© SkySQL Ab. Commercial in Confidence

Die CONNECT Storage Engine für MySQL

DOAG Conference 2013

Nürnberg

© SkySQL Ab. Commercial in Confidence

Who am I?

- Ralf Gebhardt
- Principal Sales Engineer @ SkySQL
- Joined MySQL GmbH in 2002
- Worked for MySQL@Sun and MySQL@Oracle until July 2011

Now we are one company

© SkySQL Ab. Commercial in Confidence

Agenda

- Plugin Architecture
- CONNECT Storage Engine
  - Overview
  - Architectures
  - Features
  - Table Types and Options
- Examples

© SkySQL Ab. Commercial in Confidence

The MySQL Plugin Architecture

- Plugin Architecture is a major differentiator of MySQL
- Datastores can interact with the MySQL SQL layer
- Allow advanced interaction
- Specific Create Table parameters (MariaDB)
- Auto-discovery of table structure (MariaDB)
- Condition push down
- Allow join with other storage engines
  - InnoDB / MyISAM tables

© SkySQL Ab. Commercial in Confidence

The CONNECT Storage Engine

- What is the CONNECT storage engine?
  - CONNECT is a storage engine that enables MariaDB to use external data as they were standard tables in the server
  - Data is not loaded into MariaDB
- History of the CONNECT storage engine
  - The engine has been mainly developed by Olivier Bertrand, an ex IBM database researcher, with the intent to have a more versatile way to access external data sources in various formats
  - The idea dates back in 2004 and Olivier has been in touch with MySQL and MariaDB since
  - Today CONNECT is a standard engine in MariaDB 10

