



MariaDB

The MariaDB CONNECT Storage Engine

Serge Frezefond

<http://serge.frezefond.com>

@sfrezefond



Who am I?

- Serge Frezefond
- Principal Sales Engineer @ SkySQL
- Joined MySQL Ab in 2006
- Worked for MySQL@Sun and MySQL@Oracle until July 2011



Goal of the CONNECT Storage Engine : BI on various targets

Most of the data in companies is in various external datasources (many in non relational database format) :

- relational databases: Oracle, SQL Server...
- Dbase, Firebird, SQLite
- Microsoft Access & Excel
- Distributed mysql servers
- DOS, FIX, BIN, CSV, XML
- stored per column...

Not targeted for OLTP



Behind the scene Traditional BI

Data is processed by an ETL

- Change in the data model(denormalization...)

Agregates are computed

- Need to be defined and maintained

Might need to move data out of RDBMS to other kind of datastore

- OLAP, Collumn store, Hadoop/Hbase ...

Specific tools are used to query the data

IT is involved to maintain this machinery



The CONNECT Storage Engine

- What is the CONNECT storage engine?
 - 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
 - Developed by Olivier Bertrand, an ex IBM database researcher
 - The idea dates back in 2004 and Olivier has been in touch with MySQL and MariaDB since

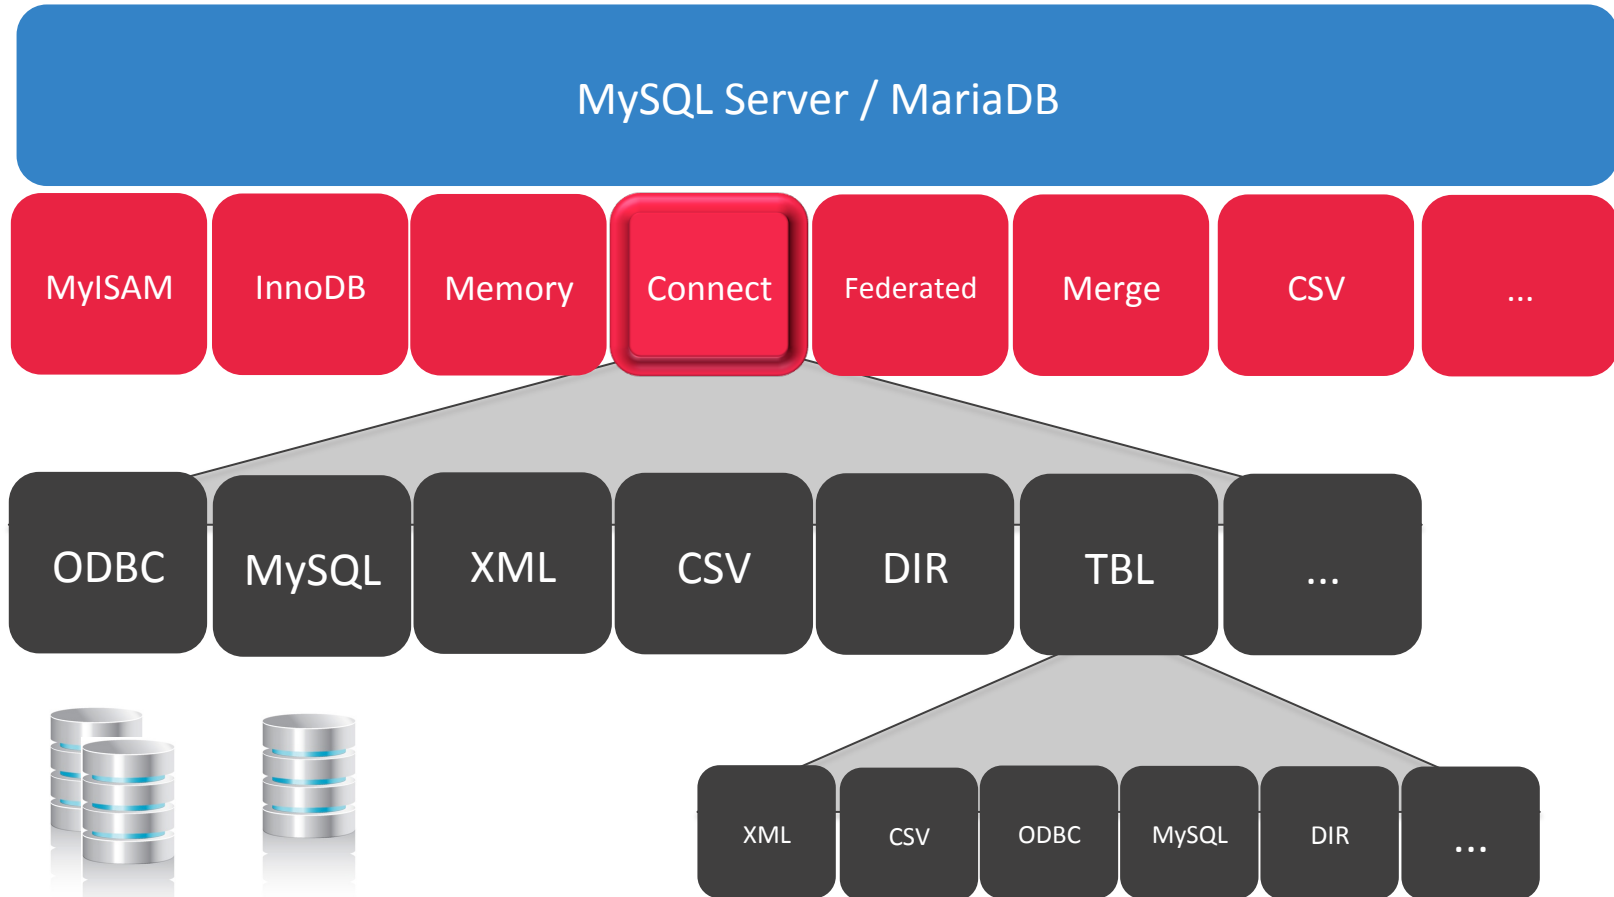


The MySQL Plugin Architecture

- 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



The CONNECT Storage Engine





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
- More powerful than CSV, FederatedX and Merge Engines
- FILE privilege is required



The CONNECT Storage Engine implements advanced features

- Condition Push down
 - Used with ODBC and MySQL to push condition to the target database. Big perf gain
- ```
set optimizer_switch='engine_condition_pushdown=on'
```
- Support MariaDB virtuals columns
  - Support of special columns :
    - Rowid, fileid, tabid, servid
  - Extensible with the OEM file type
  - Catalog table for table metadata(ODBC ...)



