

Technical Guide

GUIDE #7: When Disaster Strikes: How Does EZSTACKR and Gluster FS Recover ?

Introduction:

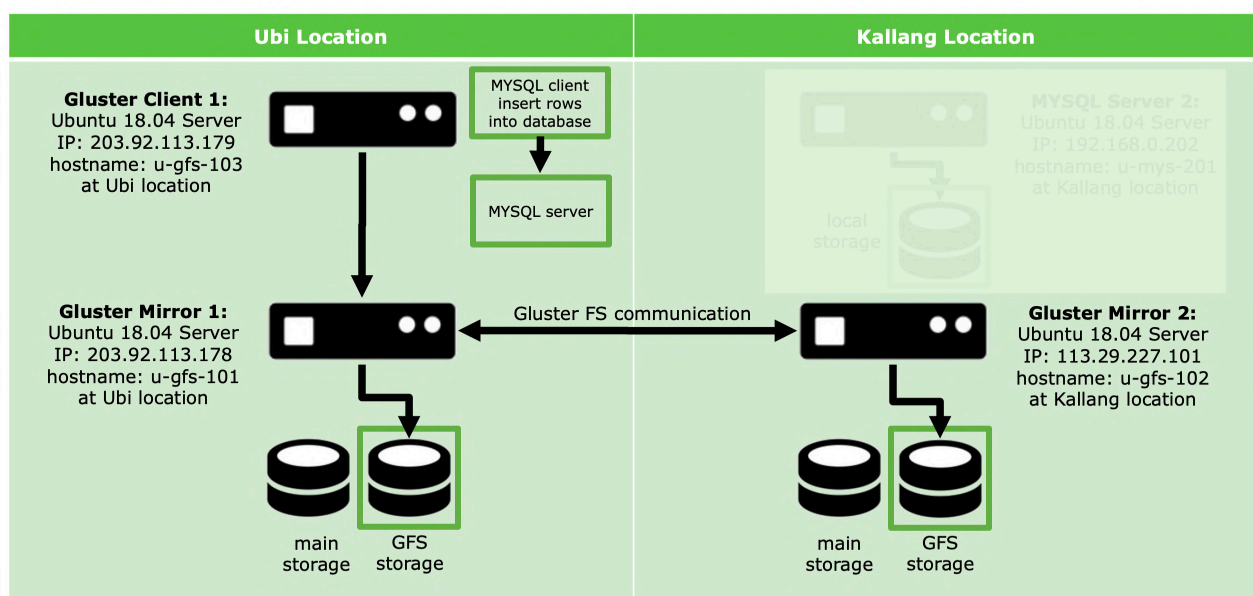
A storage mirror is a complete backup of the data that can be used if the primary data store fails. Storage mirroring facilitates high availability of data systems and the applications that use them. In this guide, a Gluster mirror node and the MYSQL service is "destroyed". Another Gluster mirror node survives. MYSQL service is down. When MYSQL service is restored, it relies on the surviving Gluster mirror node. MYSQL operations e.g. SQL inserts start up again. As SQL operations continue, the "destroyed" node is restored, it rejoins the Gluster cluster. Data is resynced. MYSQL operations continue as usual.

Storage Mirror Scenario Summary Part 2 (details in the next Technical guide):


In storage mirror scenario part 2, MYSQL is introduced into the Gluster mirror cluster.

1. Setting up the Gluster mirror cluster for MYSQL data store
2. Installing MYSQL server
3. Creating MYSQL user, database, tables and scripts
4. Inserting data into MYSQL
5. Verifying the integrity of the MYSQL data written to Gluster storage mirror pool
6. Destroying MYSQL server and destroying first Gluster node
7. Restoring MYSQL data from second Gluster node
8. Observing first Gluster node when it is repaired and rejoins the Gluster cluster


A Mirroring Scenario on Ezstackr with MYSQL



Recovery Scenario when Ubi Location is Destroyed



Copy MYSQL Data Out from Mirror and Start New MYSQL



SETTING UP THE GLUSTER MIRROR CLUSTER FOR MYSQL DATA STORE

In technical guide 6, the Gluster mirror was set up. No changes need to be made at the Gluster server mirror. Let’s review some of the commands needed to set up the Gluster mirror:

```
gluster peer probe u-gfs-102
gluster peer status
gluster pool list
gluster volume create volume1 replica 2 u-gfs-101:/gluster-storage u-gfs-102:/gluster-storage force
gluster volume start volume1
sudo gluster volume status
```

INSTALLING MYSQL SERVER

Choosing u-gfs-103 to install the MYSQL server and reviewing the file system:

```
df -H
```

Unmounting any Gluster file system from previous tech guide setup:

```
sudo umount /storage-pool
```

Making MYSQL DATA directory:

Typically, MYSQL stores its data and configuration files at /var/lib/mysql and creates this directory during installation. In this case, let’s create this directory and mount it to Gluster before MYSQL installation begins.

```
sudo mkdir /var/lib/mysql
```

Mount the Gluster volume:

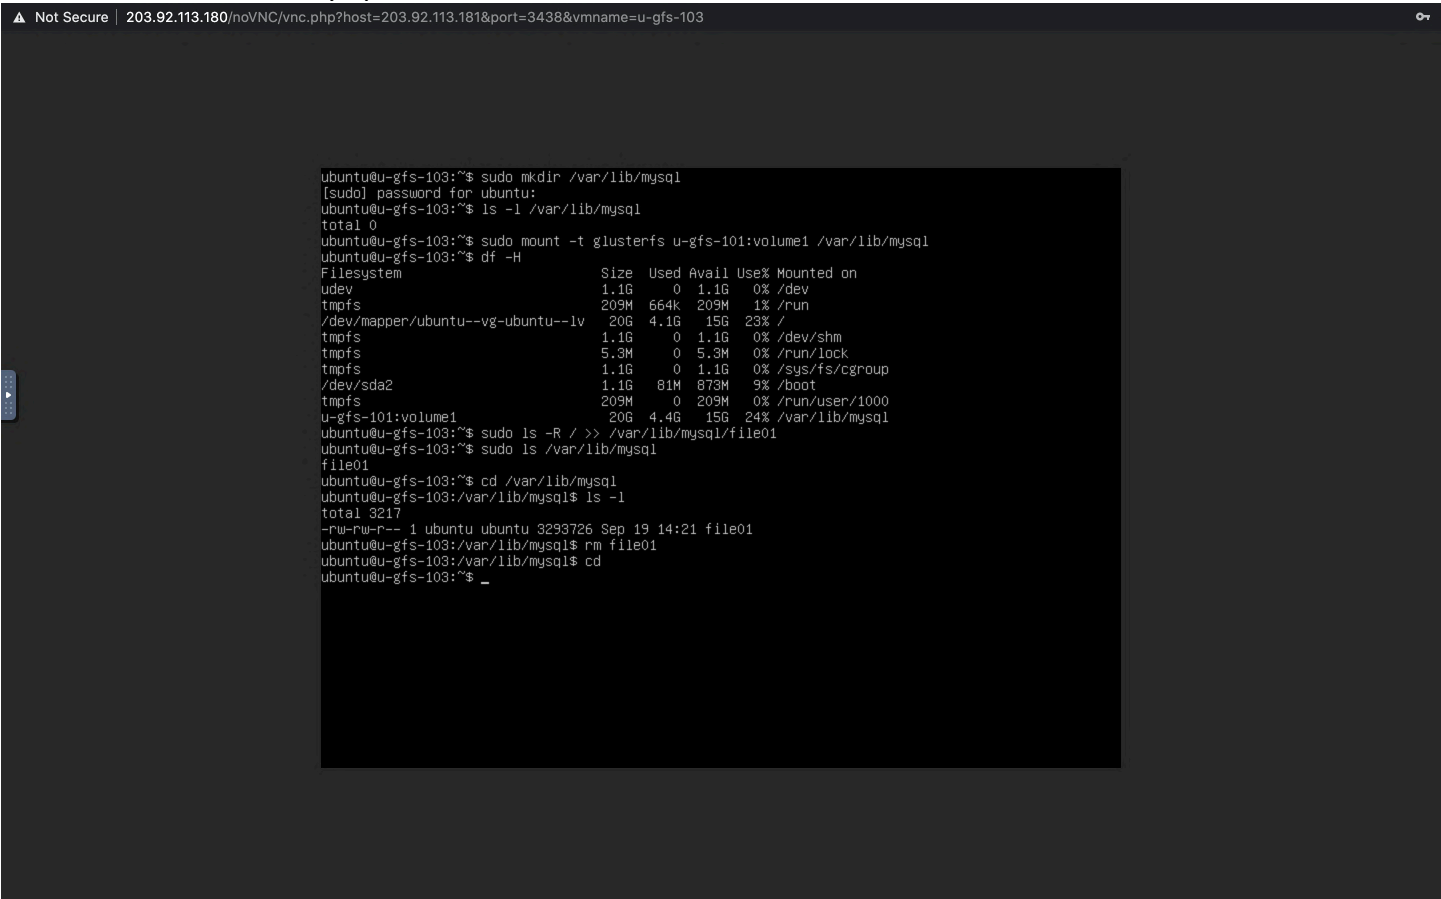
```
sudo mount -t glusterfs u-gfs-101:volume1 /var/lib/mysql
```

