

XCAT 2 MySQL Setup

03/28/10, 08:12:07 AM

1.0 Switch to MySQL on the Management Node.....	1
1.1 Install MySQL.....	1
1.2 Configure and migrate xCAT to MySQL using mysqlsetup script.....	3
1.3 Configure MySQL manually.....	3
1.4 Migrate xCAT data to MySQL.....	8
1.5 Add ODBC support.....	9
1.6 Removing MySQL xcatd database.....	13
1.7 Migrate to new level MySQL.....	14
1.8 Diagnostics.....	15
1.9 References.....	15

1.0 Switch to MySQL on the Management Node

MySQL is supported on xCAT 2.1 or later. One reason to migrate from the default SQLite database to MySQL with xCAT is for xCAT hierarchy using Service Nodes. MySQL provides the ability for remote access to the xCAT database on the Management node which is required by Service Nodes. Refer to the xCAT Service Node documentation for more information. Other programs or scenarios within your environment may also benefit from or require MySQL. This document contains steps to install MySQL, configure the server, create a database and populate it with your xCAT data.

Before using this document, you should have a general understanding of MySQL. If necessary, review the installation and tutorial sections of the MySQL 5.1 Reference Manual located at:

<http://dev.mysql.com/doc/refman/5.1/en/index.html>

Although this should be a one time setup, the MySQL documentation should also be reviewed so that you can manage and maintain the MySQL environment.

1.1 Install MySQL

Before you install verify you have enough free space. Increase if needed. On AIX, the /usr file system will be installed with MySQL, which is quite large. Check the size of the rpm (currently > 500meg). During the install you will need at least twice

the size of the rpm available, it will shrink after the gzip file is removed from the /usr directory

On AIX: The xCAT RPM called *xcat-mysql* is provided to help simplify the installation of MySQL on an AIX system. Download the *xcat-mysql-2.*.gz*. From this website:

<https://sourceforge.net/projects/xcat/files/>

Unzip and untar it in the location of your choice. The *xcat-mysql* post processing will automatically unwrap MySQL in the */usr/local* directory and will create a link for */usr/local/mysql*. It will also update the PATH environment variable in the */etc/profile* file.

```
gunzip xcat-mysql-2*.gz
tar -xvf xcat-mysql-2*.tar
./instmysql
```

On Redhat/Fedora: MySQL comes as part of the OS. Ensure that the following rpms are installed on your Management Node:

```
perl-DBD-MySQL*
mysql-server-5.*
mysql-5.*
mysql-devel-5.*
mysql-bench-5.*
mysql-5.*
mysql-connector-odbc-*
mysql-devel-5.*
```

On SLES: MySQL comes as part of the OS. Ensure that the following rpms are installed on your Management Node:

```
mysql-client-5*
libmysqlclient_r15*
libqt4-sql-mysql-4*
libmysqlclient15-5*
perl-DBD-mysql-4*
mysql-5*
```

1.2 Configure and migrate xCAT to MySQL using mysqlsetup script

As of xCAT 2.3.1, you can use mysqlsetup to perform all the operations in the Configure MySQL manually and the Migrate xCAT data to MySQL sections of the doc. See the manpage for mysqlsetup.

Note: after you setup MySQL on the Management node, you will need to update the database base to give access to your service nodes when they are defined. See mysqlsetup -u -f option. Also you can review the “GRANT ALL ” command below in the manual process, which is what the mysqlsetup script uses.

1.3 Configure MySQL manually

Stop: Did you notice as of xCAT 2.3.1, you can use the mysqlsetup script provided by xCAT to automatically accomplish all the following manual steps. See man mysqlsetup.

This section takes you through setting up the the MySQL environment, starting the server and connecting to the interactive program to create server definitions and perform queries.

This example assumes:

- Management Node: *mn20*
- xCAT database name: *xcatdb*
- Database user id used by xCAT for access: *xcatadmin*
- Database password for xcatadmin: *xcat201*

Substitute your addresses and desired database administration, password and database name as appropriate.

All of the following steps should be run logged into the Management Node as root.

1. Create a user id and group that will run the MySQL server.

On AIX, create the mysql user and groups, that will be used to run mysql:

```
mkgroup mysql  
mkuser pgrp=mysql mysql
```

Additionally on AIX, update the mysql file permissions:

```
cd /usr/local/mysql
chown -R mysql .
chgrp -R mysql .
```

On Linux: The mysql user id and group already exists, and the permissions are already correct when MySQL is installed.

