The old switcheroo...

- ▶ Doesn't use index

```
WHERE column + 2000 > 2013
```

```
WHERE FIND_IN_SET('pool',column)
```

```
WHERE CONCAT(f1,'.',f2) =  
"xxxx.com"
```

```
WHERE YEAR(date) = 2015
```

```
WHERE column LIKE '%.com'
```

- ▶ Uses index

```
WHERE has_pool = 1
```

```
UPDATE t SET has_pool =  
FIND_IN_SET('pool',column);
```

```
WHERE column_rev LIKE 'moc.%%'
```

```
UPDATE t SET  
column_rev=REVERSE(column)
```





CAT IS HERE


Wrap up

▶ Data Types

- ▶ How to select them
- ▶ Masked types

▶ Indexes

- ▶ Surrogate keys
- ▶ Indexing relations
- ▶ Using indexes



[pplu_io](#)



<https://www.linkedin.com/in/josluismartineztorres/>