



What Do You Mean, SQL Syntax Error?

Presented by:
Sheeri K. Cabral

At ODTUG Kaleidoscope 2010

Pythian
love your data

Who I Am

- MySQL DBA
- MySQL User Group
- First Oracle ACE Director for MySQL
- Lots of community stuff (videos, blog, podcast on hiatus)

Standards are great....

-that's why there are so many of them!
- Compare to ANSI/ISO SQL:2003
- Does not cover implementation details

In addition to all this....

- Storage engines
- Data types
- Index types
- Stored code uses extended SQL syntax
 - Not pl/sql, t-sql
 - 8:30 am on Wed

MySQL

- Fast
- Reliable
- Easy to use

```
SELECT TABLE_NAME  
FROM INFORMATION_SCHEMA.TABLES  
WHERE TABLE_SCHEMA IN 'sakila';
```

```
SHOW TABLES FROM sakila;
```

Easy to use

- Usually not what you want
 - Friendly by default
 - Invalid data gets truncated
 - Warning, not error
 - Implicit conversion of data
 - Can change with `sql_mode`

Case Sensitivity

- SQL reserved words are case-insensitive
 - Typically all caps (SELECT, FROM, NULL, etc)
- Most fields, indexes, names, case-insensitive

Case Sensitivity

- Determined by filesystem:
 - Databases
 - Tablespaces
 - Tables
 - Views
 - Aliases
 - Triggers
 - Log file groups
- `lower_case_table_names`

Comments

```
-- single-line comment
```

```
#no space character required
```

```
/* comment that may span  
multiple lines*/
```

Portability

- `/*! MySQL-specific */`

```
/* SHOW DATABASES; */
```

- `/*!VERSION Version-specific */`

```
SHOW /*!50002 GLOBAL */ STATUS
```

Escaping in Strings

- Backslash (\)
 - \\
 - \' and \"
 - _ and \% when using LIKE
 - \r, \n, \t
 - \0 and \Z
 - \b

Identifiers

- Databases
- Tables, views, aliases
- Fields
- Indexes
- Tablespaces
- Stored routines, triggers, events
- Servers
- Log file groups

Identifiers

- 64 chars
 - Except aliases, 255 chars
- May not end in spaces
- Quote:
 - Reserved
 - Numbers
 - punctuation

Quoting

- <http://bit.ly/mysqlquoting>
- ` for identifiers
- ' or " for strings

Dot notation

- **SELECT db.tbl.fld1 FROM [db.]tbl;**

```
use test;
```

```
SELECT tables.table_name FROM  
information_schema.tables WHERE  
table_schema='test';
```

- **Or alias instead of tbl**

```
SELECT information_schema.s.table_name FROM  
information_schema.tables AS s WHERE  
table_schema='test';
```

Aliasing

- AS
- Or not

```
SELECT s.table_name FROM tables s WHERE  
table_schema='test';
```


Character sets / Collations

- charset = alphabet
 - default: latin1
- collation = sort order
 - default: latin1_swedish_ci
 - _cs
 - _bin
 - Depends on charset
 - All ignore trailing whitespace when sorting

Charset / collation

- Server
- Database
- Table
- Field
- Strings
 - `character_set_client`
 - `character_set_connection`
 - `character_set_results`

String comparison

- Case-insensitive by default
 - Unless binary/case-sensitive collation

- LIKE can compare numbers

```
mysql> SELECT 1 LIKE 1, 1 LIKE 0, '1A' LIKE '1%';
+-----+-----+-----+
| 1 LIKE 1 | 1 LIKE 0 | '1A' LIKE '1%' |
+-----+-----+-----+
|          1 |          0 |                  1 |
+-----+-----+-----+
1 row in set (0.00 sec)
```

No Support

- CREATE TABLE
 - GLOBAL/LOCAL
 - ON COMMIT
 - REF IS
 - UNDER
 - AS SUBQUERY
- CREATE/ALTER/DROP SCHEMA
 - AUTHORIZATION

No Support

- Catalogs
 - CATALOG fields in I_S are NULL
 - SET CATALOG not supported
- Querying data samples with TABLESAMPLE
- CREATE/DROP ASSERTION
- GROUP BY
 - CUBE
 - GROUPING SETS

No Support

- Binary set functions
 - CORR, COVAR_POP, COVAR_SAMP
- COLLECT, FUSION, INTERSECTION
- Windowing functions
 - RANK, DENSE_RANK, PERCENT_RANK, CUME_DIST
- Domain/Domain constraints
 - CREATE/ALTER/DROP DOMAIN
 - DOMAIN permissions with GRANT