2. Using the mysql userid, execute the script that will create the MySQL data directory and initialize the grant tables.

```
On AIX: /usr/local/mysql/scripts/mysql_install_db --user=mysql
On Linux: /usr/bin/mysql_install_db --user=mysql
```

3. **Edit server configuration file so that it runs in ANSI-QUOTES mode.**

```
On AIX: cp /usr/local/mysql/support-files/my-large.cnf /etc/my.cnf
```

Then on both AIX and Linux under the section:

```
# The MySQL server
[mysqld]
```

add the following line:

```
sql_mode=ANSI_QUOTES
```

Note: IF you are at xCAT 2.4 or later sql_mode=ANSI_QUOTES is no longer required.

4. For large systems you may also need to increase *max_connections* to the database in the *my.cnf* file. The default is 100. Add this line to the configuration file after “*sql_mode=ANSI_QUOTES*”.

```
max_connections=300
```

For additional system variables that can be set go to:

http://dev.mysql.com/doc/refman/5.1/en/server-system-variables.html#sysvar_max_connections

5. On AIX, update the data directory ownership permission to mysql. All other mysql install directories should be owned by root.

```
cd /usr/local/mysql
```

```
chown -R root .
chown -R mysql data
```

6. Check ulimit settings.

```
ulimit -a
time(seconds) unlimited
file(blocks) 2097151
data(kbytes) 131072
stack(kbytes) 32768
memory(kbytes) 32768
coredump(blocks) 2097151
nofiles(descriptors) 2000
threads(per process) unlimited
```

If not unlimited, change the ulimit setting to unlimited, for this session. coredump is optional.

```
ulimit -m unlimited
ulimit -n unlimited
ulimit -d unlimited
ulimit -f unlimited
ulimit -s unlimited
ulimit -t unlimited
```

Edit the `/etc/security/limits` file, to make these limits stay unlimited for root through reboot or the start of the MySQL server may fail.

root:

```
fsize = -1
core = -1
cpu = -1
data = -1
rss = -1
stack = -1
nofiles = -1
```

7. **Start the MySQL server** (running as root must use the `--user` option).

On AIX: `/usr/local/mysql/bin/mysqld_safe --user=mysql &`
(may need to hit enter to get prompt back)

On Linux:

`/usr/bin/mysqld_safe --user=mysql &`

or

service mysqld start (on sles service mysql start)

If you need to stop the MySQL server:

Note the mysql root id must have been setup, see below.

On AIX: */usr/local/mysql/bin/mysqladmin -u root -p shutdown*

On Linux:

/usr/bin/mysqladmin -u root -p shutdown

or *service mysqld stop*

If command fails,

On AIX, check the */usr/local/mysql/data/mn20.err* file.

On Linux, check */var/log/mysqld.log*.

8. Setup MySQL to automatically start after a reboot of the Management Node.

On AIX, add the following line before the xcatd line (so that it starts before it) in the */etc/inittab*:

mysql:2:once:/usr/local/mysql/bin/mysqld_safe --user=mysql &

On Linux: *chkconfig mysqld on (on sles chkconfig mysql on)*

9. Set the MySQL root password in the MySQL database

On AIX: */usr/local/mysql/bin/mysqladmin -u root password 'new-password'*

On Linux: */usr/bin/mysqladmin -u root password 'new-password'*

10. Log into the MySQL interactive program.

On AIX: */usr/local/mysql/bin/mysql -u root -p*

On Linux: */usr/bin/mysql -u root -p*

11. **Create the xcatdb database** which will be populated with xCAT data later in this document.

mysql > CREATE DATABASE xcatdb;

12. Create the xcatadmin id and password

```
mysql > CREATE USER xcatadmin IDENTIFIED BY 'xcat201';
```

13. Create the lists of hosts that will have permission to access the database.

First add your Management Node (MN), where the database is running. A good name to use for your MN is the name in the master attribute of the site table.

