

Topic 8 Database Replication

Dr Diarmuid Ó Briain



Licence



This work is licensed under a Creative Commons
Attribution-ShareAlike 4.0 International License.
Full License: <http://creativecommons.org/licenses/by-sa/4.0>

Learning objectives

By the end of this topic the learner will be able to:

- Develop a replication database server with its own read-only front-end webpage..

Distributed Databases



Introduction

- Data is stored across several sites, each managed by a RDBMS that can run independently.
- The location of data and degree of individual sites impact query optimisation, concurrency control and recovery.
- Distributed data is governed the Atomicity, Consistency, Isolation, Durability (ACID) properties.

Introduction


- **Distributed Data Independence:**
 - The distribution of portions of the database to various locations should be invisible to users of the database.
 - Existing applications should continue to operate successfully:
 - when a distributed version of the DBMS is first introduced; and
 - when existing distributed data are redistributed around the system.

Introduction

- **Distributed Transaction Atomicity:**
 - Users should be able to write transactions that access and update data at several sites.
 - Transactions are atomic, all changes persist if the transaction commits, or roll-back if transaction aborts.


Types of Distributed DB - Homogeneous

- **Homogeneous** – data is distributed but all servers run the same DBMS software.




Types of Distributed DB - Heterogeneous

- Different sites run different RDBMSs separately and are connected somehow to enable access to data from multiple sites.
 - Gateway protocols - API that exposes DBMS functionality to external applications. Examples: ODBC and JDBC



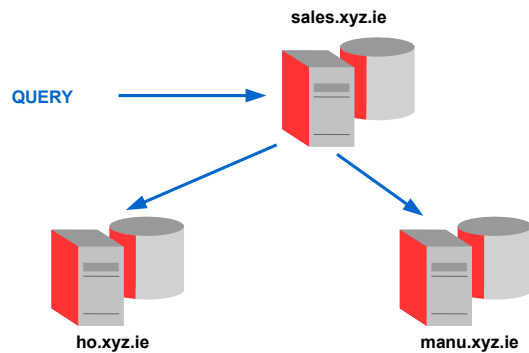
Client-Server Architecture



Connected to sales.xyz.ie

```
MariaDB [(none)]> SELECT * FROM Pricing;
MariaDB [(none)]> SELECT * FROM Pricing where Product_no = '23456';
```

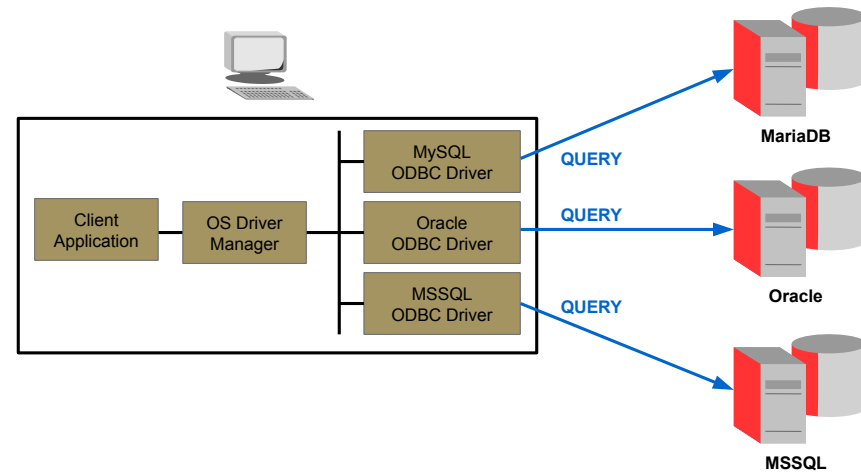
Collaborated Server Architecture




Connected to sales.xyz.ie

```
MariaDB [(none)]> SELECT * FROM Product@manu.xyz.ie;
```

Middleware Architecture



HA Cluster Architecture



HA Cluster configurations

- **Active/Active**
 - Traffic intended for the failed node is either passed onto an existing node or load balanced across the remaining nodes.
 - Homogeneous software configuration.
- **Active/passive (N+1)**
 - Provides a fully redundant instance of each node, which is only brought online when its associated primary node fails.
- **N+M**
 - In cases where a single cluster is managing many services.
 - More than one (M) standby servers are included and available.
 - Tradeoff between cost and reliability requirements.
- **N-to-1**
 - Allows the failover standby node to become the active one temporarily, until the original node back online.
- **N-to-N**
 - A combination of active/active and N+M clusters, N-to-N clusters redistribute the services, instances or connections from the failed node among the remaining active nodes.