© SkySQL Ab. Commercial in Confidence

The CONNECT Storage Engine

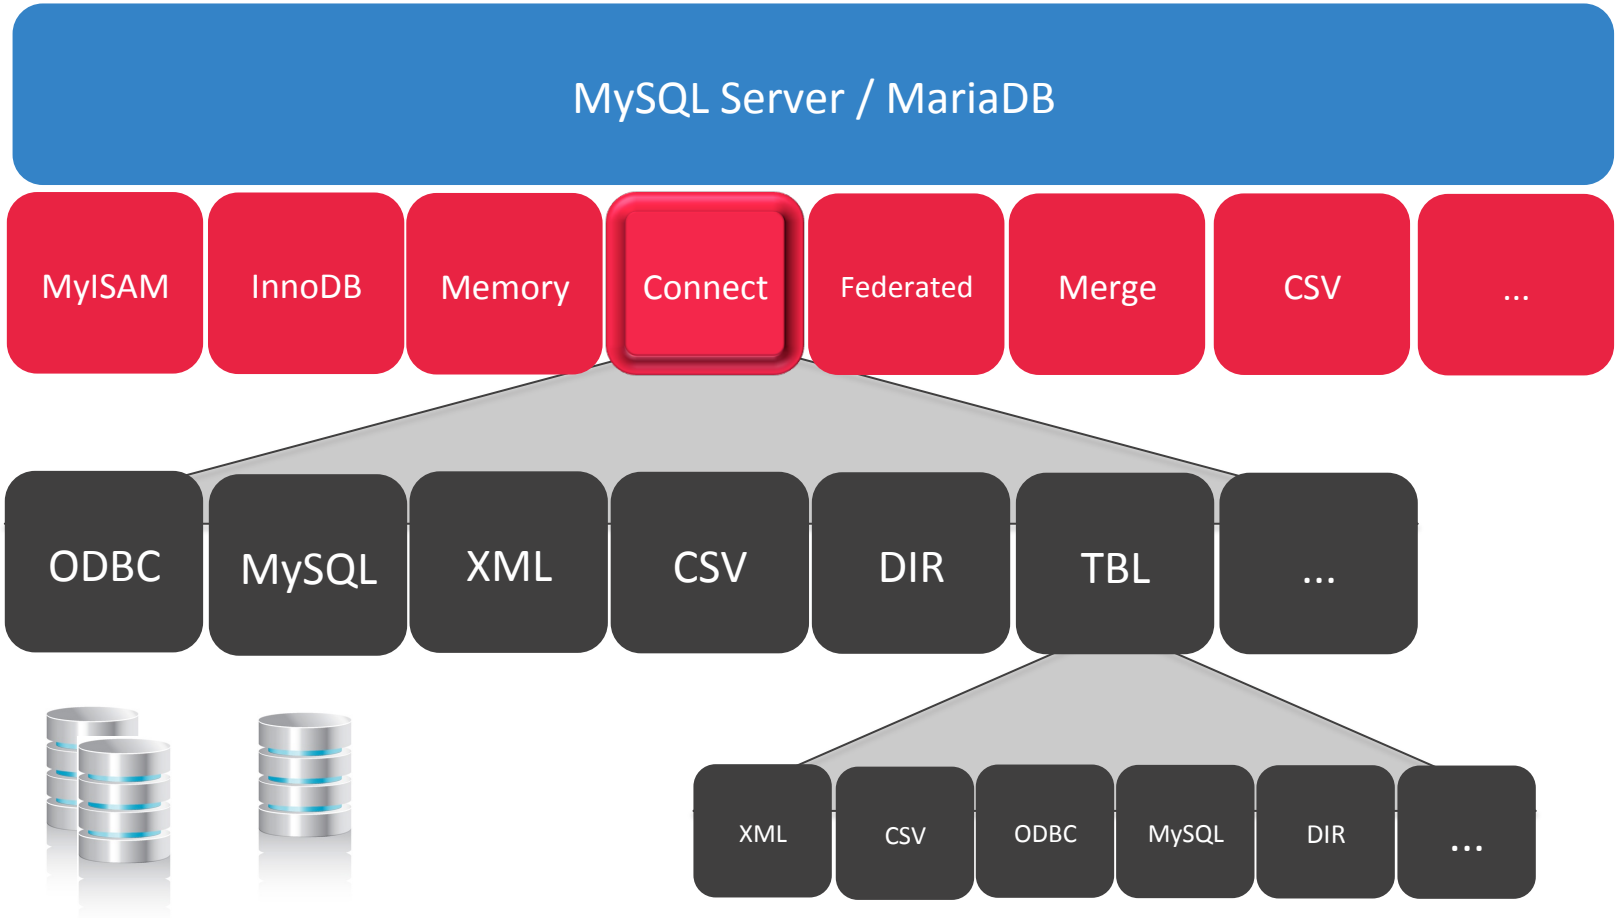

© SkySQL Ab. Commercial in Confidence

CONNECT Engine Usage

- Integrates/access data directly in many non-MariaDB formats
- Simplifies the ETL procedures in Business Intelligence and Business Analytics
- Simplifies the export/import of data from/to MariaDB, to/from other data sources
- There are strong similarities with the CSV engine and some similarities with FederatedX and Merge
- The CSV engine does not allow indexing and block reading
- The FederateX engine imposes several restrictions and performance limitations
- FILE privilege is required

© SkySQL Ab. Commercial in Confidence

CONNECT Engine - Advantages

- Multiple File Table (option multiple=[0|1|2]) processes sequentially files of the same type
- Indexing for most of the table types
- Data compression
- Block reading boosts read performance
- Condition push down
  - Valid for ODBC, MYSQL, TBL and WMI
  - set optimizer_switch='engine_condition_pushdown=on‘
    ...or...
    --engine_condition_pushdown=on
- Even more possibilities with the OEM file type

© SkySQL Ab. Commercial in Confidence

CONNECT Engine - Features

- Table “auto-creation” when the the file does not exist or it is not specified
- Large tables support (>2GB)
  - Available for FIX, BIN and VEC
  - Use ‘option_list’=‘huge=1’
- Compression - gzlib format
  - Available for DOS, FIX, BIN, CSV and FMT
- ODBC Format
  - Access to ODBC and UnixODBC data sources
  - WHERE conditions are push to the ODBC source
  - Multiple tables/files from data sources can be consolidated into a single table