Names must be resolvable hostnames or ip addresses. So in our example, if you run host mn20, make sure it returns mn20 is xxx.xx.xx.xx. If it returns a long host name sure as mn20.cluster.net is xxx.xx.xx.xx, then put both the long and short hostname in the database. We assume below the short hostname is resolved to the short hostname.

```
mysql > GRANT ALL on xcatdb.* TO xcatadmin@mn20 IDENTIFIED BY 'xcat201';
```

Next add all other nodes that need access to the database. Service Nodes are required for xCAT hierarchical support. Compute nodes may also need access depending on the application running.

```
mysql > GRANT ALL on xcatdb.* TO xcatadmin@<servicenode(s)> IDENTIFIED BY 'xcat201';
```

Note: You want to do a GRANT ALL to every ipaddress or nodename that will need to access the database. You can use wildcards as follows:

```
mysql > GRANT ALL on xcatdb.* TO xcatadmin@'%cluster.net' IDENTIFIED BY 'xcat201';
mysql > GRANT ALL on xcatdb.* TO xcatadmin@'8.113.33%' IDENTIFIED BY 'xcat201';
```

14. Verify the user table was populated.

```
mysql > SELECT host, user FROM mysql.user;
+-----+-----+
| host                | user          |
+-----+-----+
| %                   | xcatadmin    |
| 127.0.0.1           | root         |
| %cluster.net        | xcatadmin    |
| 8.113.33%           | xcatadmin    |
| mn20                | xcatadmin    |
| localhost           |              |
| localhost           | root         |
+-----+-----+
```

15. Check system variables.

```
mysql > SHOW VARIABLES;
```

16. Check the defined databases.

```
mysql > SHOW DATABASES;
```

```
+-----+  
| Database          |  
+-----+  
| information_schema |  
| mysql             |  
| test              |  
| xcatdb            |  
+-----+
```

17. The following shows you how to view the tables. At this point no tables have been defined in the xcatdb yet. Run again after the database is populated.

```
mysql > use xcatdb;  
mysql > SHOW TABLES;  
mysql > DESCRIBE <table name>;
```

18. Exit out of MySQL.

```
mysql > quit;
```

1.4 Migrate xCAT data to MySQL

If you are using the mysqlsetup script from xCAT2.3.1 or later, this section will automatically be done for you. See man mysqlsetup.

1. You must backup your xCAT data before populating the xcatdb database. There are required default entries that were created in the SQLite database when the xCAT RPMs were installed on the Management Node, and they must be migrated to the new MySQL database.

```
mkdir -p ~/xcat-dbback  
dumpxCATdb -p ~/xcat-dbback
```

Note: if you get an error, like “Connection failure: IO::Socket::SSL: connect: Connection refused at....,” make sure your xcatd daemon is running.

2. **Creating the /etc/xcat/cfgloc file** tells xcat what database to use. If the file does not exist, it uses by default SQLite, which is setup during the xCAT install by default. The information you put in the files, corresponds to the

information you setup when you configured the database. Create a file called */etc/xcat/cfgloc* and populate it with the following line:

```
mysql:dbname=xcatdb;host=mn20|xcatadmin|xcat201
```

The *dbname* is the *xcatdb* you previously created. The *host* must match what is in *site.master* for the Management Node which you entered as a resolvable hostname that could access the database with the “Grant ALL” command. The *xcatadmin* and password must match what was setup when you setup your *xcatadmin* and password when you created the *xcatadmin* and password.

Finally change permissions on the file, so only root can read, to protect the password.

```
chmod 0600 /etc/xcat/cfgloc
```

3. Restore your database to MySQL. Use bypass mode to run the command without *xcatd*. This is needed to make sure the data is restored to the MySQL database and not to SQLite.

```
XCATBYPASS=1 restorexCATdb -p ~/xcat-dbback
```

Note: if you get errors, check to make sure you updated the my.cnf files as indicated above so that it runs in ANSI-QUOTES mode. If you still have errors, you can go back to using SQLite, by moving /etc/xcat/cfgloc to /etc/xcat/cfgloc.save and restarting xcatd.

4. Start the *xcatd* daemon using the MySQL database.

On AIX: *xcatstart* (*xCAT2.4 restartxcatd*)

On Linux: *service xcatd restart*

5. Test the database

```
tabdump site
```

1.5 Add ODBC support

Note: You only need to follow the steps in this section on adding ODBC support, if you plan to develop C, C++ database applications on the database or run such applications (like LoadLeveler). Otherwise skip to the next section.