Storing data

	1	Student ID	Name	Town
Horizontal Fragment (Rows)	2	K001234	Lovelace	Nottingham
	3	K001235	Babbage	London
	4	K001236	Menabrea	Turin

Vertical Fragment (Columns)

Replication



- The storing several copies of a relation or fragment.
- Entire relation can be stored at one or more sites.
 - **High Availability (HA)**
 - **Faster Query Evaluation**

Sync/Async Replication

- Synchronous replication writes the data to both the primary and to the secondary site database at the same time.
- Asynchronous replication also writes data to both a primary and secondary site, however with this process there is a delay when data is copied from one to another.
 - Known as "**Store and Forward**" replication.

Type of Replication	Synchronous	Asynchronous
Recovery Point Objective	Zero	15 minutes to a few hours
Distance Limitations	Best if both SANs are in the same datacentre.	Anywhere with a good data connection.
Cost	Most expensive type of SAN solution.	Not as expensive as Synchronous but more expensive than basic SANs.



Securing Distributed Data

- What systems should be made accessible to users in other locations and which users in other locations should have access to those systems.
- How tightly controlled access to those systems should be.
 - For example, should a user password be required when a conversation is started by a remote user?
- Is it required that passwords flow over the wire in encrypted form?

Securing Distributed Data

- Is it required that a user profile under which a client job runs be mapped to a different user identification or password based on the name of the relational database to which you are connecting?
- What data should be made accessible to users in other locations and which users in other locations should have access to that data.
- What actions those users should be allowed to take on the data.
- Whether authorisation to data should be centrally controlled or locally controlled.

MariaDB Security

- General factors that affect security
 - Good passwords
 - Not granting unnecessary privileges to users
 - Preventing SQL injections and data corruption, etc...
- Security of the installation itself
 - Data files
 - Log files
 - Application files.

MariaDB Security

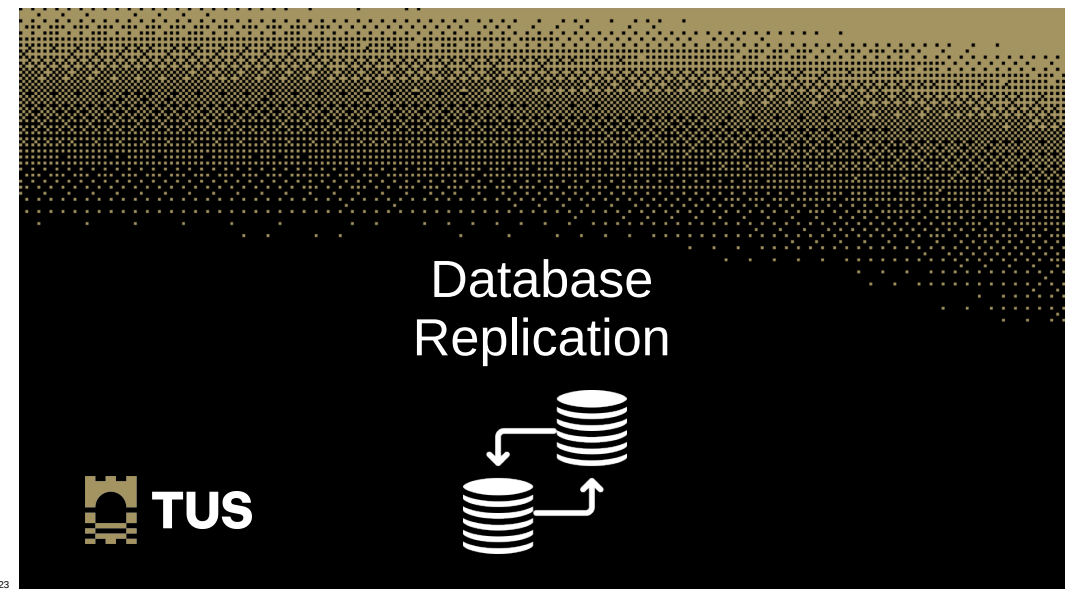
- Access control and security within the database system
 - Database users
 - Databases
- Network security of MariaDB and the system
 - Localhost only
 - Restricted list of hosts.