Checking storage size of mountpoint:

```
df -H
```

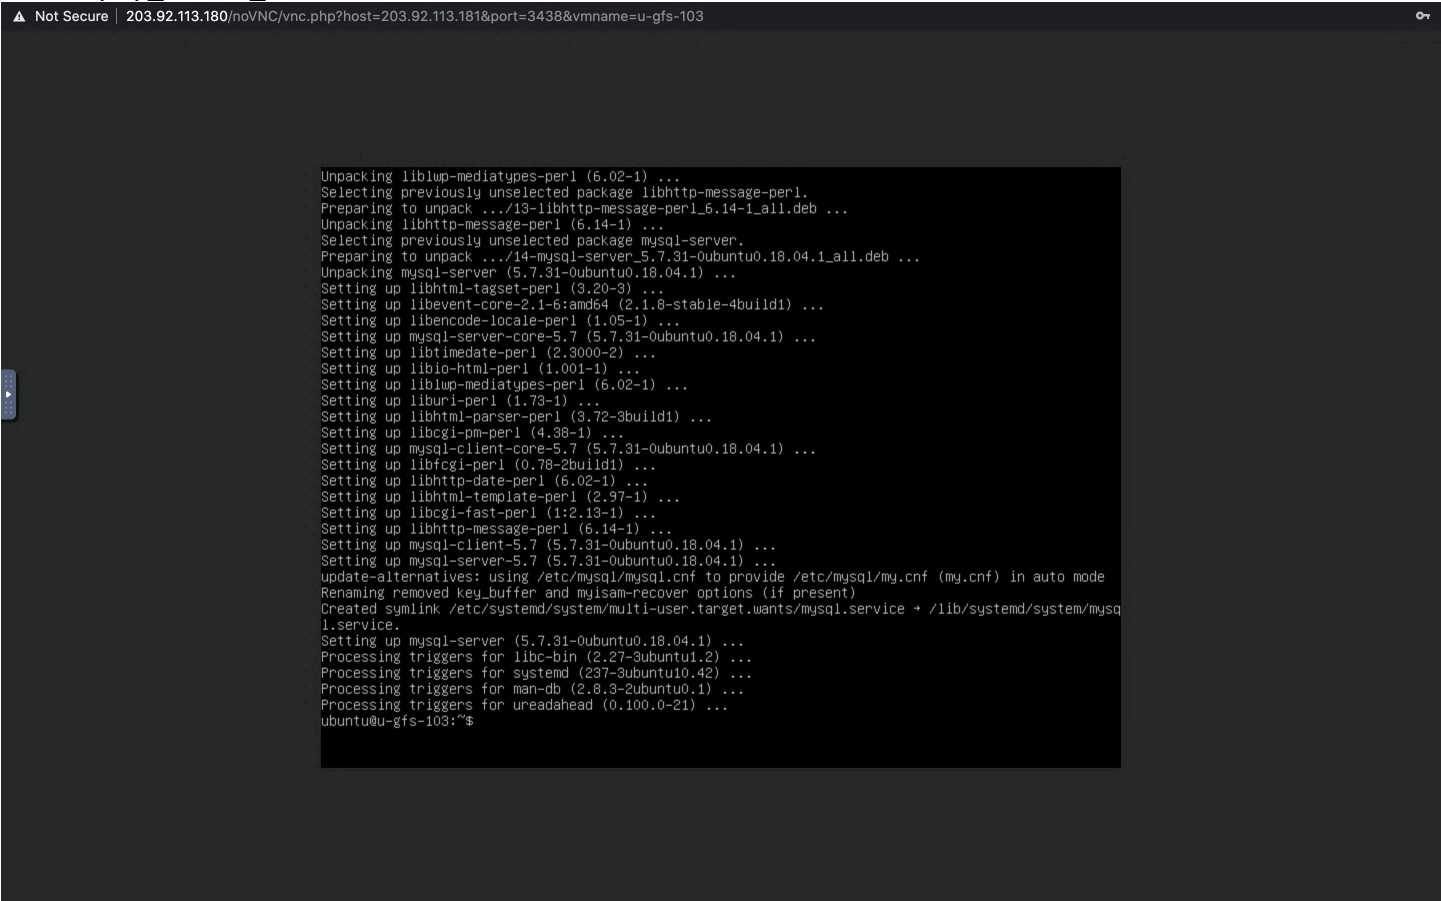
Testing /var/lib/mysql by adding a file to it:

```
sudo ls -R / >> /var/lib/mysql/file01
```

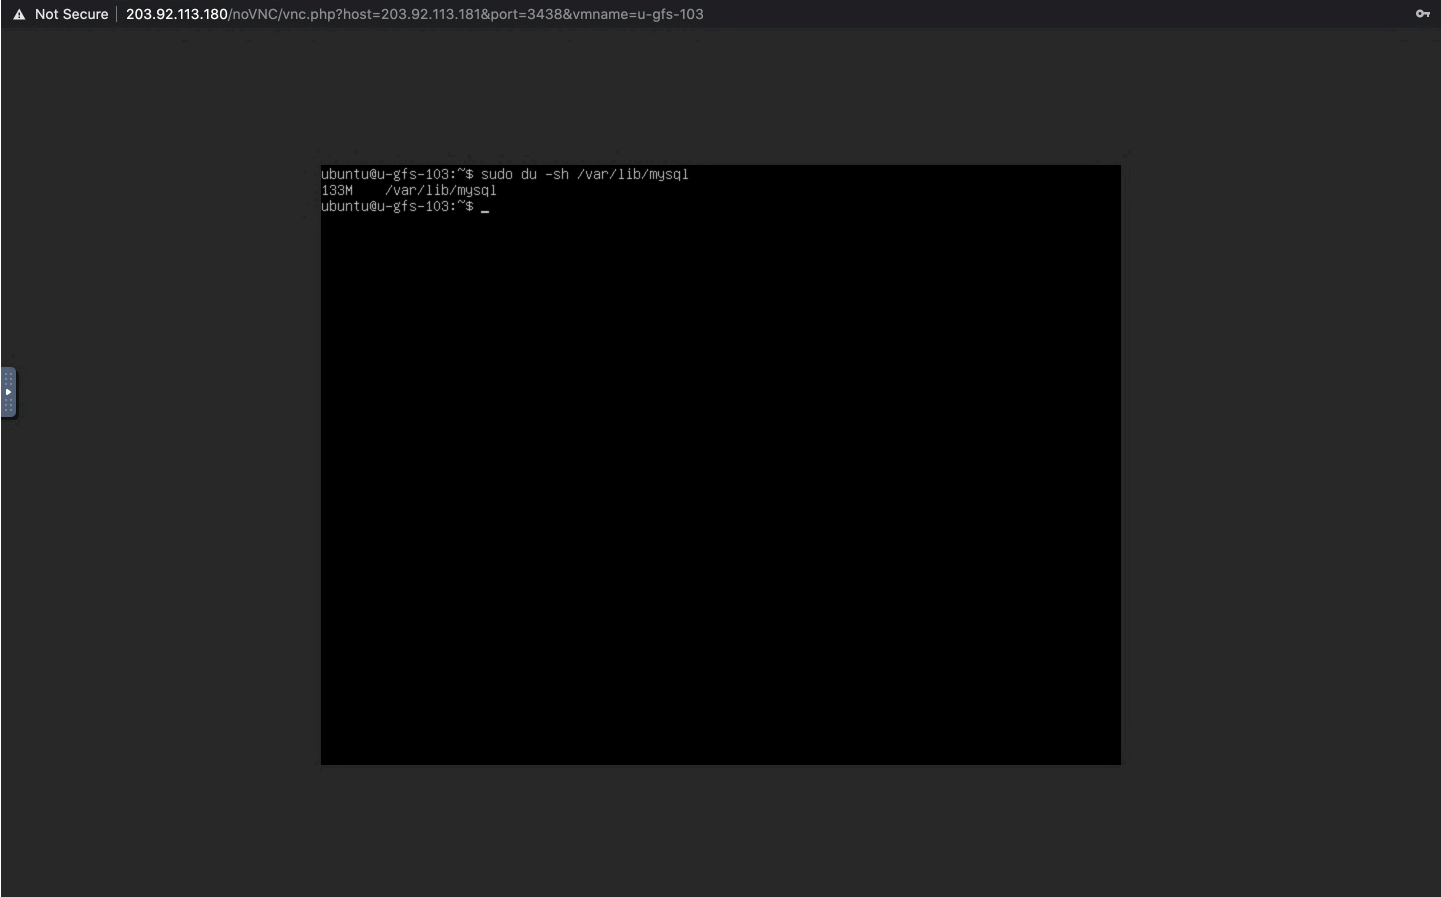


Installing MYSQL:

```
sudo apt update
sudo apt install mysql-server
sudo mysql_secure_installation
```



About 130MB of MYSQL files were installed into /var/lib/mysql:



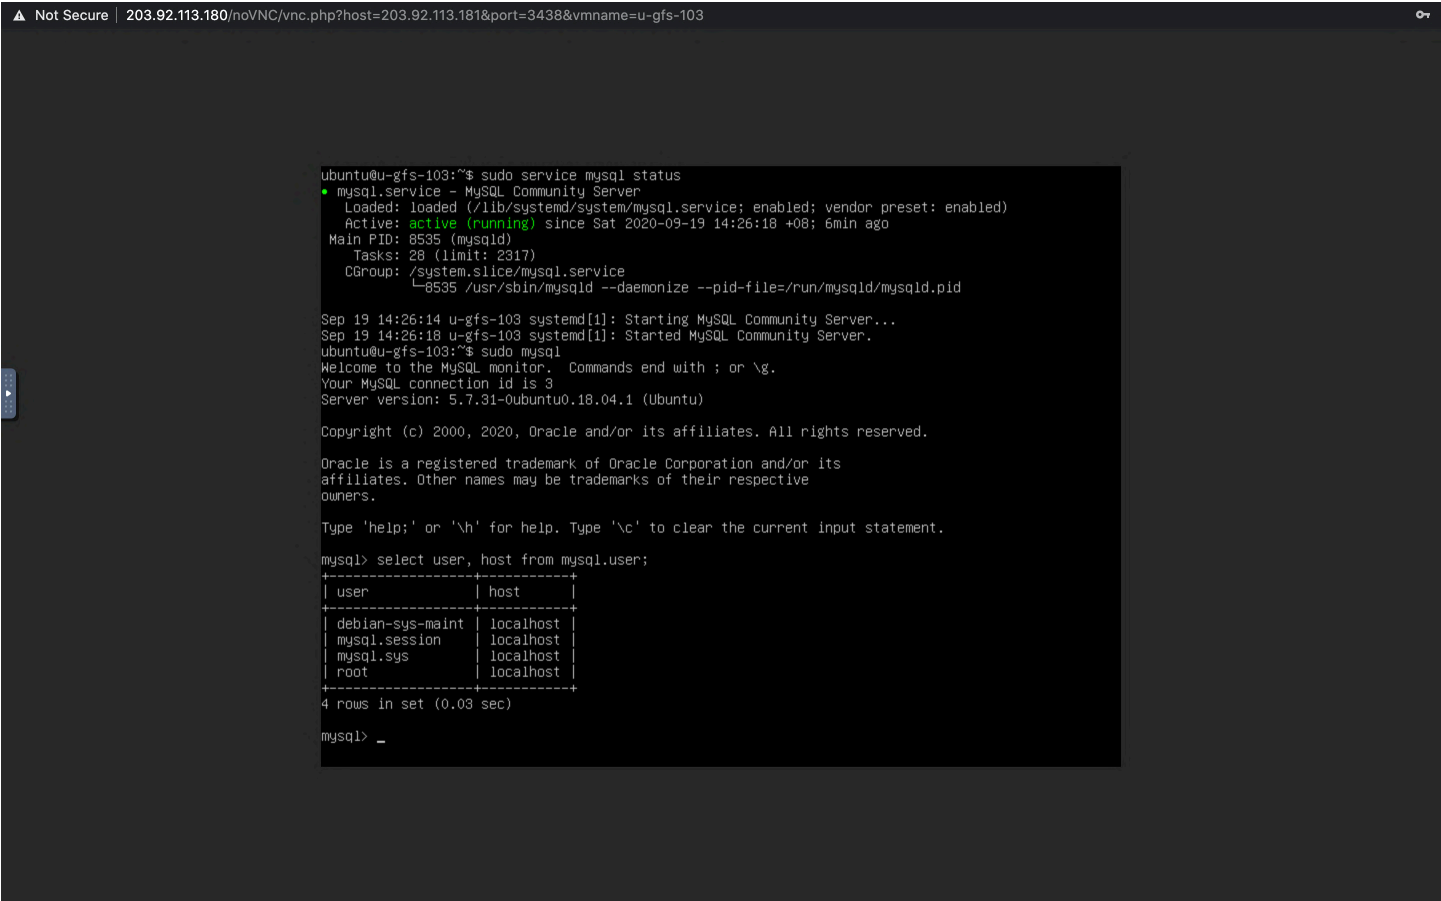
Checking MYSQL status:

sudo service mysql status
sudo mysqladmin -p -u root version

CREATING MYSQL USER, DATABASE, TABLES AND SCRIPTS

Running MYSQL commands to check database:

sudo mysql
SELECT user, plugin, host FROM mysql.user;