# CONNECT

## File table type

- DOS, FIX, BIN, FMT, CSV, INI, XML
- Support virtual tables (DIR)
- Large tables support (>2GB)
- Compression - gzlib format
- Memory file mapping
- Add read optimized indexing to files
- Multiple CONNECT tables can be created on the same underlying file
- Indexes can be shared between tables



MariaDB

# XML Table Type

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<BIBLIO SUBJECT="XML">
 <BOOK ISBN="9782212090819" LANG="fr" SUBJECT="applications">
 <AUTHOR>
 <FIRSTNAME>Jean-Christophe</FIRSTNAME>
 <LASTNAME>Bernadac</LASTNAME>
 </AUTHOR>
 <TITLE>Construire une application XML</TITLE>
 <PUBLISHER>
 <NAME>Eyrolles</NAME> <PLACE>Paris</PLACE>
 </PUBLISHER>
 <DATEPUB>1999</DATEPUB>
 </BOOK>
</BIBLIO>
```



MariaDB

# XML Table Type

```
create table xsampall (
 isbn char(15) field_format='@ISBN',
 authorln char(20) field_format='AUTHOR/LASTNAME',
 title char(32) field_format='TITLE',
 translated char(32) field_format='TRANSLATOR/@PREFIX',
 year int(4) field_format='DATEPUB')
engine=CONNECT table_type=XML file_name='Xsample.xml'
tablename='BIBLIO' option_list='rownode=BOOK,skipnull=1';
```



# XMLTable Type Query Result

```
select isbn, subject, title, publisher from xsamp2;
```

<b>ISBN</b>	<b>SUBJEC</b>	<b>TTITLE</b>	<b>PUBLISHER</b>
9782212090819	applications	Construire une application XML	Eyrolles Paris
9782840825685	applications	XML en Action	Microsoft Press

Can also generate HTML



MariaDB

## XCOL Table Type

<b>Name</b>	<b>childlist</b>
Sophie	Manon, Alice, Antoine
Valentine	Arthur, Sidonie, Prune

```
CREATE TABLE xchild (
 mother char(12) NOT NULL flag=1,
 child varchar(30) DEFAULT NULL flag=2)
ENGINE=CONNECT table_type=XCOL
tablename='children' option_list='colname=child';
```



MariaDB

## XCOL Table Type

```
select * from xchild;
```

<b>mother</b>	<b>child</b>
---------------	--------------

Sophie	Manon
--------	-------

Sophie	Alice
--------	-------

...