MySQL Security

- Secure Sockets Layer (SSL)
 - OpenSSL API.

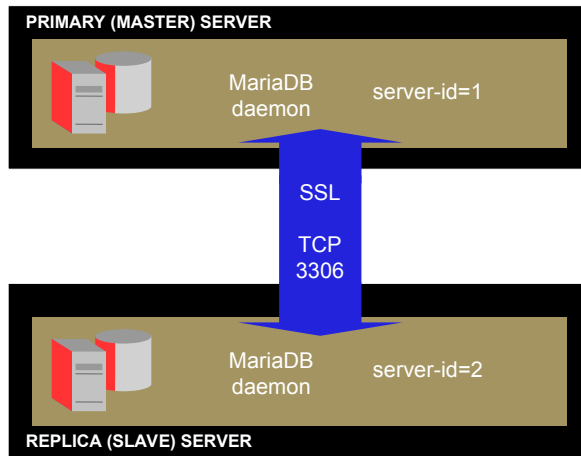
```
mysql> SHOW VARIABLES LIKE 'have_ssl';
+-----+-----+
| Variable_name | Value |
+-----+-----+
| have_ssl      | YES   |
+-----+-----+
```

- If DISABLED, the server has SSL support but not started with `--ssl-xxx` options to enable them to be used.
- Backups
 - Ensure that there are adequate and appropriate backups of the database files, configuration and log files.




The diagram features a dark background with a halftone pattern. At the top, the text "Database Replication" is written in a white, sans-serif font. Below the text, there are two stacks of three cylinders representing databases. An arrow points from the top database to the bottom database, and another arrow points from the bottom database back to the top database, forming a cycle. In the bottom left corner, the TUS logo is displayed, consisting of a stylized castle icon and the letters "TUS".

Database replication



Database replication



Install software on the Replica Server

- Install the following software on the second server:
 - `mariadb-server`
 - `mariadb-client`
 - `Python Flask`

Configure the Primary Server

- Create `replication_user` for access by the replica database

```
[Primary]~$ sudo mysql -u root

MariaDB [(none)]> GRANT REPLICATION SLAVE ON *.* TO
                    -> 'replication_user'@'%'
                    -> IDENTIFIED BY 'replication_pass';
Query OK, 0 rows affected (0.002 sec)

MariaDB [(none)]> FLUSH PRIVILEGES;
```

Configure the Primary Server

- Create `replication_user` for access by the replica database
- Note the `'%'` to allow remote access

```
MariaDB [(none)]> select User, Host from mysql.user;
+-----+-----+
| User      | Host      |
+-----+-----+
| replication_user | %         |
| admin     | localhost |
| enguser   | localhost |
| phpmyadmin | localhost |
| root      | localhost |
+-----+-----+
5 rows in set (0.002 sec)
```

Configure the Primary Server

```
[Primary]~$ sudo systemctl stop mariadb.service
```

```
[Primary]~$ cd /etc/mysql/mariadb.conf.d
[Primary]/etc/mysql/mariadb.conf.d$ sudo vi 50-server.cnf
~~~~
[mariadb]
log_bin
server_id=1
log_basename=master1
binlog_format=mixed
binlog_do_db=Eng
bind_address=0.0.0.0
~~~~
:wq!
```

Configure the Primary Server

```
[Primary]~$ sudo systemctl start mariadb.service
```

```
MariaDB [(none)]> SHOW VARIABLES LIKE 'bind%';
+-----+-----+
| Variable_name | Value      |
+-----+-----+
| bind_address  | 0.0.0.0    |
+-----+-----+
1 row in set (0.008 sec)
```

```
MariaDB [(none)]> SHOW VARIABLES LIKE 'server%';
+-----+-----+
| Variable_name | Value      |
+-----+-----+
| server_id     | 1          |
+-----+-----+
1 row in set (0.009 sec)
```