© SkySQL Ab. Commercial in Confidence

CONNECT Engine - Features

- MYSQL Format - to access local or remote MySQL tables
  - Allows to define a subset of the source columns and type conversion
  - Condition LIMIT push down
- TBL - Table List Table
  - Collection of tables seen as one
  - No limitation on the storage engine: tables can be from different storage engines (including CONNECT)
  - The tables may have different column structure
- VEC - Column Store
  - Data is stored in binary files as vectors
  - I/O optimization, as CONNECT reads only columns that are requested by the query
  - The “split format” allows to split columns in multiple files
- XML - for XML and HTML data
  - HTML attributes can be defined in the option_list

© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Install Plugin

- Install the plugin

```
INSTALL PLUGIN CONNECT SONAME 'ha_connect';
```

Note: The library may not be in the rpm on RH/ CentOS, you can find it in the standard tarball

© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Create a Table

- Standard create statement specifying `engine=CONNECT` to select engine
- CONNECT table type defined via table option `table_type=`
- More connect table options exist, some are defined `option_list=`

```
create table essai (
  num integer(4) not null,
  line char(15) not null)
engine=CONNECT table_type=MYSQL dbname=test tabname=people
option_list='user=root,host=localhost';
```


© SkySQL Ab. Commercial in Confidence

CONNECT Table Types

Type	Description
DOS	The table is contained in one or several files. The file format can be refined by some other options of the command or more often using a specific type as many of those described below. Otherwise, it is a flat text file where columns are placed at a fixed offset within each record, the last column being of variable length.
FIX	Text file arranged like DOS but with fixed length records.
BIN	Binary file with numeric values in platform representation, also with columns at fixed offset within records.
VEC	Binary file organized in vectors, in which column values are grouped consecutively, either split in separate files or in a unique file.
DBF*	File having the dBASE format.
CSV*	"Comma Separated Values" file in which each variable length record contains column values separated by a specific character (defaulting to the comma)
FMT	File in which each record contains the column values in a non-standard format (the same for each record) This format is specified in the column definition.
INI	File having the format of the initialization or configuration files used by many applications.
XML	File having the XML or HTML format.
DIR	Virtual table that returns a file list like the Unix ls or DOS dir command.

(*) Auto discovery of the table structure

© SkySQL Ab. Commercial in Confidence

CONNECT Table Types

Type	Description
ODBC*	Table extracted from an application accessible via ODBC or unixODBC. For example from another DBMS or from an Excel spreadsheet.
MYSQL*	Table accessed using the MySQL API like the FEDERATED engine.
PROXY*	A table based on another table existing on the current server.
TBL*	Accessing a collection of tables as one table (like the MERGE engine does for MyIsam tables)
XCOL*	A table based on another table existing on the current server with one of its column containing of comma separated values.
OCCUR	A table based on another table existing on the current server, several columns of the object table containing values that can be grouped in only one column.
WMI*	Windows Management Instrumentation table displaying information coming from a WMI provider. This type enables to get in tabular format all sort of information about the machine hardware and operating system (Windows only).
MAC	Virtual table returning information about the machine and network cards (Windows only).
OEM	Table of any other formats not directly handled by CONNECT but whose access is implemented by an external plugin module (DLL or Shared Library).

(*) Auto discovery of the table structure

© SkySQL Ab. Commercial in Confidence

CONNECT Table Options

Option	Description
TABLE_TYPE	The type of the external table : DOS, FIX, BIN…. Defaults to DOS.
FILE_NAME	The file (path) name for all table types based on files. Can be absolute or relative to the current data directory.
XFILE_NAME	The file (path) base name for a table index files. Can be absolute or relative to the data directory. Defaults to the file name.
TABNAME	The target table or node for ODBC, MYSQL, XML or catalog tables.
TABLE_LIST	The comma separated list of a TBL table sub-tables.
DBNAME	The target database for ODBC, MYSQL or catalog tables.
DATA_CHARSET	The character set used in the external file or data source.
SEP_CHAR	Specifies the field separator character of a CSV tables.
QCHAR	Specifies the character used for quoting some fields of a CSV table or the identifiers of an ODBC tables.
MODULE	The (path) name of the DLL or shared lib implementing the access of a non-standard (OEM) table type.