```
select count(child) from xchild; returns 10
```



# OCCUR Table Type

MariaDB

<b>Name</b>	<b>dog</b>	<b>cat</b>	<b>rabbit</b>	<b>bird</b>	<b>fish</b>
-------------	------------	------------	---------------	-------------	-------------

John	2	0	0	0	0
------	---	---	---	---	---

Bill	0	1	0	0	0
------	---	---	---	---	---

Mary	1	1	0	0	0
------	---	---	---	---	---

...

```
create table xpet (
```

```
 name varchar(12) not null,
```

```
 race char(6) not null,
```

```
 number int not null)
```

```
engine=connect table_type=occur tablename=pets
```

```
option_list='OccurCol=number,RankCol=race'
```

```
Colist='dog,cat,rabbit,bird,fish';
```





# OCCUR Table Type

MariaDB

```
select * from xpet;
```

<b>Name</b>	<b>race</b>	<b>number</b>
John	dog	2
Mary	dog	1
Mary	cat	1
Lisbeth	rabbit	2
...		



MariaDB

## PIVOT Table Type

<b>Who</b>	<b>Week</b>	<b>What</b>	<b>Amount</b>
Joe	3	Beer	18.00
Beth	4	Food	17.00
Janet	5	Beer	14.00
Joe	3	Food	12.00

...