Configure the Primary Server

- Clear existing logs

```
[Primary]~$ sudo rm /var/log/mysql/*
```

```
[Primary]~$ sudo systemctl restart mariadb.service
```

```
[Primary]~$ sudo systemctl status mariadb.service
● mariadb.service - MariaDB 10.3.29 database server
   Loaded: loaded (/lib/systemd/system/mariadb.service; enabled; vendor preset:
   Active: active (running) since Fri 2023-01-02 17:48:20 GMT; 4s ago
     Docs: man:mysqld(8)
           https://mariadb.com/kb/en/library/systemd/
   Process: 19333 ExecStartPre=/usr/bin/install -m 755 -o mysql -g root -d /var/r
   Process: 19334 ExecStartPre=/bin/sh -c systemctl unset-environment _WSREP_STA
   Process: 19336 ExecStartPre=/bin/sh -c [ ! -e /usr/bin/galera_recovery ] && VA
   Process: 19433 ExecStartPost=/bin/sh -c systemctl unset-environment _WSREP_STA
   Process: 19435 ExecStartPost=/etc/mysql/debian-start (code=exited, status=0/SU
 Main PID: 19401 (mysqld)
   Status: "Taking your SQL requests now..."
   Tasks: 32 (limit: 2059)
  CGroup: /system.slice/mariadb.service
         └─19401 /usr/sbin/mysqld
```


Configure the Primary Server

- Lock the Primary Database
- Note the **File** and **Position** values

```
[Primary]~$ sudo mysql -u root

MariaDB [(none)]> FLUSH TABLES WITH READ LOCK;
Query OK, 0 rows affected (0.003 sec)

MariaDB [(none)]> SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| master1-bin.000005 |      635 | Eng          |                   |
+-----+-----+-----+-----+
1 row in set (0.001 sec)
```

Dump the Primary Database

```
[Primary]~$ sudo mysqldump -u root -p Eng > Eng.sql
Enter password: rootpass

[Primary]~$ ls
Eng.sql
```

Configure the Replica Server

- Confirm the connection to the Primary Server

```
[Replica]~$ sudo mysql -h 192.168.0.27 -u replication_user -p
Enter password: replication_pass
Welcome to the MariaDB monitor.  Commands end with ; or \g.
Your MariaDB connection id is 32
Server version: 10.6.12-MariaDB-0ubuntu0.22.04.1-log Ubuntu 22.04

Copyright (c) 2000, 2018, Oracle, MariaDB Corporation Ab and others.

Type 'help;' or '\h' for help. Type '\c' to clear the current input
statement.

MariaDB [(none)]>
```

Configure the Replica Server

- Create local version of the database

```
MariaDB [(none)]> CREATE DATABASE Eng;
Query OK, 1 row affected (0.000 sec)

MariaDB [(none)]> SHOW DATABASES;
+-----+
| Database |
+-----+
| Eng      |
| information_schema |
| mysql   |
| performance_schema |
+-----+
4 rows in set (0.000 sec)
```

Configure the Replica Server

- Create a local user for the database

```
MariaDB [(none)]> CREATE USER 'replicausер'@'localhost' IDENTIFIED BY 'replicapass';
Query OK, 0 rows affected (0.013 sec)
```

```
MariaDB [(none)]> GRANT ALL ON Eng.* TO 'replicausер'@'localhost';
Query OK, 0 rows affected (0.12 sec)
```

Configure the Replica Server

- Replica database configuration

```
[Replica]~$ sudo systemctl stop mariadb.service
```

```
[Replica]~$ cd /etc/mysql/mariadb.conf.d
/etc/mysql/mariadb.conf.d$ sudo vi 50-server.cnf
~~~~
~~~~
[mariadb]
server_id = 2
replicate_do_db = Eng
~~~~
~~~~
:wq!
```