© SkySQL Ab. Commercial in Confidence

CONNECT Table Options

Option	Description
SUBTYPE	The subtype of an OEM table type.
CATFUNC	The catalog function used by a catalog table.
OPTION_LIST	Used to specify all other options not yet directly defined.
MAPPED	Specifies whether “file mapping” is used to handle the table file.
HUGE	To specify that a table file can be larger than 2GB.
COMPRESS	True if the data file is compressed. Defaults to NO.
SPLIT	True for a VEC table when each columns are in separate files.
READONLY	True if the data file must not be modified or erased..
SEPINDEX	When true, indexes are saved in separate files.
LRECL	The file record size (often calculated by default).
BLOCK_SIZE	The number of rows each block of FIX, BIN, DBF OR VEC tables contains. For an ODBC table this is the RowSet size option.

© SkySQL Ab. Commercial in Confidence

CONNECT Table Options

Option	Description
MULTIPLE	Used to specify multiple file tables.
HEADER	Applies to CSV, VEC and HTML files. Its meaning depends on the table type.
QUOTED	The level of quoting used in CSV table files.
ENDING	End of line length. Default to 1 for Unix/Linux and 2 for Windows.

© SkySQL Ab. Commercial in Confidence

CONNECT - Table List Table (TBL)

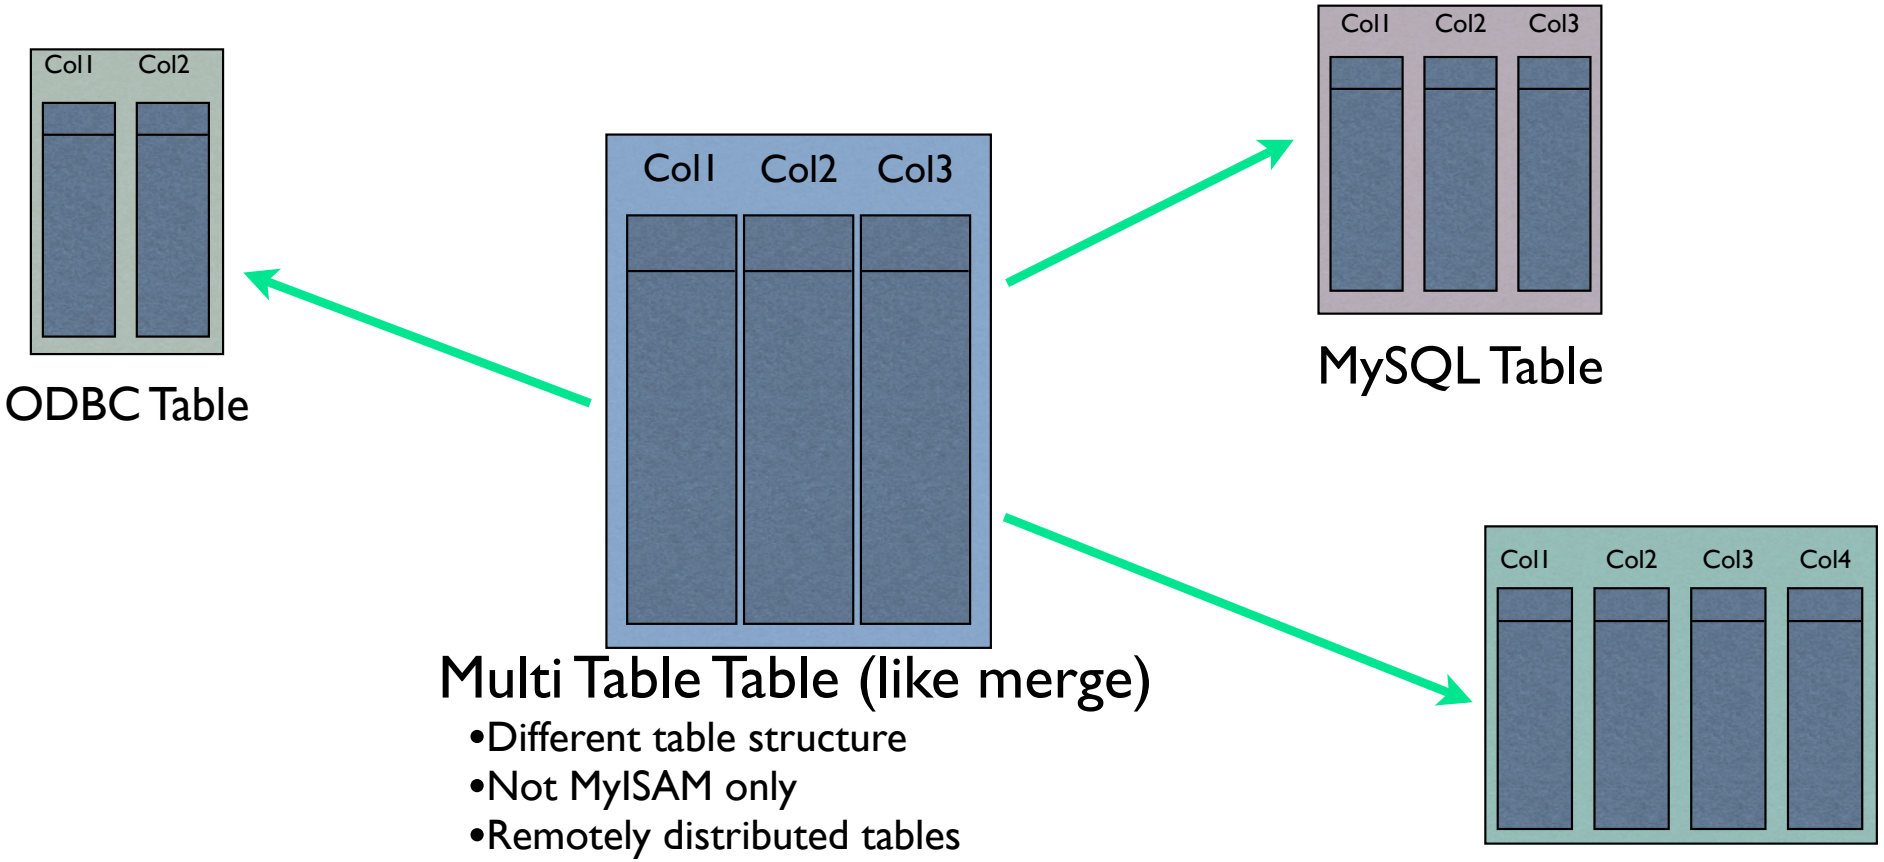

© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

- Create a HTML table and export data

```
CREATE TABLE employees engine=connect table_type=XML file_name='/var/
lib/mysql_connect/employees.html' header=yes
option_list='name=TABLE,coltype=HTML,attribute=border=1;cellpadding=5,
headattr=bgcolor=yellow‘ SELECT emp_no, birth_date, first_name,
last_name, hire_date FROM employees.employees;
Query OK, 300024 rows affected (20.99 sec)
Records: 300024  Duplicates: 0  Warnings: 0
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