No Support

- Sequences
 - CREATE/ALTER/DROP SEQUENCE
 - Do have per-table AUTO_INCREMENT
 - SEQUENCE permissions with GRANT
- In field definitions
 - GENERATED
 - ALWAYS
 - BY DEFAULT
 - AS IDENTITY

No Support

- Transliterations
 - CREATE/DROP TRANSLATION
 - TRANSLATION permission with GRANT
- DECLARE for cursors only
 - No embedded declaration
 - No embedded SQL
 - No embedded authorization declaration
 - No temporary table declarations

No Support

- Descriptor areas
 - ALLOCATE/DEALLOCATE DESCRIPTOR
 - GET/SET DESCRIPTOR
- Recursive queries
 - SEARCH DEPTH/BREADTH FIRST BY
 - CYCLE
- Updatable cursors
- WHERE CURRENT OF in UPDATE/DELETE

No Support

- Connection management
 - CONNECT TO
 - SET CONNECTION
 - DISCONNECT
- Session management
 - SET ROLE
 - SET TIME ZONE
 - SET SESSION is for session variables

No Support

- SET SCHEMA
 - In mysql, use or \u
 - Client option when connecting
- GET DIAGNOSTICS
 - Monitor error log
 - SHOW ERRORS
 - SHOW WARNINGS

DECLARE

- DECLARE for cursors only
- No embedded:
 - declaration, SQL, authorization declaration
- Not for temporary table declarations
 - CREATE/DROP TEMPORARY TABLE
- DECLARE extended
 - Variables, conditions, handlers
 - 8:30 am Wed

User-defined types

- Not supported
- So no:
 - TYPE permissions for GRANT
 - CREATE/DROP CAST
 - CREATE/DROP ORDERING FOR
 - CREATE/ALTER/DROP TYPE
 - CREATE/ALTER/DROP TRANSFORM

Not Fully Supported

- CREATE TABLE LIKE
 - tbl_name only
 - Not INCLUDING/EXCLUDING DEFAULTS
 - Not INCLUDING/EXCLUDING IDENTITY
- User-defined, only in C
 - metadata (system view)
 - functions

Cursors

- All are dynamic
 - prepared at runtime
 - Not fully dynamic
 - Cached in temp tables
- Performance usually worse than other DBMS

Foreign Key Constraints

- Syntax accepted
 - Only tables using transactional storage engines implement foreign key checking
 - Others disregard, no errors
- Invalid: REFERENCES ARE [NOT] CHECKED
 - SET SESSION FOREIGN_KEY_CHECKS=0
 - No way to disable individual constraints
- MATCH clause disregarded

Prepared Statements

- Local scope
 - Each session compiles its own
 - Not a big performance boost
- Basic PREPARE
 - Only statement + name
 - No ATTRIBUTES

Prepared Statements

- No dynamically prepared statements
 - No IMMEDIATE keyword
- No DESCRIBE for input/output parameters

DESCRIBE

- Translates to SHOW COLUMNS FROM

```
DESCRIBE tblname;
```

```
SHOW COLUMNS FROM tblname;
```

```
DESCRIBE tblname fldname;
```

```
SHOW COLUMNS FROM tblname WHERE  
Field='fldname';
```

```
DESCRIBE tblname 'fldname';
```

```
SHOW COLUMNS FROM tblname WHERE Field  
LIKE 'fldname';
```

String Manipulation

- `CAST()`, `CONVERT()` to any data type
 - Cannot convert to `REAL`, `BIGINT`
- `TRIM()`
 - Standard = 1 leading/trailing char
 - MySQL = multiple
- `||` is OR, not `CONCAT()`
- `CONCAT()` can take >2 values

GRANT/REVOKE syntax

- No WITH ADMIN OPTION
 - SUPER privilege
 - WITH GRANT OPTION
- No WITH HEIRARCHY OPTION
- No GRANTED BY
- No CHARACTER SET/COLLATION privileges
- Roles not supported
 - DROP ROLE

Users

- A user is unique to user@host
 - multiple hosts per username
 - wildcard, localhost
 - Fields cannot be associated with a user
- RENAME USER for one or more renames
- If no privileges granted, USAGE ON *.*
- REVOKE ALL...does not revoke USAGE
- DROP USER revokes all including USAGE

Transaction Support

- Only in supported storage engines
- SET [LOCAL] TRANSACTION not supported
- Transaction modes
 - Not args to START TRANSACTION
 - Isolation level can be set
 - Server can be made read_only
 - READ WRITE by setting user permissions
 - No equivalent to DIAGNOSTICS SIZE