Configure the Replica Server

- Clear the replica logs

```
[Replica]~$ sudo rm /var/log/mysql/*
```

```
[Replica]~$ sudo systemctl restart mariadb.service
```

```
[Replica]~$ sudo systemctl status mariadb.service
```

```
• mariadb.service - MariaDB 10.5.11 database server
  Loaded: loaded (/lib/systemd/system/mariadb.service; enabled; vendor prese>
  Active: active (running) since Tue 2021-08-31 20:07:29 BST; 5s ago
    Docs: man:mariadb(8)
          https://mariadb.com/kb/en/library/systemd/
  Process: 3401 ExecStartPre=/usr/bin/install -m 755 -o mysql -g root -d /var>
  Process: 3402 ExecStartPre=/bin/sh -c systemctl unset-environment _WSREP_ST>
  Process: 3404 ExecStartPre=/bin/sh -c [ ! -e /usr/bin/galera_recovery ] && >
  Process: 3475 ExecStartPost=/bin/sh -c systemctl unset-environment _WSREP_S>
  Process: 3477 ExecStartPost=/etc/mysql/debian-start (code=exited, status=0/>
 Main PID: 3463 (mariadb)
  Status: "Taking your SQL requests now..."
    Tasks: 12 (limit: 14110)
  Memory: 67.3M
    CPU: 252ms
  CGroup: /system.slice/mariadb.service
          └─3463 /usr/sbin/mariadb
```

Configure the Replica Server

- Confirm the configuration

```
MariaDB [(none)]> SHOW VARIABLES LIKE 'server_id';
```

```
+-----+-----+
| Variable_name | Value |
+-----+-----+
| server_id     | 2     |
+-----+-----+
1 row in set (0.002 sec)
```

```
MariaDB [(none)]> SHOW VARIABLES LIKE 'replicate_do_db%';
```

```
+-----+-----+
| Variable_name | Value |
+-----+-----+
| replicate_do_db | Eng   |
+-----+-----+
1 row in set (0.002 sec)
```

Configure the Replica Server

- Copy primary database to the replica

```
[Replica]~$ sftp ada@192.168.0.23
ada@192.168.0.23's password: ada_pass
Connected to 192.168.0.23.
sftp> ls
Eng.sql

sftp> get Eng.sql
Fetching /home/ada/Eng.sql to Eng.sql
/home/ada/Eng.sql          100% 2807   440.0KB/s   00:00
sftp> quit

[Replica]~$ ls
Eng.sql

[Replica]~$ mysql -u root Eng < ./Eng.sql
```

Configure the Replica Server

- Confirm the table import

```
MariaDB [(none)]> USE Eng;
Database changed

MariaDB [Eng]> SHOW TABLES;
+-----+
| Tables_in_Eng |
+-----+
| EngHobbies     |
| EngProject     |
+-----+
2 rows in set (0.001 sec)

MariaDB [Eng]> SELECT Firstname, Lastname FROM EngProject;
+-----+-----+
| Firstname | Lastname |
+-----+-----+
| Ada       | Lovelace |
+-----+-----+
1 row in set (0.001 sec)
```

Configure the Replica Server

- Link Replica database to the Primary database

```
[Replica]~$ mysql -u root -p
Enter password: rootpass

MariaDB [(none)]> CHANGE MASTER TO
-> MASTER_HOST='192.168.0.23',
-> MASTER_USER='replication_user',
-> MASTER_PASSWORD='replication_pass',
-> MASTER_LOG_FILE='master1-bin.000002',
-> MASTER_LOG_POS=330;

Query OK, 0 rows affected (0.11 sec)

Start Replica (Slave).
MariaDB [(none)]> START SLAVE;
Query OK, 0 rows affected (0.003 sec)
```

Unlock the Primary Database

```
[Primary]~$ mysql -u root

MariaDB [(none)]> UNLOCK TABLES;
Query OK, 0 rows affected (0.001 sec)
```