```
MariaDB [connect_examples]> select * from employees limit 5;
+--------+------------+------------+-----------+------------+
| emp_no | birth_date | first_name | last_name | hire_date  |
+--------+------------+------------+-----------+------------+
|  10001 | 1953-09-02 | Georgi     | Facello   | 1986-06-26 |
|  10002 | 1964-06-02 | Bezalel    | Simmel    | 1985-11-21 |
|  10003 | 1959-12-03 | Parto      | Bamford   | 1986-08-28 |
|  10004 | 1954-05-01 | Chirstian  | Koblick   | 1986-12-01 |
|  10005 | 1955-01-21 | Kyoichi    | Maliniak  | 1989-09-12 |
+--------+------------+------------+-----------+------------+
5 rows in set (44.32 sec)
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

```
MariaDB [connect_examples]> select * from employees limit 5;
+-----+
| e
+
|
|
|
|
|
|
+
5
```

```
[root@Sky1 mysql_connect]# head -17 employees.html
<?xml version="1.0" encoding="UTF-8"?>
<!-- Created by CONNECT Version 1.01.0008 August 18, 2013 -->
<TABLE border="1" cellpadding="5">
     <TR>
          <TD>10001</TD>
          <TD>-515379600</TD>
          <TD>Georgi</TD>
          <TD>Facello</TD>
          <TD>520124400</TD>
     </TR>
     <TR>
          <TD>10002</TD>
          <TD>-176173200</TD>
          <TD>Bezalel</TD>
          <TD>Simmel</TD>
          <TD>501379200</TD>
     </TR>
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

- Create a CSV table and an index

```
CREATE TABLE employees_csv ENGINE=connect table_type=CSV
file_name='/var/lib/mysql_connect/employees.csv'
SELECT emp_no, birth_date, first_name, last_name, hire_date FROM
employees.employees;
Query OK, 300024 rows affected (1.36 sec)
Records: 300024  Duplicates: 0  Warnings: 0
```

```
CREATE UNIQUE INDEX employees_csv
    ON employees_csv(emp_no);
Query OK, 0 rows affected (0.11 sec)
Records: 0  Duplicates: 0  Warnings: 0
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

```
EXPLAIN SELECT * FROM employees_csv WHERE emp_no = 10005;
+------+-------------+---------------+-------+---------------+---------------+---------+-------+------+-------+
| id   | select_type | table         | type  | possible_keys | key           | key_len | ref   | rows | Extra |
+------+-------------+---------------+-------+---------------+---------------+---------+-------+------+-------+
|    1 | SIMPLE      | employees_csv | const | employees_csv | employees_csv | 4       | const |    1 |       |
+------+-------------+---------------+-------+---------------+---------------+---------+-------+------+-------+
1 row in set (0.00 sec)
```

```
+---------------+-------+---------------+---------------+
| table         | type  | possible_keys | key           |
+---------------+-------+---------------+---------------+
| employees_csv | const | employees_csv | employees_csv |
+---------------+-------+---------------+---------------+
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

- Create a ODBC table

```
create table Customer (
  CustomerID varchar(5),
  CompanyName varchar(40),
  ContactName varchar(30),
  ContactTitle varchar(30),
  Address varchar(60),
  City varchar(15),
  Region varchar(15),
  PostalCode varchar(10),
  Country varchar(15),
  Phone varchar(24),
  Fax varchar(24))
engine=connect table_type=ODBC block_size=10
tabname='Customers'
Connection='DSN=MS Access Database;DBQ=C:/Program
Files/Microsoft Office/Office/1033/FPNWIND.MDB;';
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

- Create a ODBC table without column definition

```
create table Customer engine=connect table_type=ODBC
  block_size=10 tabname='Customers'
  Connection='DSN=MS Access Database;DBQ=C:/Program Files/Microsoft
Office/Office/1033/FPNWIND.MDB;';
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

- Create a ODBC table with a subset of the columns defined in the data source table
  - Column names must exit in data source table

```
create table empodbc (
  EMP_NO smallint(5) not null,
  FULL_NAME varchar(37) not null),
  PHONE_EXT varchar(4) not null,
  HIRE_DATE date,
  DEPT_NO smallint(3) not null,
  JOB_COUNTRY varchar(15),
  SALARY double(12,2) not null)
engine=CONNECT table_type=ODBC tabname='EMPLOYEE'
connection='DSN=firebird';
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

- Create a MySQL table (“Federated”)
  - Syntax as used for Federated

```
create table essai (
  num integer(4) not null,
  line char(15) not null)
engine=CONNECT table_type=MYSQL
connection='mysql://root@localhost/test/people';
```