```
create table pivex
```

```
Engine=connect table_type=pivot
tablename=expenses;
```



## PIVOT Table Type

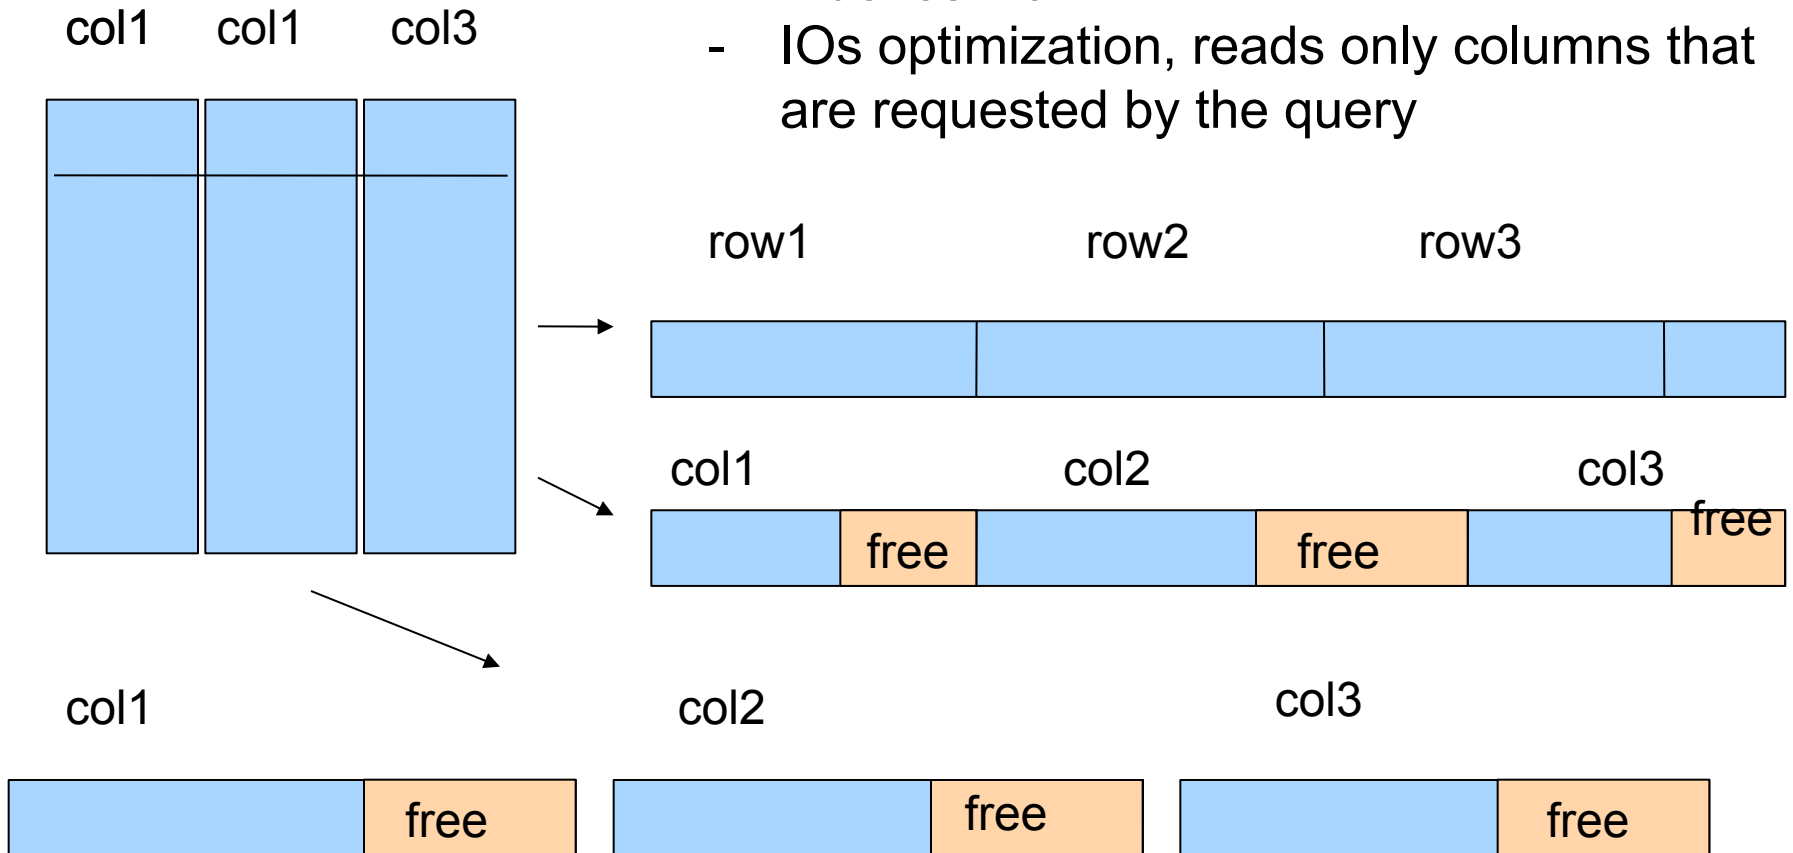
```
select * from pivex;
```

<b>Who</b>	<b>Week</b>	<b>Beer</b>	<b>Car</b>	<b>Food</b>
Beth	3	16.00	0.00	0.00
Beth	4	15.00	0.00	17.00
Beth	5	20.00	0.00	12.00
Janet	3	18.00	19.00	18.00
...				

# Connect Storage Engine

## VEC table / Column store

- 1 or per column file
- Indexes work
- IOs optimization, reads only columns that are requested by the query





# CONNECT Storage Engine ODBC table type

Allow to access to any ODBC datasource.

- Excel, Access, Firebird, SQLite
- SQL Server, Oracle, DB2
- Supports insert, update, delete and any other commands
- Multi files ODBC: consolidated monthly excel datasheet
- Access to ODBC and UnixODBC data sources
- WHERE conditions are push to the ODBC source



## ODBC table type Access db example

```
create table customers engine=connect
table_type=ODBC block_size=10
tablename='Customers'
```

```
Connection='DSN=MS Access Database;DBQ=C:/
Program Files/Microsoft Office/Office/1033/
FPNWIND.MDB;';
```



# ODBC database access From a linux box

- UnixODBC must be used as an ODBC Driver manager.
- The ODBC driver of the target database must be installed
  - For Oracle, DB2
    - install Oracle Database instant Client with ODBC supplement



MariaDB

ODBC access database  
any command to ODBC target

```
create table crlite (
 command varchar(128) not null,
 number int(5) not null flag=1,
 message varchar(255) flag=2)
engine=connect table_type=odbc
connection='Driver=SQLite3 ODBC
Driver;Database=test.sqlite3;NoWCHAR=yes'
option_list='Execsrc=1';
```





MariaDB

## ODBC Database Access

Any command to ODBC target