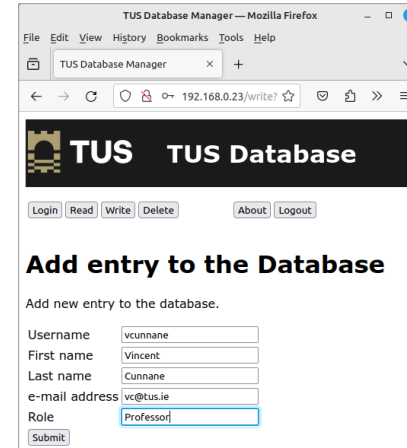
Confirm the connection

```

MariaDB [(none)]> SHOW SLAVE STATUS \G;
***** 1. row *****
Slave_IO_State: Waiting for master to send event
Master_Host: 192.168.0.23
Master_User: replication_user
Master_Port: 3306
Connect_Retry: 60
Master_Log_File: master1-bin.000003
Read_Master_Log_Pos: 330
Relay_Log_File: mysqld-relay-bin.000003
Relay_Log_Pos: 557
Relay_Master_Log_File: master1-bin.000003
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
Replicate_Do_DB: Eng
~~~~~: ~~~~~

```

Add a new entry to the Primary



TUS Database Manager — Mozilla Firefox

TUS Database

Login Read Write Delete About Logout

Add entry to the Database

Add new entry to the database.

Username

First name

Last name

e-mail address

Role

Confirm the connection

Primary

```
MariaDB [Eng]> SELECT * FROM EngProject;
```

Student_no	Username	FirstName	LastName	Email	Role
0	cbabage	Charles	Babbage	charles@bababage.com	Hardware
1	alovelace	Ada	Lovelace	ada@lovelace.com	Programmer
2	lmenabrea	Luigi	Menabrea	luigi@menabrea.it	Politician
3	cquinn	Edel	Quinn	edel@quinn.net	Nurse
4	vcampbell	Veronica	Campbell	vcampbell@setu.ie	Professor

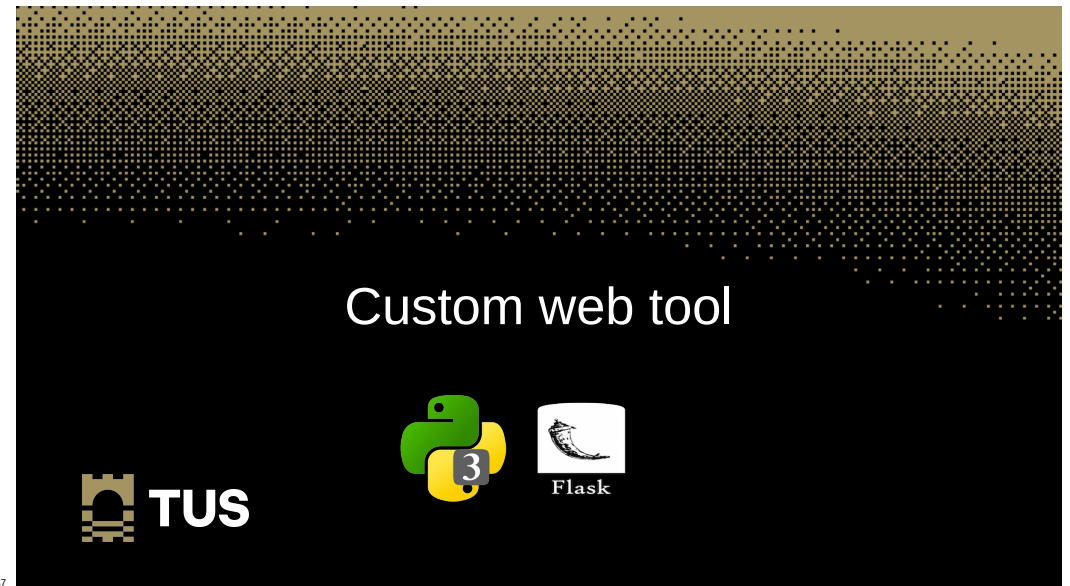
5 rows in set (0.000 sec)

Replica


```
MariaDB [Eng]> SELECT * FROM EngProject;
```


Student_no	Username	FirstName	LastName	Email	Role
0	cbabage	Charles	Babbage	charles@bababage.com	Hardware
1	alovelace	Ada	Lovelace	ada@lovelace.com	Programmer
2	lmenabrea	Luigi	Menabrea	luigi@menabrea.it	Politician
3	cquinn	Edel	Quinn	edel@quinn.net	Nurse
4	vcampbell	Veronica	Campbell	vcampbell@setu.ie	Professor


5 rows in set (0.001 sec)



Custom web tool

 TUS

 3

 Flask

Custom tool to review the Replica

- Link Replica database to the Primary database.