Creating MYSQL user:

sudo mysql
CREATE USER 'ubuntu' IDENTIFIED BY 'ubuntu';
GRANT ALL PRIVILEGES ON *.* TO 'ubuntu' WITH GRANT OPTION;

Check status of storage available and status of /var/lib/mysql

df -H
du -sh /var/lib/mysql
sudo ls -l /var/lib/mysql
sudo ls -l /var/lib/mysql/personmaster

Logging in as MYSQL user:

mysql -u ubuntu -pubuntu
CREATE DATABASE personmaster;
USE personmaster;
CREATE TABLE personinfo (passport varchar(20), fullname varchar(20), sex varchar(1), birthdate varchar(10));
quit;

Checking status of storage available and status of /var/lib/mysql

df -H

Checking files created by MYSQL:

sudo ls -l /var/lib/mysql
sudo ls -l /var/lib/mysql/personmaster
[a small file of about 100K called personinfo.ibd is created in /var/lib/mysql/personmaster]

INSERTING DATA INTO MYSQL

Creating a Unix script "insertmysql.sh" that calls MYSQL to insert records:

```
RANDOM=$((RANDOM%100000))
for i in $(seq 100000)
do
    rm insert.mysql

    echo use personmaster;" >> insert.mysql

    TMP1=$((RANDOM%100000))
    TMP2=$((RANDOM%100000))
    PASSPORT=$((TMP1*100000+TMP2))
    FULLNAME=$(cat /dev/urandom|tr -dc 'a-zA-Z0-9'|fold -w 18| head -n 1)
    SEX=$((RANDOM%2))
    DIFFYEAR=$((1999-1930+1))
    FINALEYEAR=$(( ($RANDOM%$DIFFYEAR)+1930 ))
    DIFFMTH=$((12-1+1))
    FINALMTH=$(( ($RANDOM%$DIFFMTH)+1 ))
    DIFFDAY=$((30-1+1))
    FINALDAY=$(( ($RANDOM%$DIFFDAY)+1 ))
    BIRTHDATE=$FINALEYEAR "-" $FINALMTH "-" $FINALDAY

    echo $PASSPORT $FULLNAME $SEX $BIRTHDATE
```

```
echo insert into personinfo ("passport, fullname, sex, birthdate")
values "("'"'$PASSPORT'"', '"'$FULLNAME'"', '"'$SEX'"',
'"'$BIRTHDATE'"'"')"";" >> insert.mysql

TMP1=$RANDOM
TMP2=$RANDOM
PASSPORT=P$TMP1$TMP2
FULLNAME=$(cat /dev/urandom|tr -dc 'a-zA-Z0-9'|fold -w 18| head -n 1)
SEX=F
DIFFYEAR=$((1999-1930+1))
FINALYEAR=$(( ($ ($RANDOM%$DIFFYEAR)) +1930))
DIFFMTH=$((12-1+1))
FINALMTH=$(( ($ ($RANDOM%$DIFFMTH)) +1))
DIFFDAY=$((30-1+1))
FINALDAY=$(( ($ ($RANDOM%$DIFFDAY)) +1))
BIRTHDATE=$FINALYEAR "-" $FINALMTH "-" $FINALDAY

echo $PASSPORT $FULLNAME $SEX $BIRTHDATE
echo insert into personinfo ("passport, fullname, sex, birthdate")
values "("'"'$PASSPORT'"', '"'$FULLNAME'"', '"'$SEX'"',
'"'$BIRTHDATE'"'"')"";" >> insert.mysql

echo "quit" >> insert.mysql

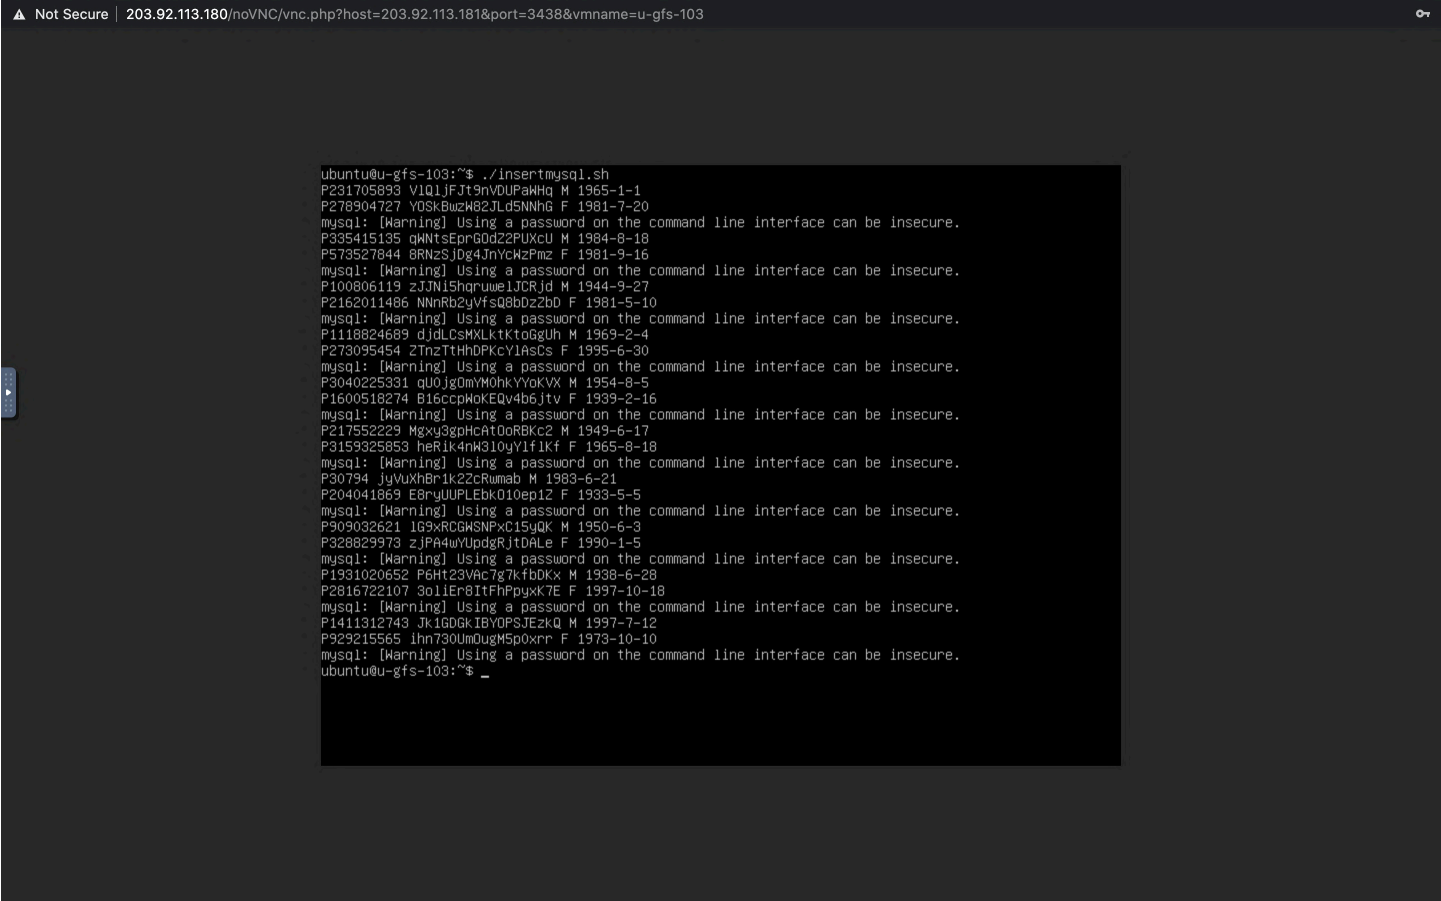
mysql -u ubuntu -pubuntu < insert.mysql

sleep 0
done
```

Running Unix script "insertmysql.sh" that calls MYSQL to insert records:

./insertmysql.sh

[because MYSQL is launched with stdin from a file containing SQL statements, there is a warning message]