```
select * from crlite where command =
'update lite set birth = "2012-07-14" where ID = 2';
```

Can be wrapped in a procedure :

```
create procedure send_cmd(cmd varchar(255))
select * from crlite where command = cmd;
call send_cmd('drop tlite');
```



# CONNECT Storage Engine

## MYSQL table type vs. Federated(X)

- Condition LIMIT push down
- Implements condition push down
- Autodiscovery of table structure
- Can define the subset of columns we want to see and type conversion
- Access local or remote MySQL tables


# Connect Storage Engine MYSQL table type (a proxy table)

same syntax as federatedx :

```
create Table lineitem1
```

```
ENGINE=CONNECT TABLE_TYPE=MYSQL
```


```
connection='mysql://proxy:pwd1@node1:3306/dbt3/lineitem3';
```



# MYSQL table Type

## Remote Query Execution

```
create Table lineitem1
ENGINE=CONNECT TABLE_TYPE=MYSQL
SRCDEF='select l_suppkey, sum(l_quantity) qt from
dbt3.lineitem3 group by l_suppkey'
connection='mysql://proxy:pwd1@node1:3306/dbt3/lineitem3';
```

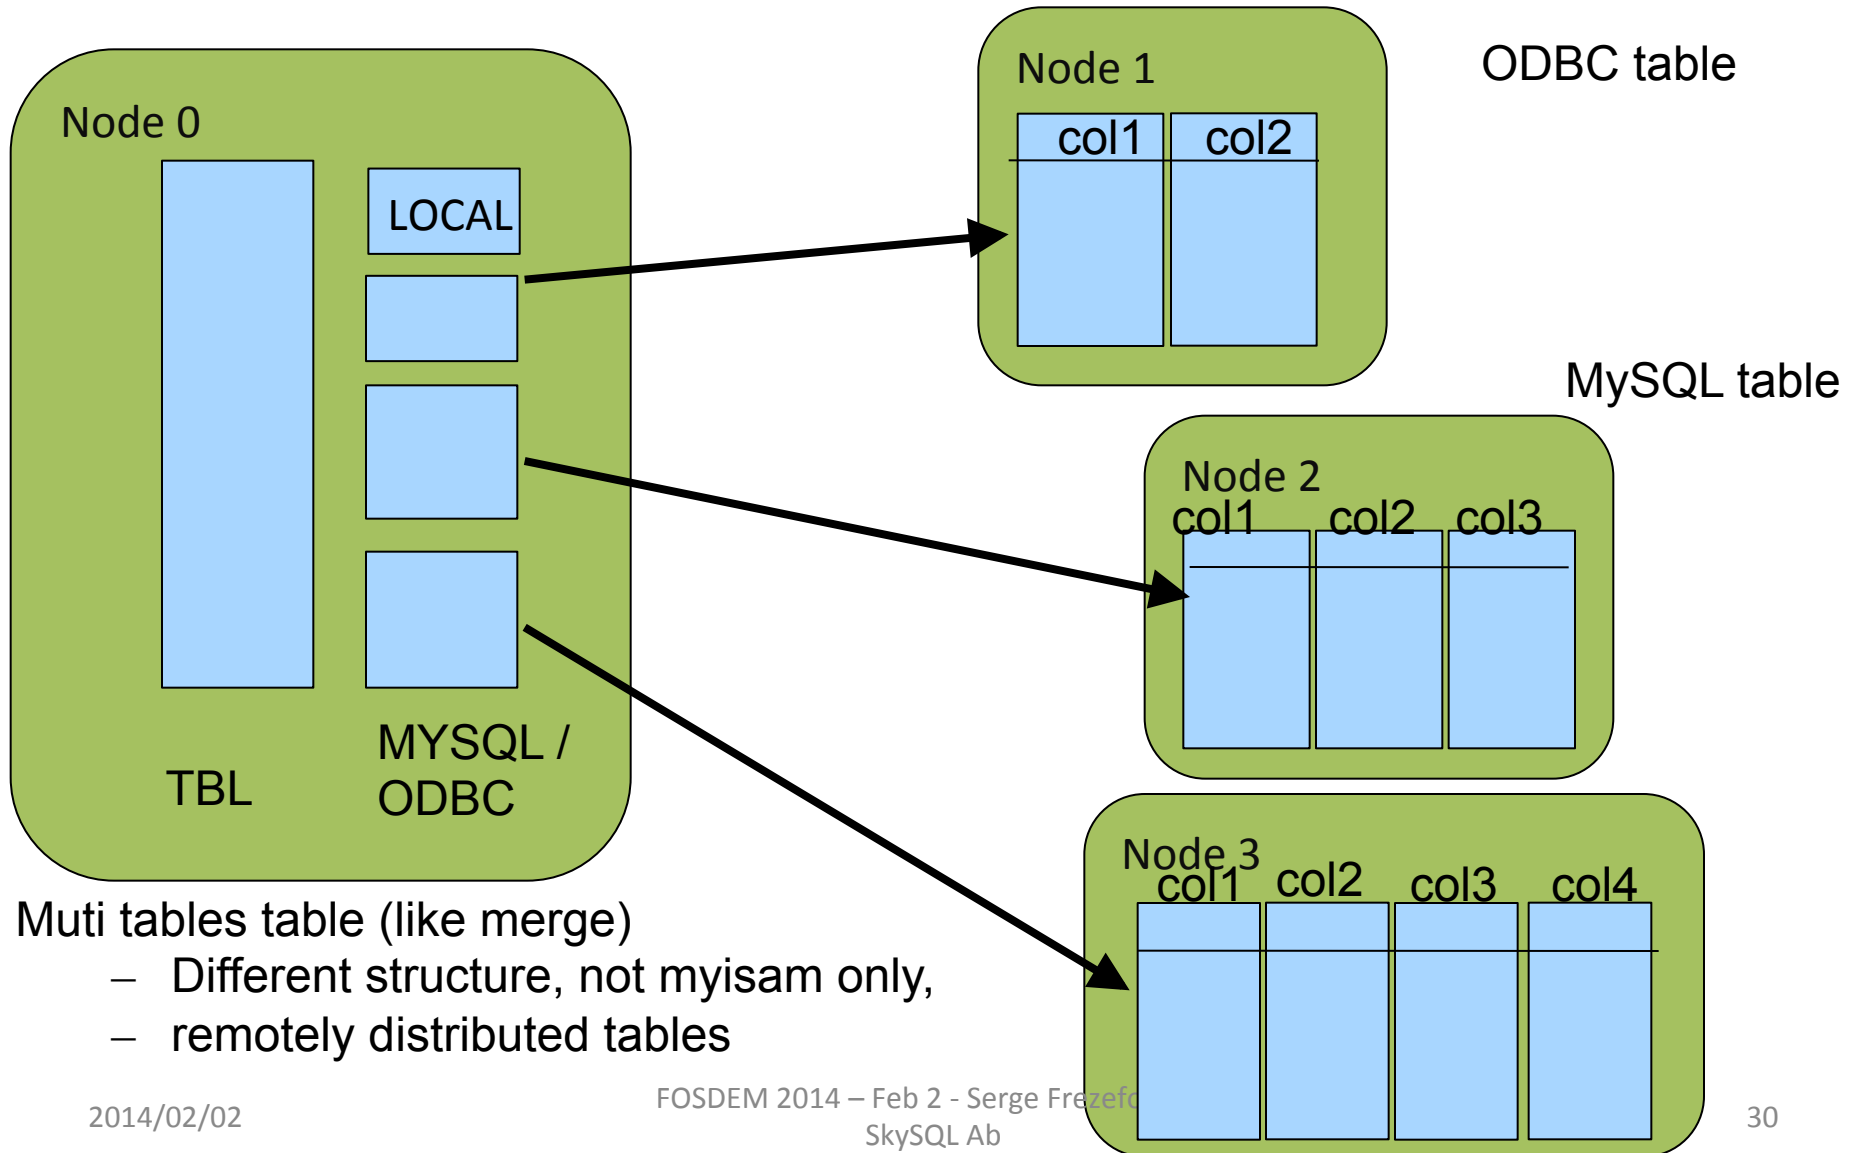




## Connect Storage Engine TBL - Table List Table

- Table list table : Collection of tables seen as one
- Tables can be from different storage engines (Not only MyISAM tables)
- Tables may have different column structure
- Underlying tables can be remote / Distributed architecture (ODBC, MySQL)

# TBL Table Type (// Merge)



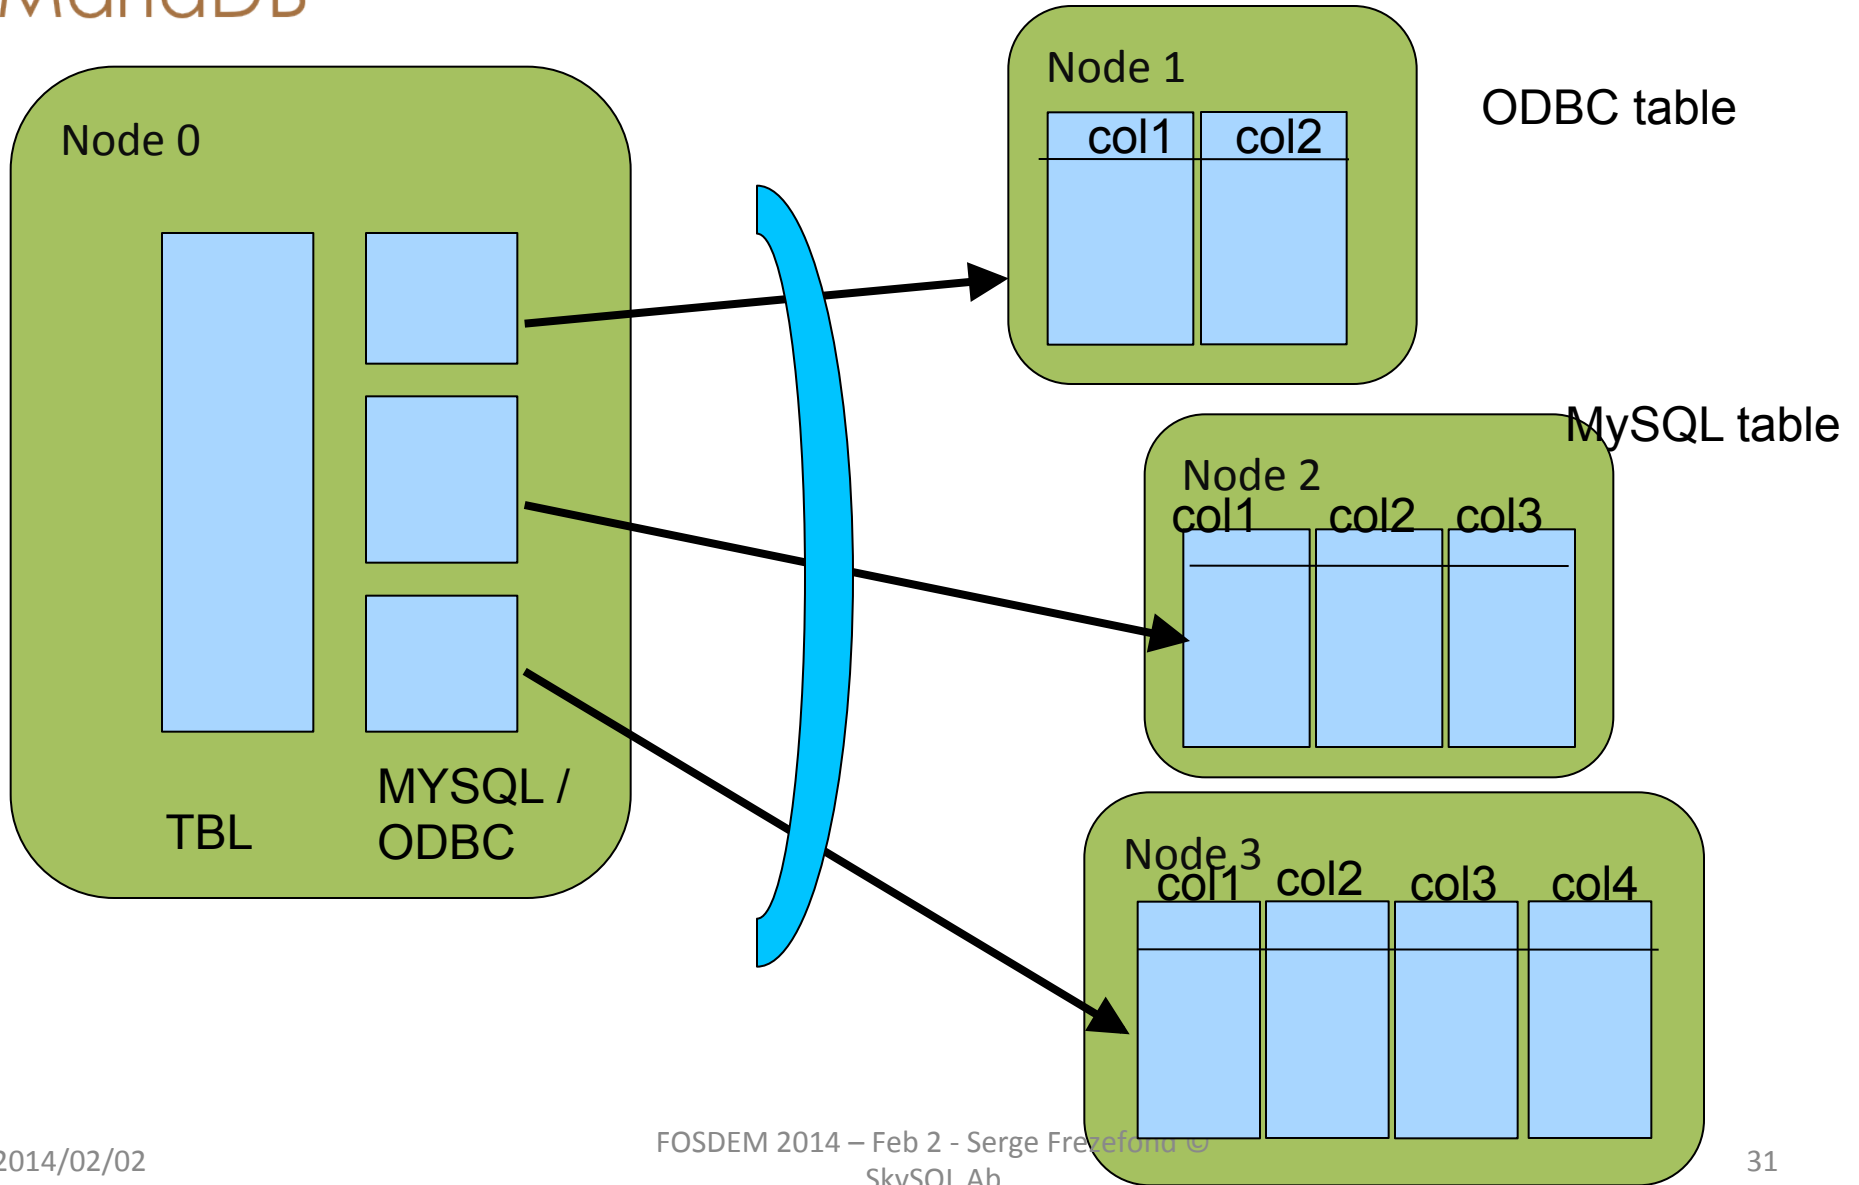
Muti tables table (like merge)

- Different structure, not myisam only,
- remotely distributed tables



# Parallel execution on distributed sharded tables

MariaDB