```
[Replica]~$ mysql -u root

MariaDB [(none)]> CREATE USER 'enguser'@'localhost' IDENTIFIED BY 'engpass';
Query OK, 0 rows affected (0.010 sec)

MariaDB [(none)]> GRANT ALL ON Eng.* TO 'enguser'@'localhost';
Query OK, 0 rows affected (0.00 sec)
```

Create a venv and install packages

```
~$ python3 -m venv ~/.venv
~$ source ~/.venv/bin/activate
(.venv) [Replica]~$

(.venv) [Replica]~$ python3 -m pip install flask mariadb pyyaml
```

Get sys.path for the packages

```
(.venv) [Replica]~$ python3
Python 3.10.12 (main, Nov 20 2023, 15:14:05) [GCC 11.4.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import sys
>>> sys.path
['', '/usr/lib/python3.10.zip', '/usr/lib/python3.10', '/usr/lib/python3.10/lib-dynload',
'/home/labuser/.venv/lib/python3.10/site-packages']
>>> quit()
(.venv) [Replica]~$ cp ~/web_repl/tools/mariadb_conn.py
/home/labuser/.venv/lib/python3.10/site-packages

(.venv) [Replica]~$ deactivate
[Replica]~$
```

- Copy mariadb_conn.py to the venv sys.path

```
[Replica]~$ cp mariadb_conn.py /home/labuser/.venv/lib/python3.10/site-packages
```

Move replica to Apache2 root and define rights

```
[Replica]~$ sudo cp -r ~/web_repl /var/www/html/web
[Replica]~$ sudo chown -R www-data: /var/www/html/web

~$ sudo usermod -a -G ada www-data
~$ sudo usermod -a -G www-data ada

~$ cat /etc/group | grep ^www
www-data:x:33:ada
~$ cat /etc/group | grep ^ada
ada:x:1000:www-data
```

Apache2 site configuration

```
[Replica]~$ cd /etc/apache2/sites-available/  
/etc/apache2/sites-available$ sudo mv 000-default.conf 000-default.conf.orig  
/etc/apache2/sites-available$ cat <<EOM | sudo tee 000-default.conf  
  
<VirtualHost *:80>  
    ServerAdmin webmaster@localhost  
  
    WSGIScriptAlias / /var/www/html/web/app.wsgi  
    <Directory /var/www/html/web>  
        Order allow,deny  
        Allow from all  
    </Directory>  
  
    ErrorLog /error.log  
    CustomLog /access.log combined  
  
</VirtualHost>  
EOM
```

Adjust the app.wsgi file

```
[Replica]~$ sudo vi /var/www/html/web/app.wsgi  
import sys  
  
sys.path.insert(0, "/home/labuser/.venv/lib/python3.10/site-packages")  
sys.path.insert(0, "/var/www/html/web")  
  
from init import app as application
```

Install WSGI library, enable WSGI

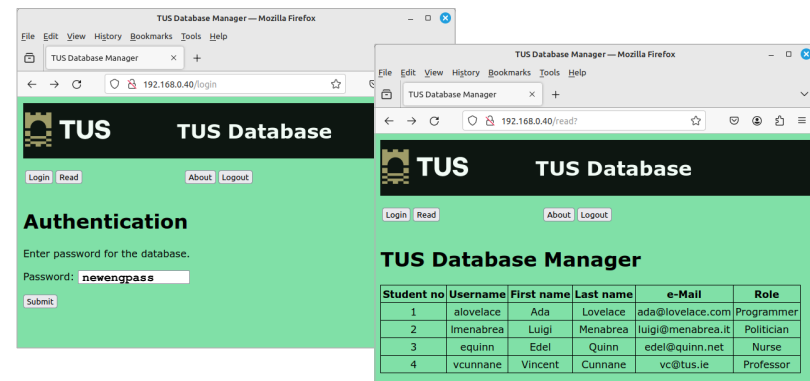
```
[Replica]~$ sudo apt install apache2-utils libapache2-mod-wsgi-py3  
[Replica]~$ sudo a2enmod wsgi  
Enabling module wsgi.
```

- Enable and Start Apache2