Check Constraints

- Constraints by:
 - data type, foreign/primary/unique keys
- Other check constraints not supported
 - No SET CONSTRAINTS
 - CONSTRAINT....CHECK allowed, ignored
- Updatable views can have check constraints
 - Wed 8:30 am

Upsert Statements

- MERGE unsupported
- Can have upsert on one condition:
 - Primary/unique key exists

Upsert Statements

- **INSERT...ON DUPLICATE KEY UPDATE**
 - INSERT/UPDATE triggers fire
- **REPLACE**
 - If conflict, DELETE current record, INSERT new
 - INSERT trigger fires, DELETE fires if used
- **INSERT IGNORE**
 - If conflict, keep current record
 - No triggers fire

Extensions

Aliases

- START TRANSACTION
 - BEGIN, BEGIN WORK
- DEALLOCATE PREPARE stmt_prep
 - DROP PREPARE stmt_prep
- MOD(num1,num2)
 - num1 % num2

ALTER TABLE

- Many offline statements (require table rebuild)
- Online statements:
 - ADD/DROP INDEX for variable-width fields
 - Renaming a field with CHANGE COLUMN
 - Changing a field's default value
 - Adding items to the end of an ENUM/SET

ALTER TABLE

- Multiple actions
 - Comma-separated list
- ADD INDEX
 - ADD FULLTEXT INDEX
 - ADD SPATIAL INDEX

```
CHANGE COLUMN fld_name  
new_fld_name  
new_fld_definition [FIRST/AFTER]
```

ALTER TABLE

CONVERT TO CHARACTER SET

```
charset_name [COLLATION coll_name]
```

- **DISABLE/ENABLE KEYS**

- **IGNORE**

- You will lose duplicate rows

```
ALTER IGNORE TABLE foo ADD UNIQUE  
KEY(non_unique_field);
```

ALTER TABLE

MODIFY COLUMN fld_name new_fld_definition

– CHANGE COLUMN TO change name,order

- ORDER BY fld_list
 - Rebuilds table with new order
 - New inserts/updates do not follow this order
- RENAME [TO] new_tbl

CREATE

- CREATE *object* IF NOT EXISTS obj_name
 - DATABASE
 - TABLE
 - EVENT
- CREATE OR REPLACE VIEW
- CREATE INDEX
 - Named index
 - Index storage (USING HASH)
 - Column prefixes

Table Definition Extensions

- For use with CREATE/ALTER TABLE

<http://dev.mysql.com/doc/refman/5.1/en/create-table.html>

- CHARACTER SET
- COLLATE
- COMMENT='comment string'
 - Also for fields and indexes
- ENGINE=storage_engine_name

Table Definition Extensions

- AUTO_INCREMENT
- KEY_BLOCK_SIZE
 - May or may not be used by storage engine
 - Default is 0, storage engine uses its own default
- KEY_BLOCK_SIZE
- Partitioning commands
- ROW_FORMAT

MyISAM Table Definition Extensions

- DATA DIRECTORY='/path/to/dir'
- INDEX DIRECTORY='/path/to/dir'
- PACK_KEYS
- For MyISAM tables over 256 Tb
 - AVG_ROW_LENGTH
 - MAX_ROWS
- DELAY_KEY_WRITE
 - 11 am Tuesday (tomorrow)

Other Table Definition Extensions

- CONNECTION='connection string'
 - For FEDERATED tables
- INSERT_METHOD
 - MERGE table
 - FIRST, LAST, NO
- MIN_ROWS (MEMORY)

DROP

- DROP *object* IF EXISTS obj_name
 - DATABASE
 - TABLE
 - EVENT
 - PROCEDURE, FUNCTION, TRIGGER, VIEW
 - SERVER
- DROP TABLE tbl1, tbl2;
- DROP INDEX

DML

- INSERT, REPLACE, UPDATE, DELETE
- LOW_PRIORITY
- IGNORE
 - Not REPLACE
- LIMIT x
 - DELETE, UPDATE
- ORDER BY
 - DELETE, UPDATE
 - Usually used with LIMIT

INSERT

- Readability

```
INSERT INTO tbl SET fld1=val1,  
fld2=val2, fld3=val3;
```

- Extended insert

```
INSERT INTO tbl (fld1, fld2, fld3...)  
VALUES (val1, val2, val3),  
(val4, val5, val6);
```


Removing data

```
TRUNCATE [TABLE] tbl_name;
```

- **Multi-table DELETE syntaxes:**

```
DELETE tbl1,tbl2 FROM tbl_expr [WHERE ...]
```

```
DELETE FROM tbl1,tbl2 USING tbl_expr [WHERE ...]
```