As of xCAT 2.3.1 or later, you can use `mysqlsetup` command in xCAT to perform the operations in steps 2 and 3 below. See manpage for `mysqlsetup`.

1. Install ODBC package and MySQL connector.

On AIX: You need the *unixODBC* and *mysql-connector-odbc* packages included in the *dep-aix-xxxx.tar.gz* file. The *mysql-connector-odbc* package is included in *xcat-mysql-xxxx.tar.gz*. Both *.gz* files were downloaded when was set up the xCAT Management Node. To install the packages use the following commands:

```
rpm -i unixODBC-*  
rpm -i mysql-connector-odbc-*
```

On Linux: These packages come as part of the OS. Please make sure the following packages are installed on your management node. Please refer to chapter 2.2.18 of the [xCAT Top Doc](#) for how to get packages from the distro.

For RedHat and Fedora:

```
rpm -i unixODBC-*  
rpm -i mysql-connector-odbc-*
```

For SLES:

```
rpm -i unixODBC-*  
rpm -i mysql-client-*  
rpm -i libmysqlclient*  
rpm -i MyODBC-unixODBC-* (Please note that MyODBC-unixODBC rpm  
can be found in SDK CD 1 for SLES 11)
```

2. To configure ODBC you need to make changes to the *odbcinst.ini* and *odbc.ini* files so that ODBC works with the xCAT database. First update the *odbcinst.ini* file with the correct *libmyodbc#.so* name.

On AIX, RH and Fedora:

First find the ODBC driver. For AIX, RH and Fedora:

```
rpm -ql mysql-connector-odbc
```

```
vi /etc/odbcinst.ini
[MySQL]
Description = ODBC for MySQL
Driver      = /usr/lib/libmyodbc3.so
```

On SLES:

First find the ODBC driver. For SLES:

```
rpm -ql MyODBC-unixODBC
```

```
vi /etc/unixODBC/odbcinst.ini
[MySQL]
Description = ODBC for MySQL
Driver      = /usr/lib64/libmyodbc3.so
```

Then update the *obdc.ini* files with the DSN information for ODBC. Use **SERVER**, and **DATABASE** name as defined in the */etc/xcat/cfgloc* file. **DRIVER** and **PORT** are fixed.

On AIX, RH and Fedora:

```
vi /etc/odbc.ini
[xCATDB]
Driver    = MySQL
SERVER    = mn20
PORT      = 3306
DATABASE  = xcatdb
```

On SLES:

```
vi /etc/unixODBC/odbc.ini
[xCATDB]
Driver    = MySQL
SERVER    = mn20
PORT      = 3306
DATABASE  = xcatdb
```

3. Put the *xcatadmin* id and password for *xcatdb* database on the root's home directory so that user root will not have to specify them when accessing the database through ODBC. The **SERVER**, **DATABASE**, **USER** and **PASSWORD** must match was put in the */etc/xcat/cfgloc* file.

```
vi ~/.odbc.ini

[xCATDB]
SERVER = mn20
DATABASE = xcatdb
USER    = xcatadmin
PASSWORD = xcat201
```

Update the permissions for root only read of the file.

```
chmod 0600 .odbc.ini
```

4. Configure the Service Node. **Skip this step if there are no service nodes in the cluster.** If there are service nodes in the cluster you need to install unixODBC and MySQL connector on them and modify the ODBC configuration files just as we did in step 1 and 2. xCAT has utilities to install additional software on the nodes. To install ODBC and MySQL on to the service nodes, refer to the following documents for details:

AIX: [xCAT2onAIXUpdates.pdf](#)

Linux: [xCAT2-updatenode.pdf](#)

Then sync the *odbcinst.ini* and *odbc.ini* files to the service nodes. The service is the node group name for all the service nodes.