## Importing /exporting MySQL data in various formats

### Importing file data into MySQL tables

– Here for example from an XML file :

*–create table biblio select \* from xsampall2;*

Exporting data from MySQL: Here f we export to XML  
format :

*create table handout engine=CONNECT table\_type=XML*

*file\_name='handout.htm' header=yes*

*option\_list='name=TABLE,coltype=HTML,attribute=border=1;cellpadding=5'*

*select plugin\_name handler, plugin\_description description from*

*information\_schema.plugins where plugin\_type = 'STORAGE ENGINE';*





## Ideas / Roadmap

- Alter table improvement
- ODBC type improvement
- MySQL table type improvement
- Batch key access(MRR/BKA)
- Partition based TBL type(Like Spider)
- Adaptative query ( // MySQL Cluster) ?
- JSON File format
- Transactional / XA support



# CONNECT is open source

## You can help

- It is 100 % open source
- Sources on MariaDB launchpad
- Open Bug database
- Public Roadmap
- Test cases are released
- Improvement request / worklog
- Well Documented

## Try it



MariaDB

## Conclusion

- The MariaDB Connect Storage Engine :
  - Open MariaDB to BI and data analysis
  - Simplify heterogeneous data integration
  - Brings real value to MariaDB users
  - Illustrates openness of MariaDB community
  - Supported by SkySQL / MariaDB



Serge Frezefond [serge.frezefond@skysql.com](mailto:serge.frezefond@skysql.com)  
**@sfrezefond**  
**<http://serge.frezefond.com>**

**Documentation:**

**<https://mariadb.com/kb/en/connect/>**

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.