- **tbl_list = comma-separated list**
- **tbl_expr = any expr that returns a table**
 - JOIN, subquery, etc

Loading data

- 11 am Tuesday (tomorrow)
- Batch inserts with INSERT DELAYED
 - MyISAM, ARCHIVE, BLACKHOLE, MEMORY
- LOAD DATA INFILE
- LOAD XML INFILE

Exporting data

- 11 am Tuesday (tomorrow)
- `SELECT ... INTO OUTFILE`
- `SELECT ... INTO DUMPFILE`

LIMIT

- Already seen for UPDATE, DELETE
- Can be used in SELECT too
- LIMIT num
- LIMIT offset, num
- SELECT SQL_CALC_FOUND_ROWS ...

SQL_CALC_FOUND_ROWS

- Use `SELECT FOUND_ROWS()` to get count

```
mysql> SELECT SQL_CALC_FOUND_ROWS
rental_date, inventory_id, customer_id
FROM sakila.rental LIMIT 1\G
***** 1. row *****
  rental_date: 2005-05-24 22:53:30
inventory_id: 367
  customer_id: 130
1 row in set (0.03 sec)
```

```
mysql> SELECT FOUND_ROWS() \G
***** 1. row *****
FOUND_ROWS(): 16044
1 row in set (0.00 sec)
```

Locking

- **SELECT....LOCK IN SHARE MODE**
 - Read (shared) lock on rows
- **SELECT...FOR UPDATE**
 - Write (exclusive) lock on rows

Intermediate Temporary Table

- Use in-memory temporary table for results
 - `SELECT SQL_SMALL_RESULT fld1, fld2 FROM...`
- Use on-disk temporary table for results
 - `SELECT SQL_BIG_RESULT fld1, fld2 FROM...`

Memory

- Similarly for any SELECT statement...
 - SELECT SQL_BUFFER_RESULT ...
 - sql_buffer_result
- Query cache
 - SELECT SQL_NO_CACHE fld1,fld2 FROM....
 - SELECT SQL_CACHE fld1,fld2 FROM....

GROUP BY

- WITH ROLLUP

```
mysql> SELECT table_schema,engine,count(*) FROM
information_schema.tables
WHERE engine IS NOT NULL and table_schema NOT IN
('information_schema','mysql')
GROUP BY engine,table_schema WITH ROLLUP;
```

table_schema	engine	count(*)	
sakila	InnoDB	15	
test	InnoDB	2	
NULL	InnoDB	17	<-- InnoDB total
sakila	MyISAM	1	
test	MyISAM	2	
NULL	MyISAM	3	<-- MyISAM Total
NULL	NULL	20	<-- Total

7 rows in set (0.00 sec)

DO

- Performs query without returning result

```
mysql> SELECT SLEEP(5);
```

```
+-----+
```

```
| SLEEP(5) |
```

```
+-----+
```

```
|          0 |
```

```
+-----+
```

```
1 row in set (5.00 sec)
```

```
mysql> DO SLEEP(5);
```

```
Query OK, 0 rows affected (5.01 sec)
```

Other SELECT Extensions

- GROUP BY fld1 ASC
 - default
- GROUP BY fld1 DESC
- UNION ... ORDER BY
- DISTINCTROW
 - Alias of DISTINCT

Other SELECT Extensions

- **SELECT HIGH_PRIORITY**
 - low_priority_updates, SELECT LOW_PRIORITY
- **PROCEDURE ANALYSE()**
 - <http://dev.mysql.com/doc/refman/5.1/en/procedure-analyse.html>
- **EXPLAIN SELECT**
 - To get an idea of the query plan
 - <http://technocation.org/files/doc/EXPLAIN.pdf>

Variables

```
mysql> SET @num:=100;  
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> SELECT @num, @num:=@num+50;
```

```
+-----+-----+  
| @num | @num:=@num+50 |  
+-----+-----+  
| 100  | 150  |  
+-----+-----+
```

```
1 row in set (0.00 sec)
```

- Can use for running totals/averages

Other extensions

- FLUSH
- KILL
- Index caching
- Replication commands
- SHOW
 - Mostly metadata in INFORMATION_SCHEMA
 - <http://dev.mysql.com/doc/refman/5.1/en/show.html>

Table Maintenance Extensions

- ANALYZE TABLE
- OPTIMIZE TABLE
- CHECK TABLE
- REPAIR TABLE
- CHECKSUM TABLE

Questions, Comments, Feedback?

- Get the 80-page PDF supplement to this presentation!
- Win a copy of the MySQL Administrator's Bible