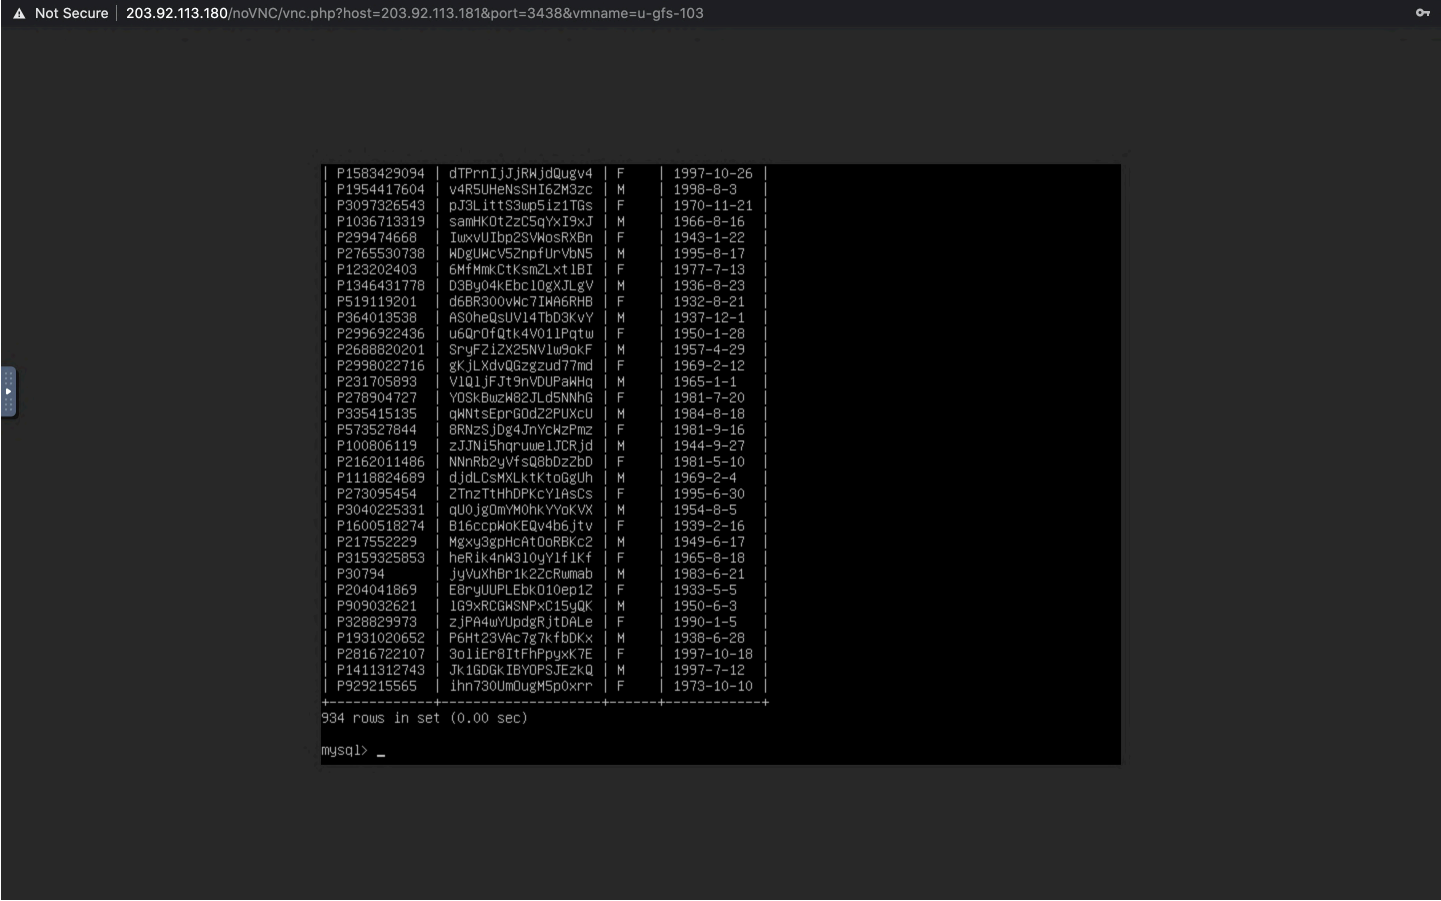


Verifying data is inserted into MYSQL:

mysql -u ubuntu -pubuntu

USE personmaster;

SELECT * from personinfo;



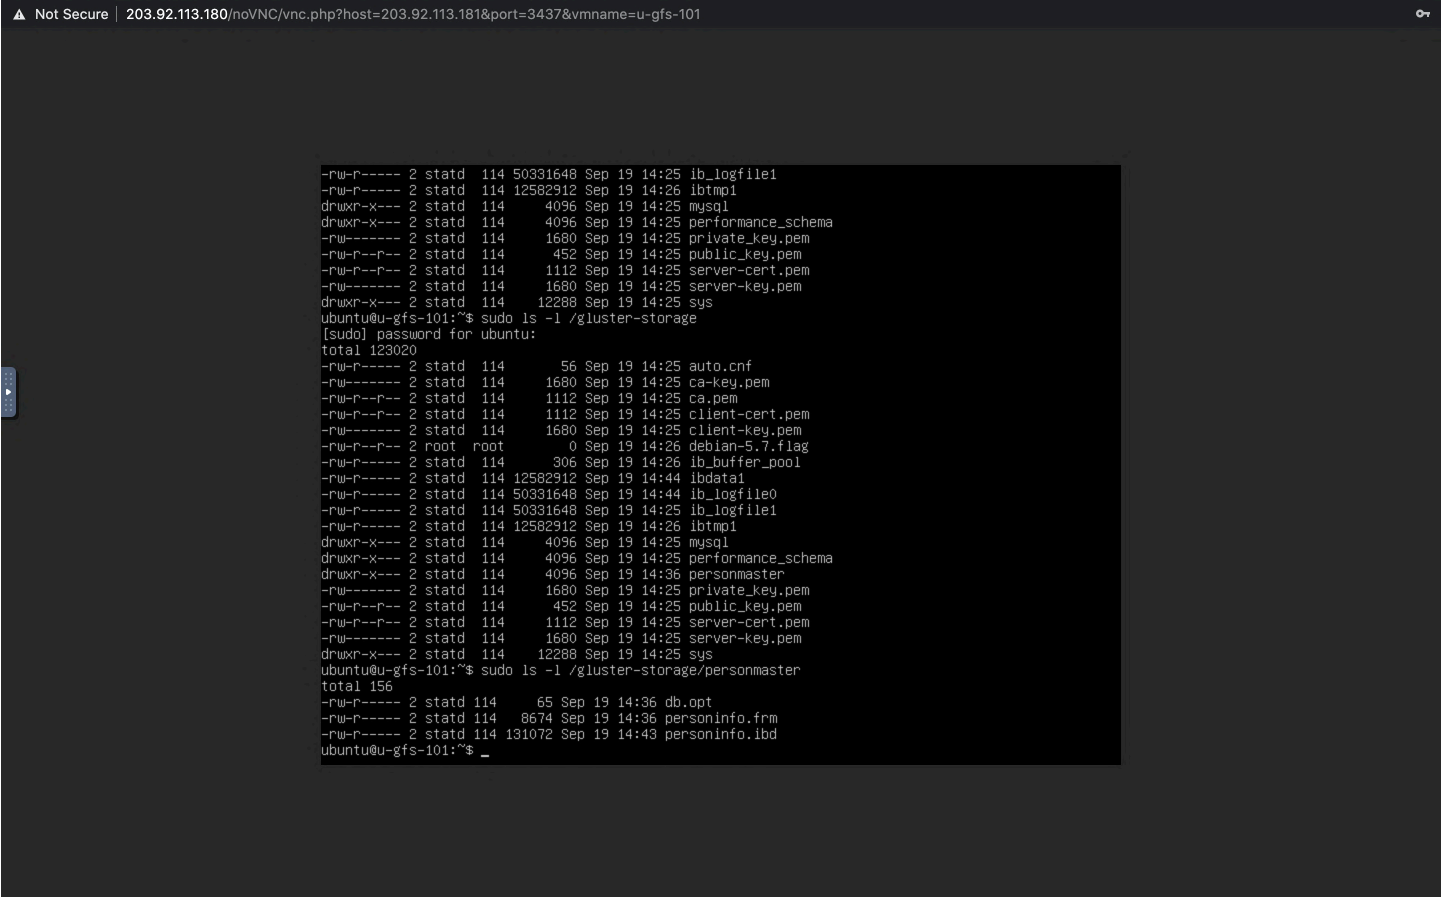
Check status of storage available and status of /var/lib/mysql

df -H

sudo ls -l /var/lib/mysql

sudo ls -l /var/lib/mysql/personmaster

[the file personinfo.ibd grows rapidly in /var/lib/mysql/personmaster]



Reviewing important SQL statements:

mysql -u ubuntu -pubuntu
USE personmaster;
INSERT INTO personinfo (passport, fullname, sex, birthdate) values ('1', '1', '1', '1');
SELECT * from personinfo;
DELETE from personinfo;

Reviewing important MYSQL statements:

Starting MYSQL:

sudo service mysql start
sudo service mysql status

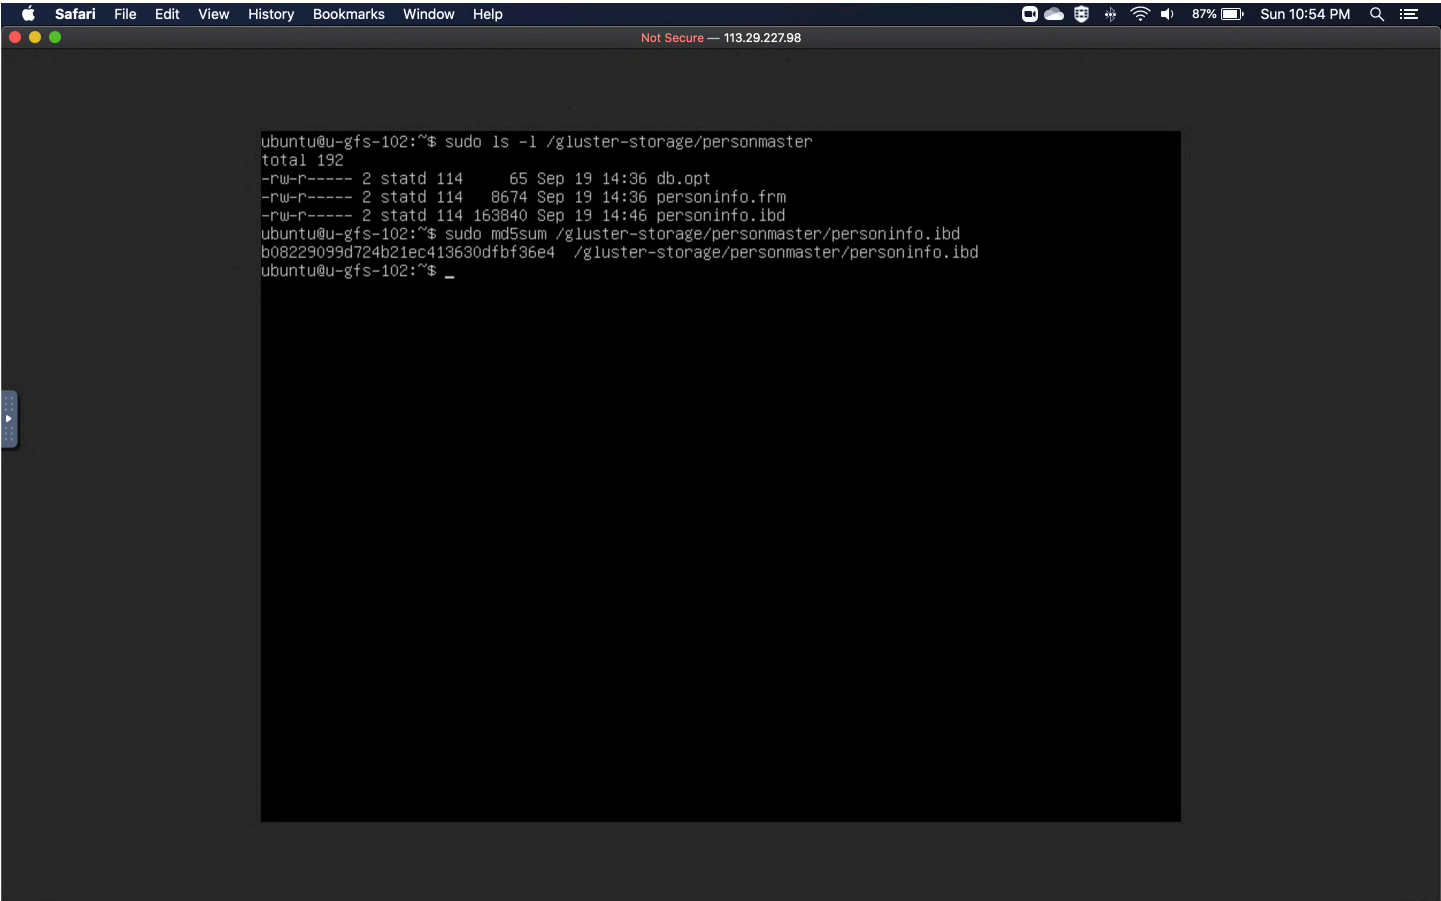

Stopping MYSQL:

sudo service mysql stop
sudo service mysql status

Reviewing important GLUSTER / MYSQL interactions:

Should you shutdown u-gfs-103, u-gfs-101 and u-gfs102 and start them up again, is important to start up u-gfs-101 and u-gfs-102 first, wait for Gluster to complete startup before starting u-gfs-103. When u-gfs-103 starts, you might need to remount /var/lib/mysql, and start MYSQL

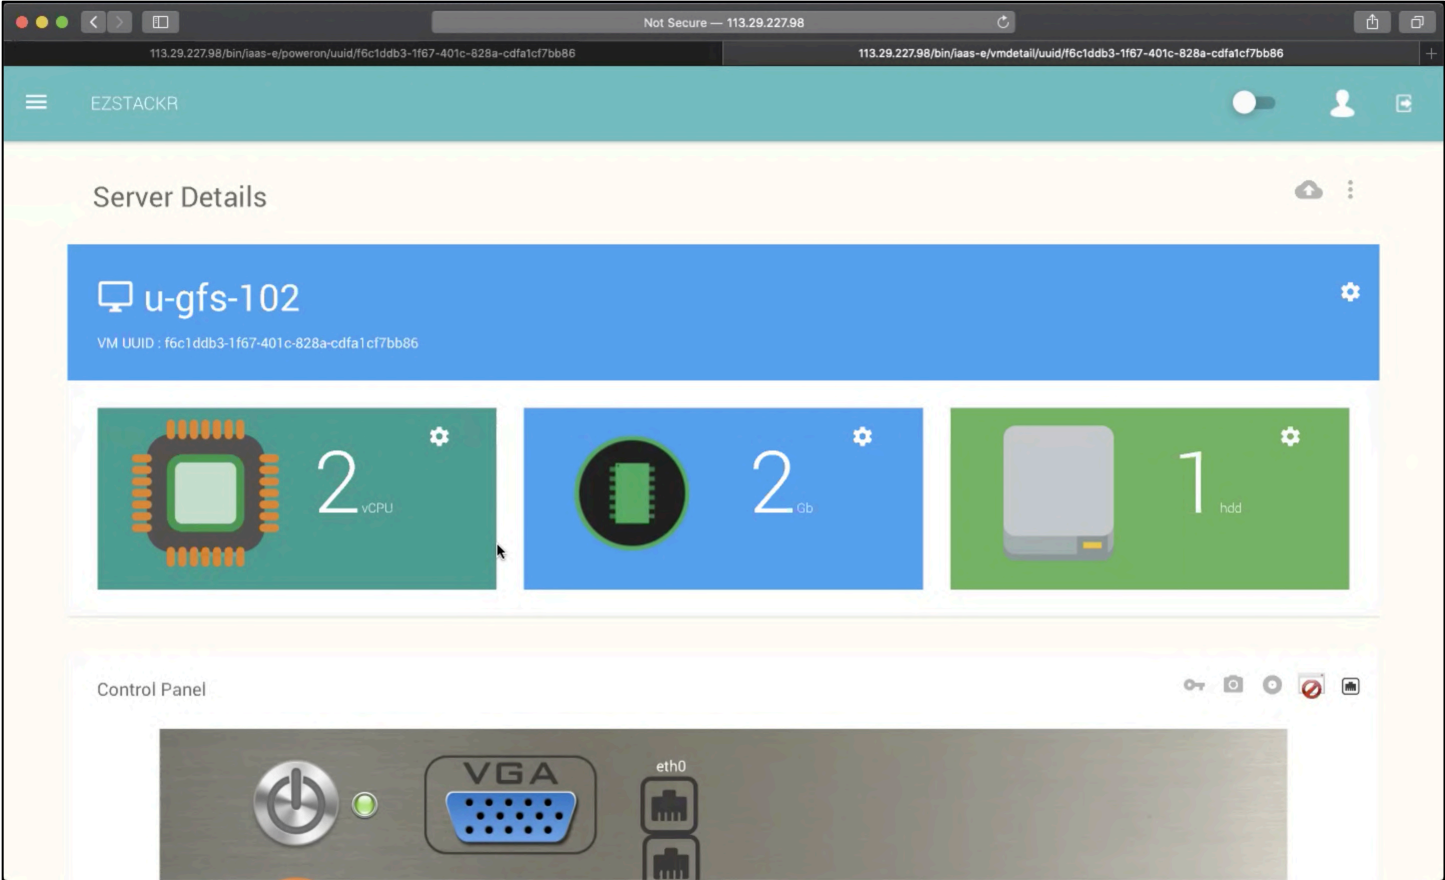
VERIFYING THE INTEGRITY OF THE MYSQL DATA WRITTEN TO GLUSTER STORAGE MIRROR POOL



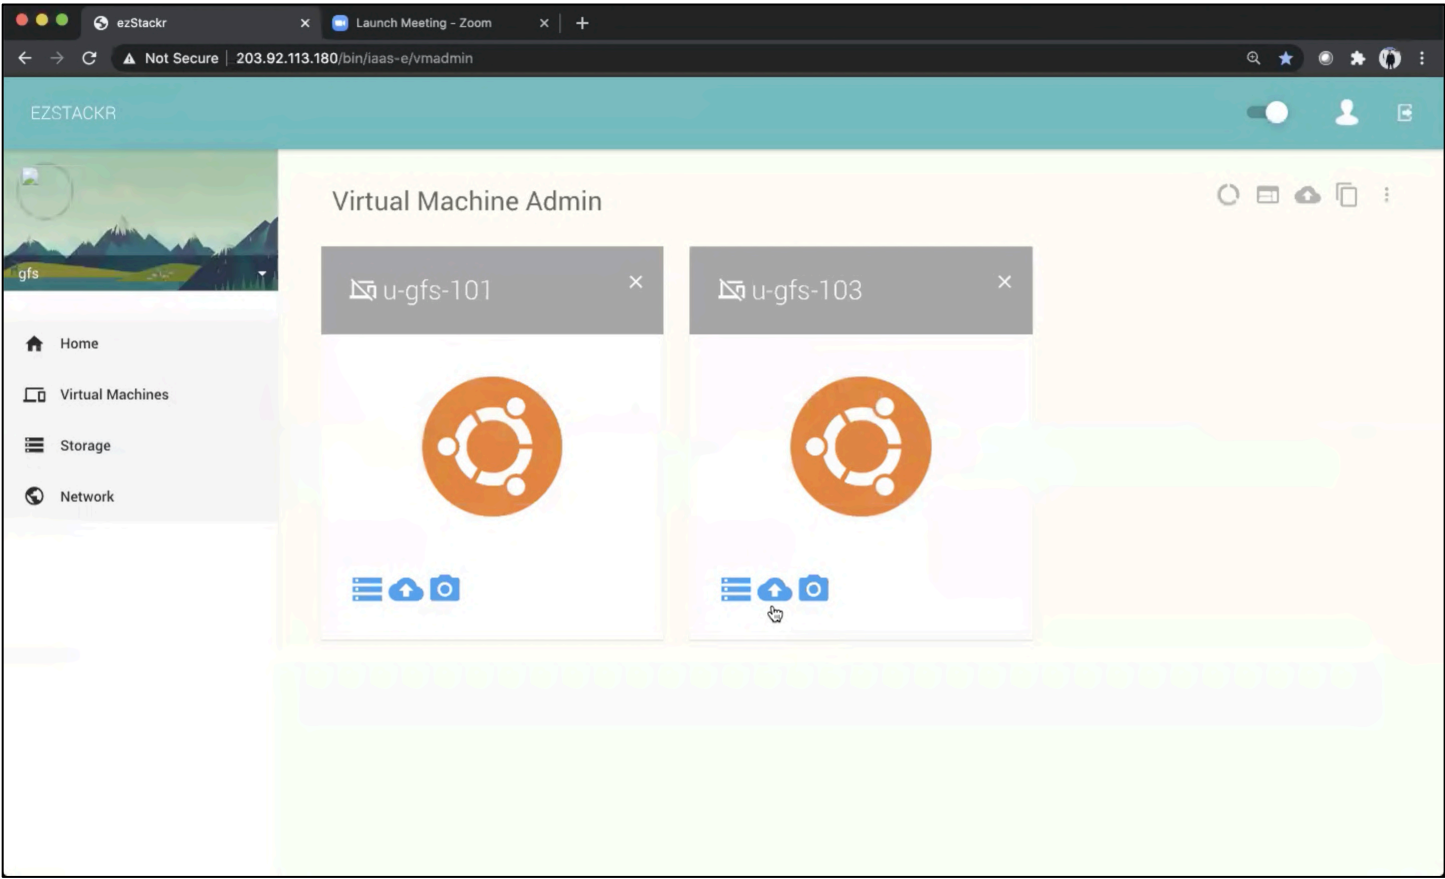
DESTROYING MYSQL SERVER AND DESTROYING FIRST GLUSTER NODE

Review u-gfs-102:

Let’s keep this VM alive

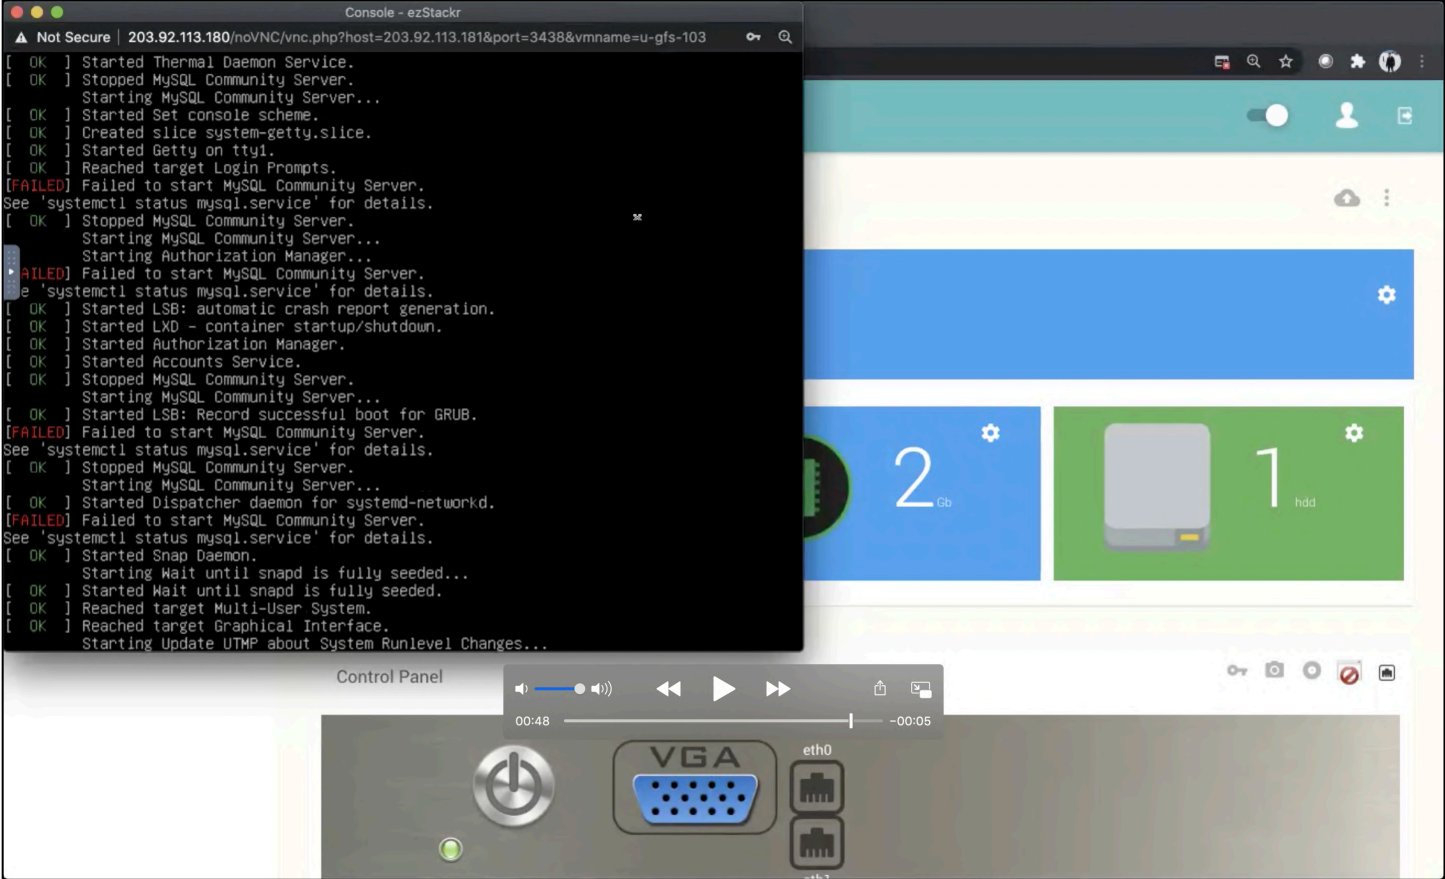


Review u-gfs-101, u-gfs-103:
Let's shutdown these 2 VMs



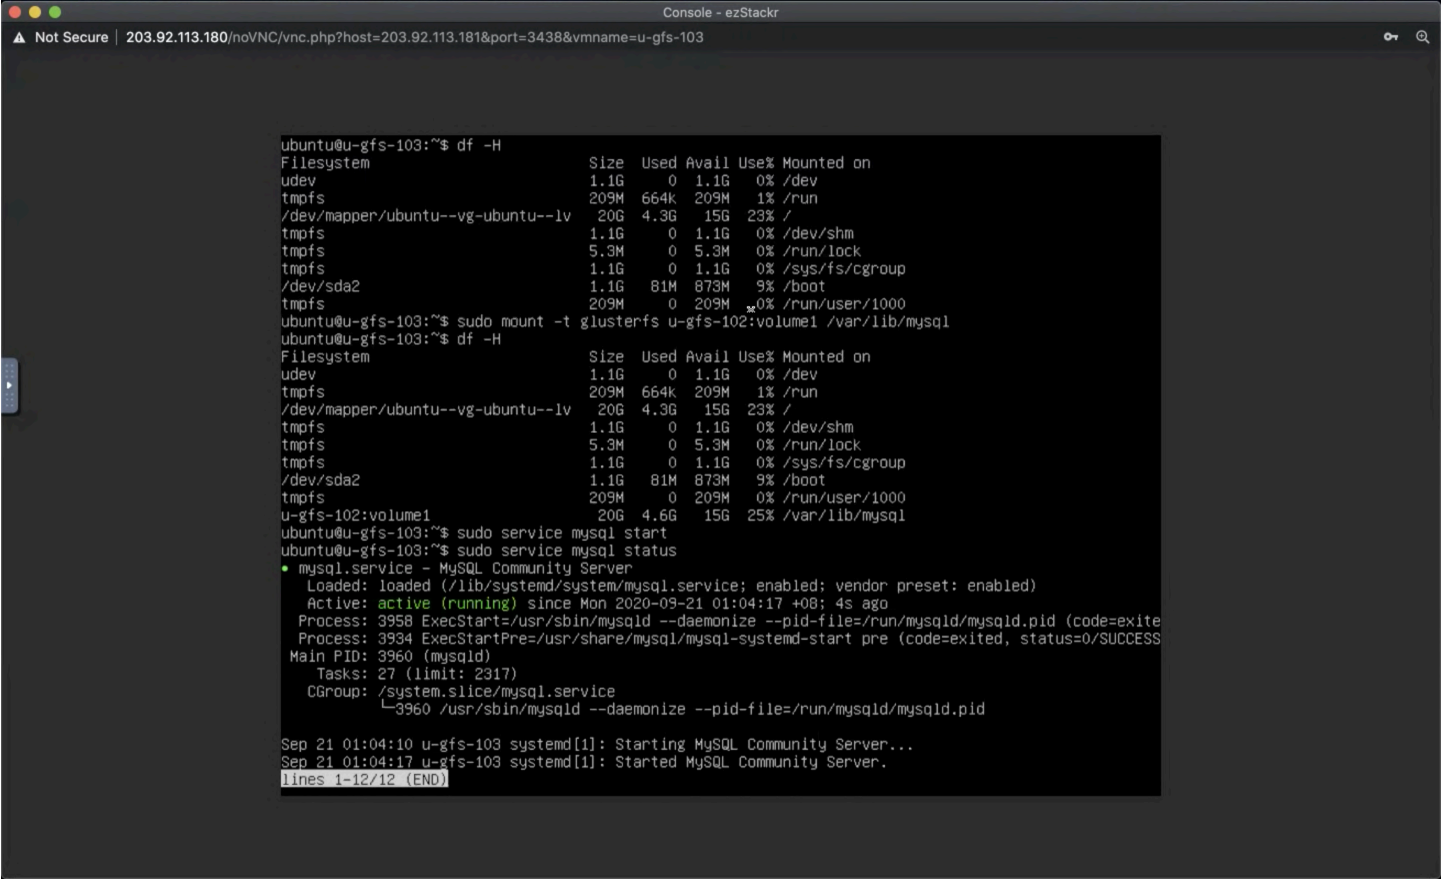
RESTORING MYSQL DATA FROM SECOND GLUSTER NODE

Review u-gfs-103:
While u-gfs-101 is down, launch u-gfs-103. At startup, MYSQL fails to start as database is not mounted

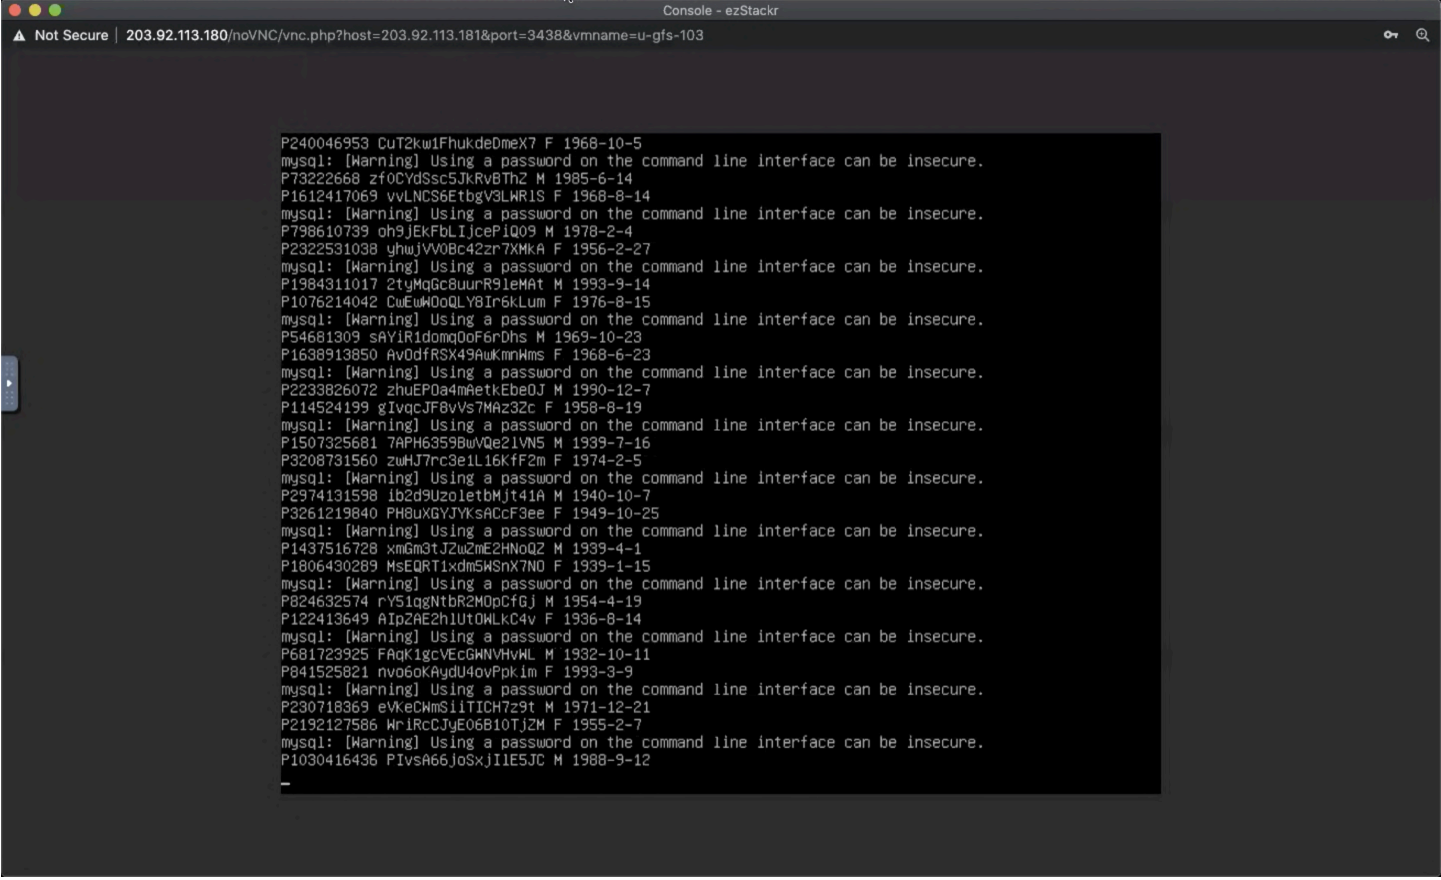


Mount the Gluster volume on u-gfs-102 (this is the surviving Gluster node):
sudo mount -t glusterfs u-gfs-102:volume1 /var/lib/mysql