  - CONNECT standard options

```
create table essai (
  num integer(4) not null,
  line char(15) not null)
engine=CONNECT table_type=MYSQL dbname=test tabname=people
connection='mysql://root@localhost';
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

```
CREATE TABLE MyFileSys (
  PATH varchar(256) NOT NULL flag=1,
  FNAME varchar(256) NOT NULL,
  FTYPE char(4) NOT NULL,
  SIZE double(12,0) NOT NULL flag=5)
engine=CONNECT table_type=DIR file_name='/var/lib/mysql/*'
option_list='subdir=1';
```

```
SELECT * FROM MyFileSys;
+----------------------------------+-----------------------------------------------------+-------+-----------+
| PATH                             | FNAME                                               | FTYPE | SIZE      |
+----------------------------------+-----------------------------------------------------+-------+-----------+
| /var/lib/mysql/mysql/            | plugin                                              | .MYI  |      2048 |
| /var/lib/mysql/mysql/            | spider_tables                                       | .MYD  |         0 |
| /var/lib/mysql/mysql/            | time_zone_transition                                | .MYD  |         0 |
| /var/lib/mysql/mysql/            | time_zone_transition                                | .MYI  |      1024 |
| /var/lib/mysql/mysql/            | slow_log                                            | .CSM  |        35 |
| /var/lib/mysql/mysql/            | plugin                                              | .frm  |      1518 |
| /var/lib/mysql/mysql/            | db                                                  | .MYI  |      9216 |
| /var/lib/mysql/mysql/            | help_relation                                       | .MYI  |     19456 |
| /var/lib/mysql/mysql/            | proxies_priv                                        | .MYI  |     10240 |
...
| /var/lib/mysql/                  | Sky1                                                | .pid  |         5 |
| /var/lib/mysql/                  | multi-master                                        | .inf  |         0 |
+----------------------------------+-----------------------------------------------------+-------+-----------+
209 rows in set (0.00 sec)
```


© SkySQL Ab. Commercial in Confidence

CONNECT Engine – Example

```
SELECT * FROM MyFileSys;
+----------------------+--------------------------+-------+-----------+
| PATH                 | FNAME                    | FTYPE | SIZE      |
+----------------------+--------------------------+-------+-----------+
| /var/lib/mysql/mysql/ | plugin                   | .MYI  |      2048 |
| /var/lib/mysql/mysql/ | spider_tables            | .MYD  |         0 |
| /var/lib/mysql/mysql/ | time_zone_transition     | .MYD  |         0 |
| /var/lib/mysql/mysql/ | time_zone_transition     | .MYI  |      1024 |
| /var/lib/mysql/mysql/ | slow_log                 | .CSM  |        35 |
| /var/lib/mysql/mysql/ | plugin                   | .frm  |      1518 |
| /var/lib/mysql/mysql/ | db                       | .MYI  |      9216 |
| /var/lib/mysql/mysql/ | help_relation            | .MYI  |     19456 |
| /var/lib/mysql/mysql/ | proxies_priv             | .MYI  |     10240 |
...
| /var/lib/mysql/      | Sky1                     | .pid  |         5 |
| /var/lib/mysql/      | multi-master             | .inf  |         0 |
+----------------------+--------------------------+-------+-----------+
209 rows in set (0.00 sec)
```


© SkySQL Ab. Commercial in Confidence

CONNECT - Where??

- It is 100 % open source
- It is available on MariaDB launchpad
- Open Bug database
- Public Roadmap
- Released test cases
- Improvement request / worklog
- Documentation https://mariadb.com/kb/en/connect/

© SkySQL Ab. Commercial in Confidence

Questions

© SkySQL Ab. Commercial in Confidence

Web, Doc & Knowledge Base:
www.mariadb.org
www.mariadb.com

Downloads:
www.skysql.com/downloads
https://mariadb.com/resources/
https://downloads.mariadb.org/

Ralf Gebhardt
ralf.gebhardt@skysql.com

MySQL is a registered trademark of Oracle and/or its affiliates. Other names may be trademarks of their respective owners
SkySQL is not affiliated with MySQL.