On AIX:

```
xhcp service -v /etc/odbcinst.ini /etc/odbcinst.ini
xhcp service -v /etc/odbc.ini /etc/odbc.ini
xhcp service -v /.odbc.ini /.odbc.ini
```

On RH and Fedora:

```
xhcp service -v /etc/odbcinst.ini /etc/odbcinst.ini
xhcp service -v /etc/odbc.ini /etc/odbc.ini
xhcp service -v /root/.odbc.ini /root/.odbc.ini
```

On SLES

```
xhcp service -v /etc/unixODBC/odbcinst.ini /etc/unixODBC/odbcinst.ini
xhcp service -v /etc/unixODBC/odbc.ini /etc/unixODBC/odbc.ini
xhcp service -v /root/.odbc.ini /root/.odbc.ini
```

5. Test the ODBC connection.

On AIX, as root: */usr/local/bin/isql -v xcatdb*

or as non-root user: */usr/local/bin/isql -v xcatdb xcatadmin xcat201*

On Linux, as root: */usr/bin/isql -v xcatdb*

or as non-root user: */usr/bin/isql -v xcatdb xcatadmin xcat201*

```
/usr/local/bin/isql -v xcatdb
```

```
+-----+
| Connected! |
```

```

|
| sql-statement
| help[tablename]
| quit
+-----+

```

SQL> SHOW TABLES;

```

+-----+
| Tables_in_xcatdb |
+-----+
| .
| networks
| nodelist
| site
| .
| .
+-----+

```

SQL> DESCRIBE site;

```

+-----+-----+-----+-----+-----+-----+
| Field      | Type          | Null  | Key  | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| key        | varchar(128) | NO    | PRI  |          |       |
| value      | text          | YES   |      | NULL    |       |
| comments   | text          | YES   |      | NULL    |       |
| disable    | text          | YES   |      | NULL    |       |
+-----+-----+-----+-----+-----+-----+

```

SQL > quit;

1.6 Removing MySQL xcatd database

1. To remove the database, first do a backup.

```

mkdir -p ~/xcat-dbback
dumpxCATdb -p ~/xcat-dbback

```

2. stop the xcatd daemon
3. Now remove the database.

```

On AIX: /usr/local/mysql/bin/mysql -u root -p
On Linux: /usr/bin/mysql -u root -p

```

4. Remove the database.

```
mysql> drop database xcatdb;
```

5. remove /etc/xcat/cfgloc file (points xCAT to MySQL)

6. Install the MySQL database into SQLite

```
XCATBYPASS=1 restorexCATdb -p ~/xcat-dbback
```

7. start xcatd

8. If you wish to remove all MySQL use the rpm -e command at this point.

1.7 Migrate to new level MySQL

When migrating to a new xCAT level of MySQL go through the entire setup again. This is best to stay on your current level, even though a new one has been made available. In the future, we will be changing the install of MySQL to be more automated so this will not be the case. To summarize do the following:

1. Backup your database. Refer to section *1.3 Migrate xCAT data to MySQL*.
2. Stop the MySQL. Refer to section *1.2 Configure MySQL*.
3. Stop xcatd daemon.

On AIX: *xcatstop (xCAT2.4 stopsrc -s xcatd)*

On Linux: *service xcatd stop*

4. Unlink the previous version of MySQL.

```
cd /usr/local
```

```
rm mysql
```

5. Download the latest MySQL as indicated section *1.1 Install MySQL*.
6. Follow the entire install process outlined in sections *1.1 Install MySQL* and *1.2 Configure MySQL*. You do not need to create the mysql id or group on AIX, since they already exist. You will need to create the /etc/my.cnf file.
7. Restore your database and start xcatd as you did in section *1.3 Migrate xCAT data to MySQL*.
8. You are now running on the new database level. You can remove the old level by going to /usr/local and removing the *mysql-5.0.67-aix5.3-powerpc-64bit* directory. It takes up a lot of space under /usr/local. Be sure your new level is

running and your database is restored. The old directory does contain the mysql data.

1.8 Diagnostics

- During restore to the MySQL database, if you see the following error message on the creation of tables:
#1071 - Specified key was too long; max key length is 1000 bytes

Check the Default char set of xcatdb database and change to Latin1, if needed:

- Log into the MySQL interactive program
- `mysql>use xcatdb;`
- `mysql>SHOW CREATE DATABASE xcatdb;`
if the default character set is not Latin1, then
- `mysql>ALTER DATABASE xcatdb DEFAULT CHARACTER SET latin1;`
- `mysql> quit`
- Restore you xcatdb again, or at least the tables that got errors.

1.9 References

- <http://www.pantz.org/software/mysql/mysqlcommands.html>
- <http://dev.mysql.com/doc/refman/5.0/en/tutorial.html>