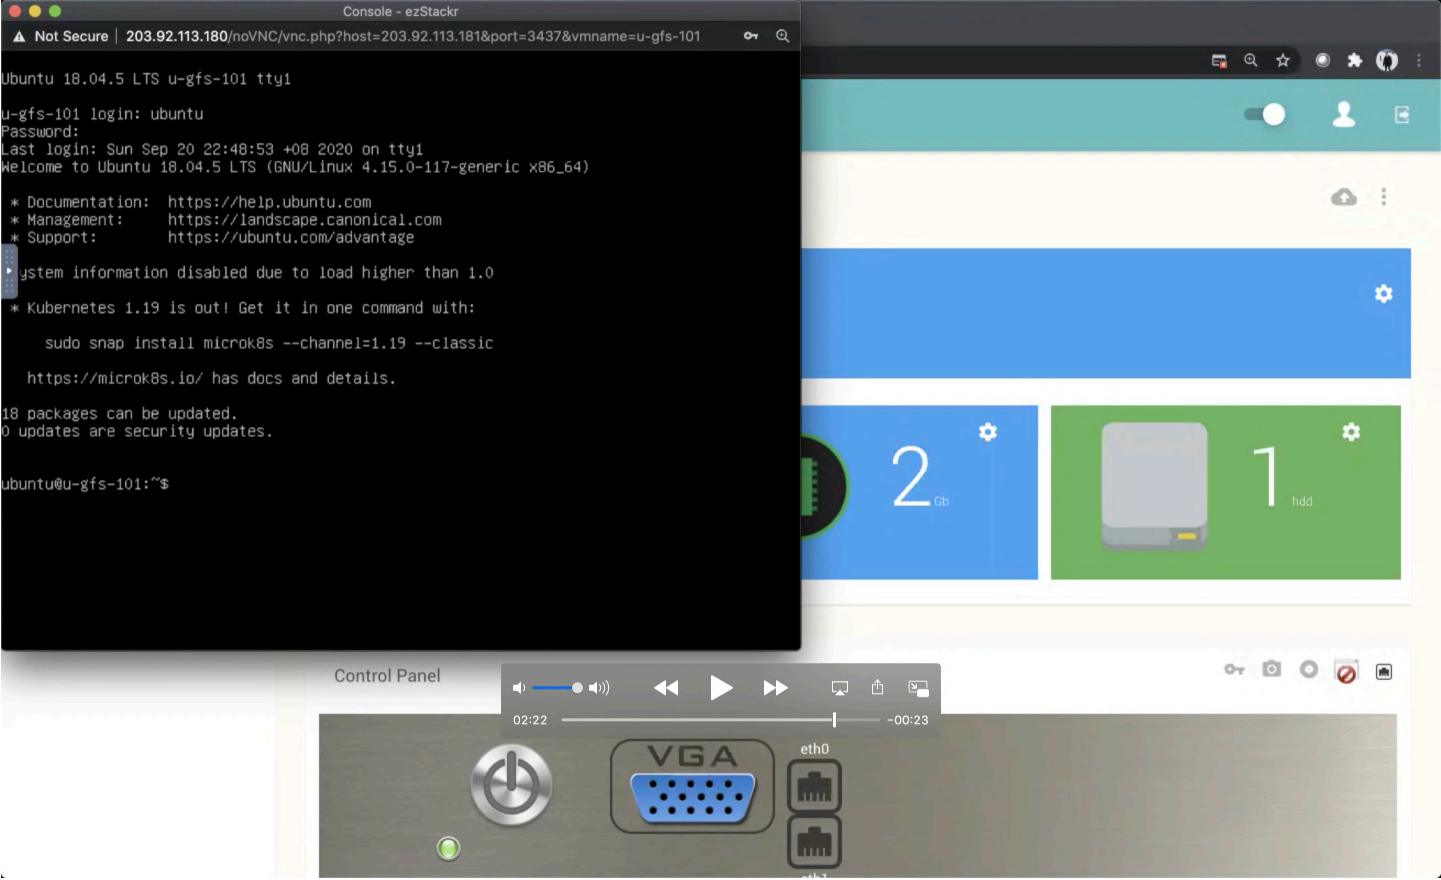
Starting MYSQL:
sudo service mysql start
sudo service mysql status



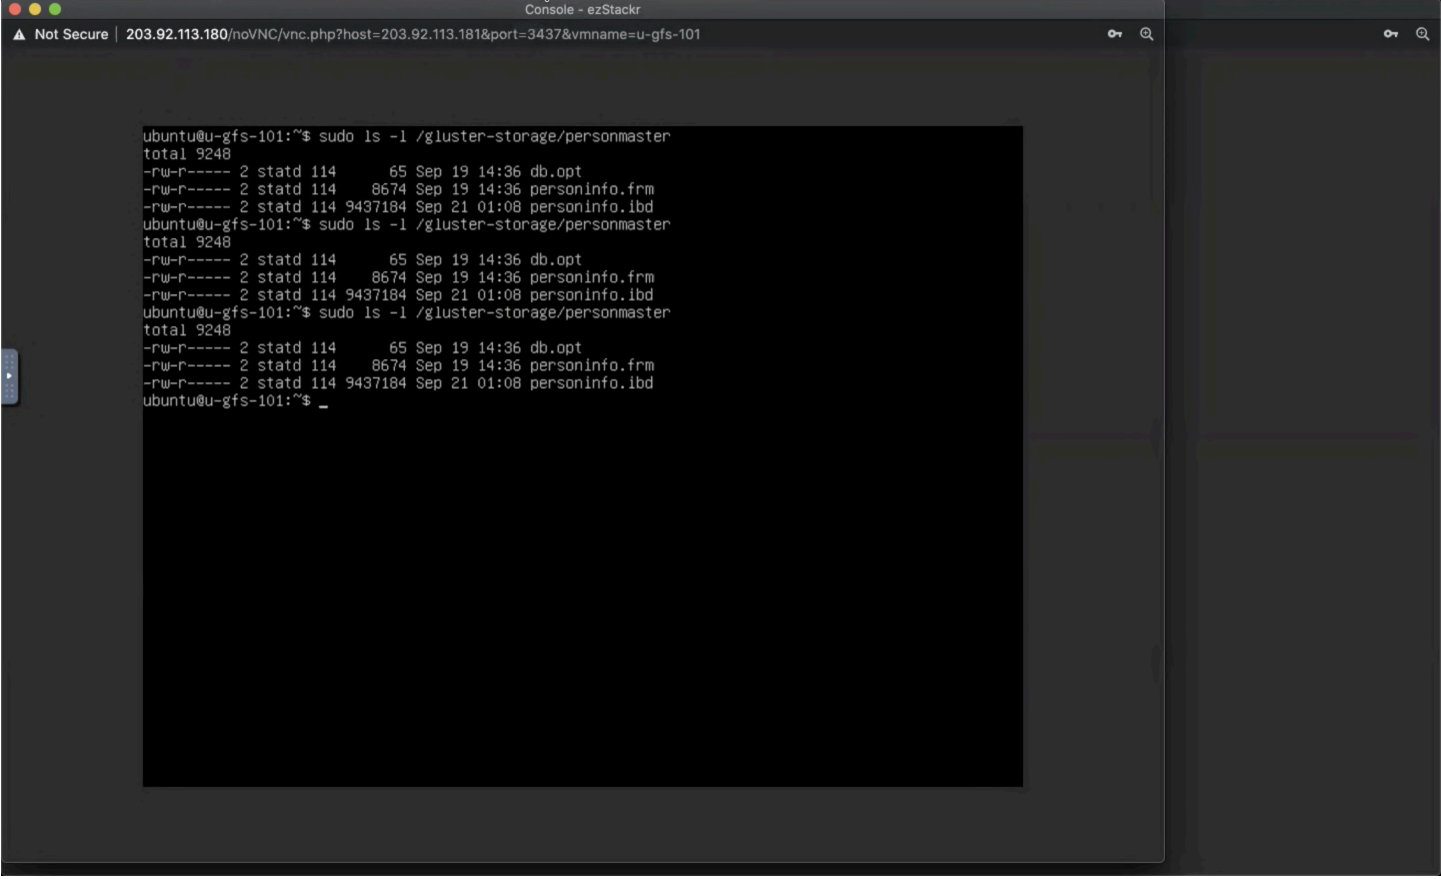
Running Unix script "insertmysql.sh" that calls MYSQL to insert records (records inserted into u-gfs-102):
./insertmysql.sh
[because MYSQL is launched with stdin from a file containing SQL statements, there is a warning message]



OBSERVING FIRST GLUSTER AFTER IT IS REPAIRED AND REJOINS GLUSTER CLUSTER
Relaunching u-gfs-101 (this Gluster mirror node will rejoin the cluster when server is up):
Logging into u-gfs-101