```
[Replica]~$ sudo systemctl enable apache2  
[Replica]~$ sudo systemctl start apache2
```

Custom tool to review the Replica

- Review the Replica database

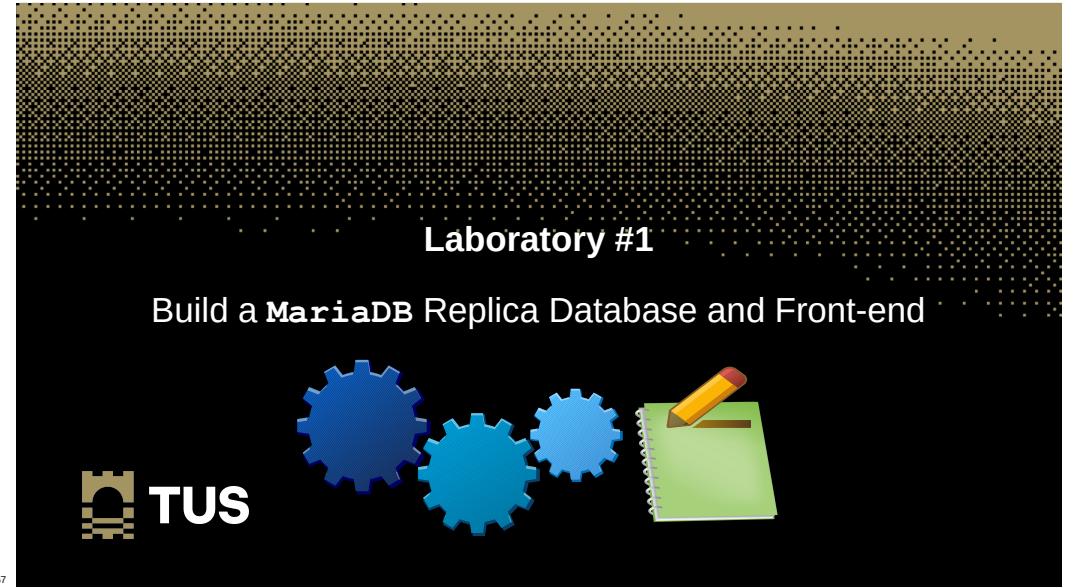


The screenshot shows two overlapping browser windows of the TUS Database Manager. The background window is at the login page, and the foreground window shows the main interface with a table of data.

Student no	Username	First name	Last name	e-Mail	Role
1	alovelace	Ada	Lovelace	ada@lovelace.com	Programmer
2	lmenabrea	Luigi	Menabrea	luigi@menabrea.it	Politician
3	equinn	Edel	Quinn	edel@quinn.net	Nurse
4	vcunnane	Vincent	Cunnane	vc@tus.ie	Professor


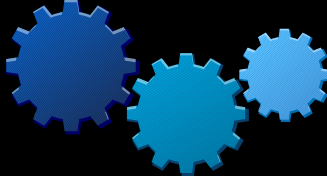

Recovery in the case of failure

- In the case of database replication failure:
 - Stop both databases
 - Clear the logs on the both servers
 - Restart the Primary
 - Restart the Replica database
 - Confirm connection from Replica to Primary.



Laboratory #1

Build a **MariaDB** Replica Database and Front-end

 **TUS**  

Laboratory #1.1 – Create a Replica Database

Create a Replica database using MariaDB

- Create a replica address database using MariaDB on a different server that mirrors the database created in the previous topic.
- Establish replication to this database to the other database such that this is the Replica (Slave) and the database from the previous topic is the Primary (Master).
- Document each stage.

Laboratory #1.2 – Create database front-end

Create a custom interface to the database

- Create a web based interface to the database built in the previous topic that has read-only access to the data.
- Document each stage.

Learning Objectives

- Develop a replication database server with its own read-only front-end webpage. ✓



TUS
Oibiceil Teicneolaíochta na Sionainne:
Lár Tíre, An Bhaile na Lár
Technological University of the Shannon:
Midlands Midwest

EUR ING Dr Diarmuid Ó Briain
Innealtóir Cairte agus
Léachtóir Sinsearach

E diarmuid.obriain@tus.ie | W tus.ie
Campas Maolais, Páirc Maolais,
Luimneach, V94 EC5T, Éire

CEng, FIEI CISSP®

Thank you

TUS