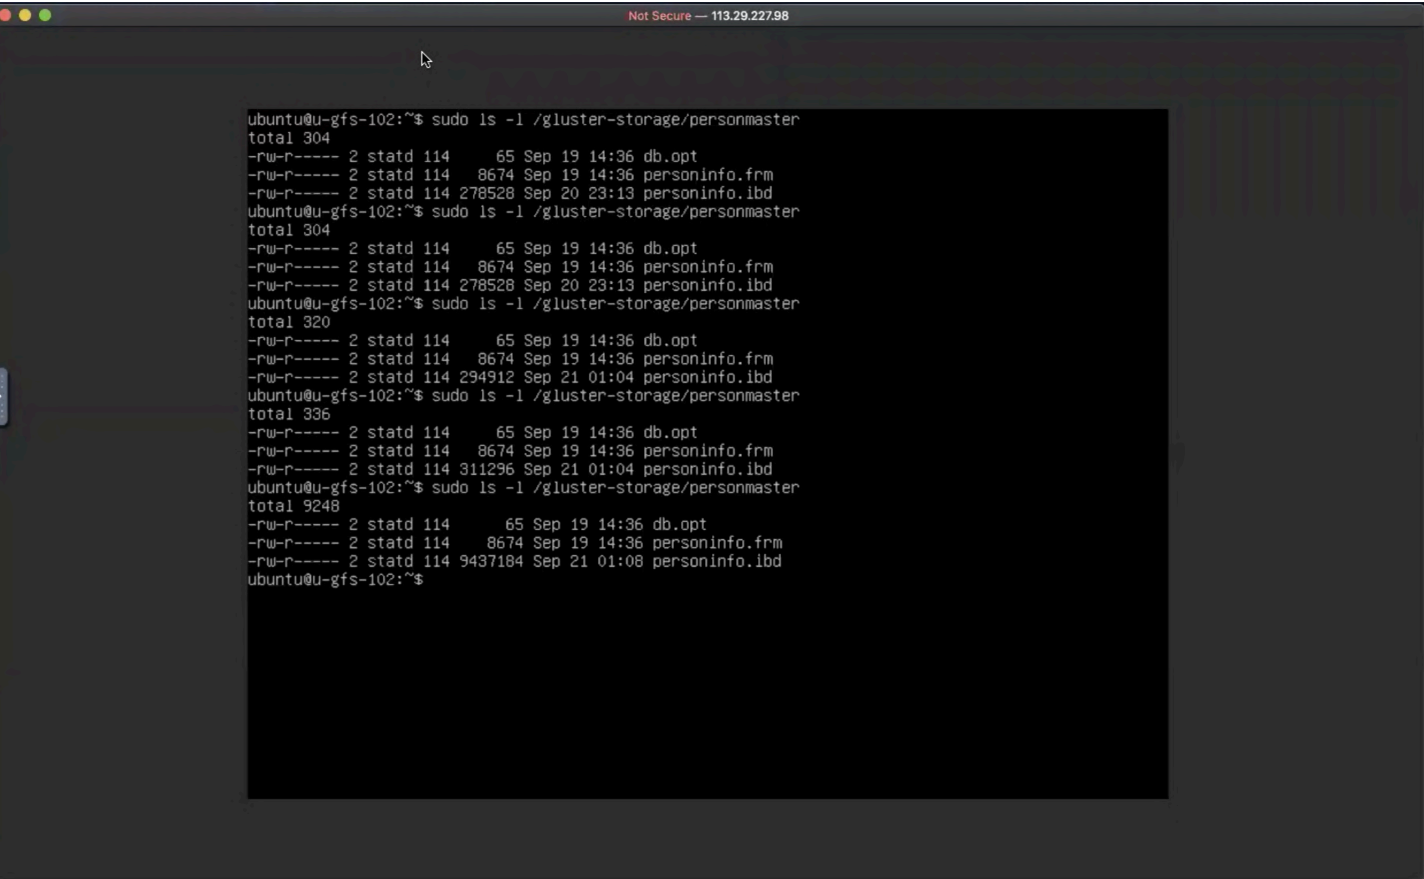


Reviewing u-gfs-101 after it has restarted (MYSQL database personinfo.ibd has been updated):
ls -l /gluster-storage/personmaster




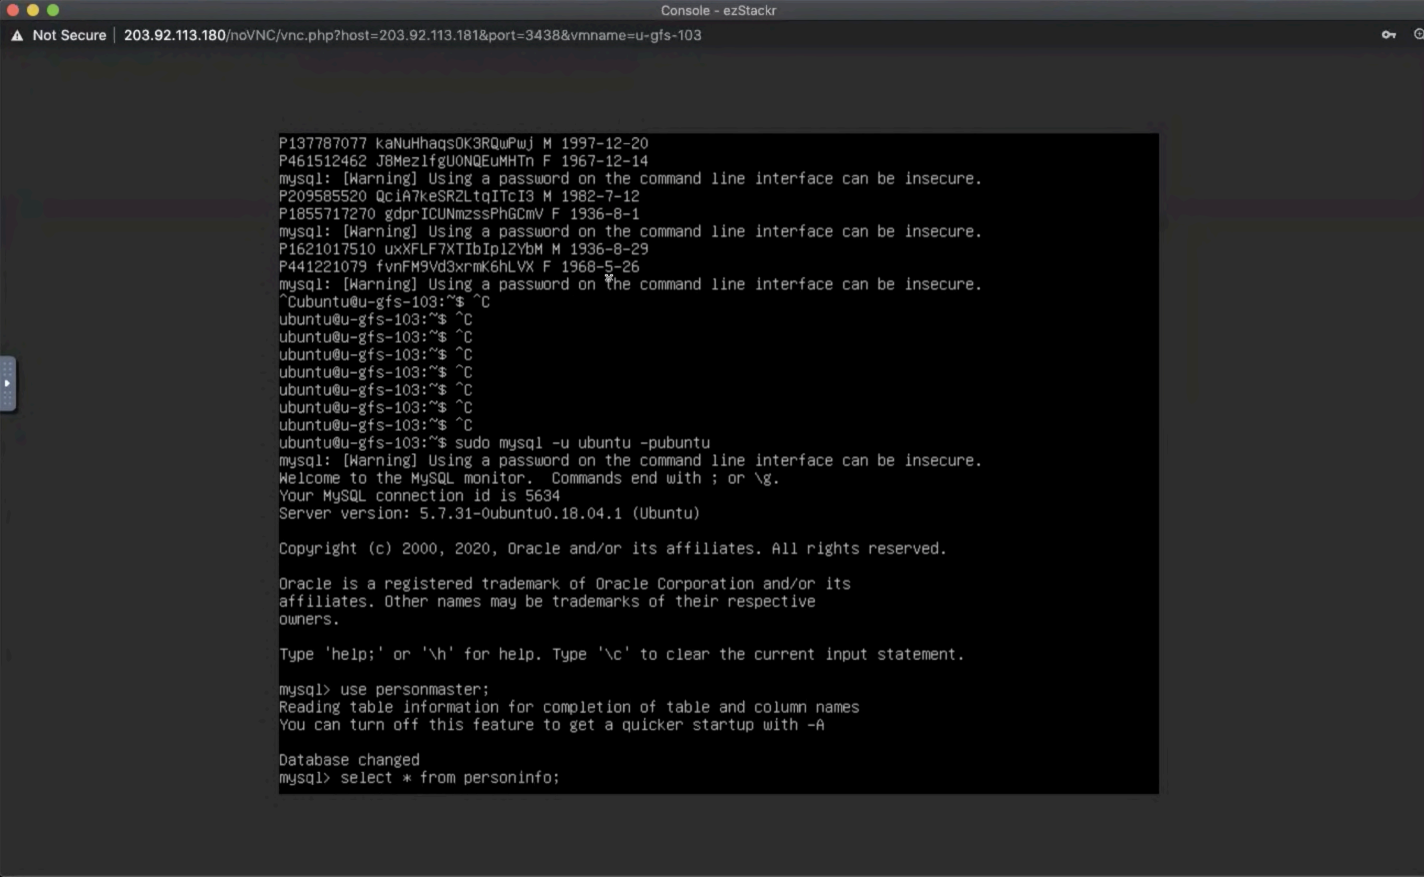
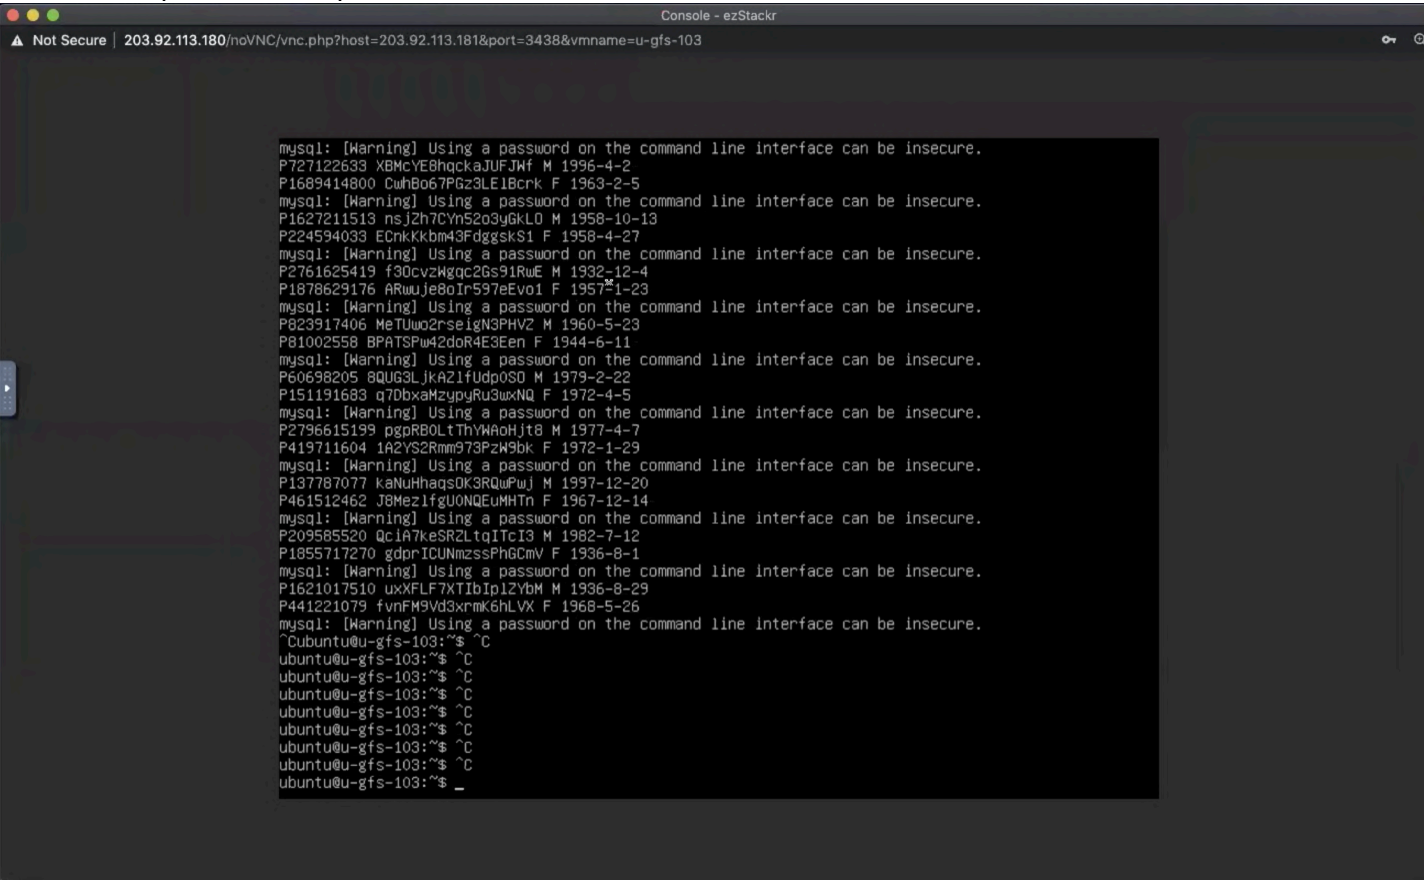
Reviewing u-gfs-102 (MYSQL database personinfo.ibd has also been updated):

ls -l /gluster-storage/personmaster



Stopping database inserts and logging into MYSQL to review the data in database:

Ctrl^C to stop the SQL script



END OF TEST

This concludes a successful storage mirror test.

Contact Us

LGA Telecom Pte Ltd
33 Ubi Avenue 3
#08-53 Vertex (Tower A)
Singapore 408868

Tel (65) 6892 2308
Email: sales@lgatelecom.net
Website: www.lgatelecom.net

© 2020 LGA Telecom Pte Ltd. All rights reserved.