




## Chapter 27.

# SECTION VIII: MySQL DATA REPLICATION

## ONE-WAY Data Replication In MySQL On Windows

### Getting Started

To begin the process of data replication three nodes with following specifications are used.

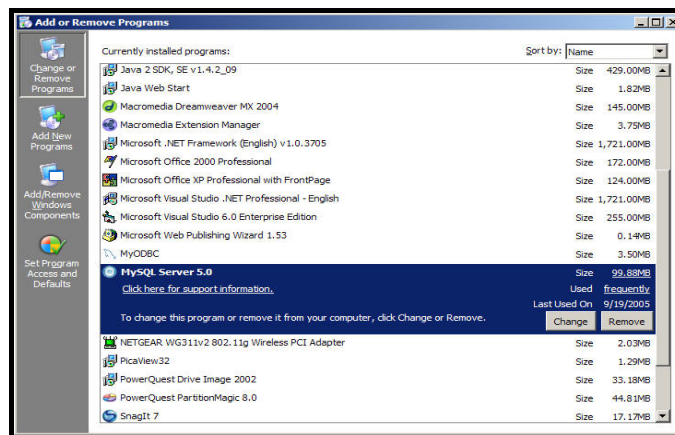
					
<b>Node Type</b>	MASTER	<b>Node Type</b>	SLAVE 1	<b>Node Type</b>	SLAVE 2
<b>Host Name</b>	Jasmine	<b>Host Name</b>	Lily	<b>Host Name</b>	Daffodil
<b>IP</b>	192.168.0.4	<b>IP</b>	192.168.0.101	<b>IP</b>	192.168.0.5
<b>Processor</b>	Athlon 1.7	<b>Processor</b>	Athlon 2.4	<b>Processor</b>	Athlon 2.6
<b>RAM</b>	512	<b>RAM</b>	512	<b>RAM</b>	512
<b>O/S</b>	Windows XP	<b>O/S</b>	Windows XP	<b>O/S</b>	Windows XP
<b>MySQL Version</b>	5.0.9	<b>MySQL Version</b>	5.0.9	<b>MySQL Version</b>	5.0.9

**Table 27.1**

All the three nodes i.e. MASTER (Jasmine), SLAVE 1 (Lily) and SLAVE 2 (Daffodil) have MySQL 5.x installed. Refer to table 27.1.

### Steps To Setup ONE-WAY Replication

Ensure all the nodes have **same** (or compatible) versions of MySQL installed. This can be verified using Add/Remove Programs under Control Panel:



**Diagram 27.1:** Verifying MySQL 5.0.x installation on Windows

Verify this on all the nodes (in this material Jasmine, Lily and Daffodil were verified). As far as replication is concerned, any MySQL 4.1.x version and any 4.0.x version are identical, because they all use the same binary log format. Thus, any servers from these versions are compatible and replication between them should work seamlessly. The exception to this compatibility is that versions from MySQL 4.0.0 to 4.0.2 were very early development versions **that should not be used anymore**.

The original binary log format was developed in MySQL 3.23. It changed in MySQL 4.0 and again in MySQL 5.0. This has consequences when servers are upgraded in a replication setup.

		MASTER	MASTER	MASTER
		3.23.33 and up	4.0.3 and up or any 4.1.x	5.x
SLAVE	3.23.33 and up	✓	✗	✗
SLAVE	4.0.3 and up	✓	✓	✗
SLAVE	5.x	✓	✓	✓

**Table 27.2**

Table 27.2 indicates the MASTER/SLAVE replication compatibility between different versions of MySQL.

Even though compatible versions can be used across nodes using the same and the latest version on all the nodes is recommended.

**1. On The MASTER Node (Jasmine):**

Stop the MySQL service on the MASTER node if already running as:

```
<System Prompt> net stop mysql
```

Ensure the following lines exist in the my.ini under **C:\Program Files\MySQL\MySQL Server 5.0** on the MASTER node i.e. Jasmine.

The file **my.ini** can be opened using **notepad** as follows: (Refer to diagram 27.2)

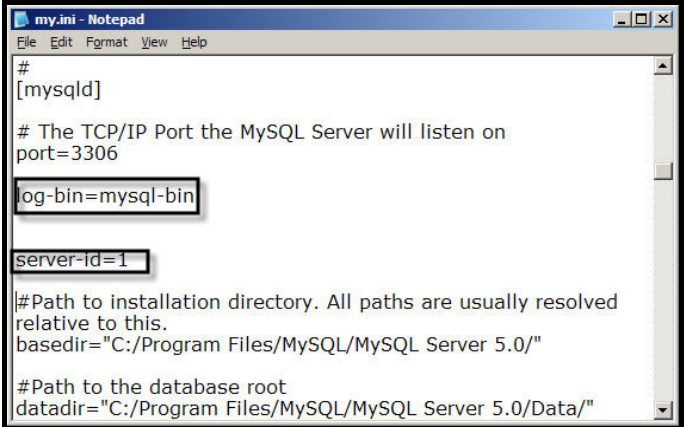
```
<System Prompt> notepad C:\Program Files\MySQL\MySQL Server 5.0\my.ini
```



**Diagram 27.2:** Opening my.ini file in notepad

The file **my.ini** on the MASTER node (Jasmine) should contain the following two lines under the **mysqld** section: (Refer to diagram 27.3)

```
[mysqld]  
log-bin=mysq-bin  
server-id=1
```



**Diagram 27.3:** Contents of the file my.ini on the **MASTER** node (Jasmine)

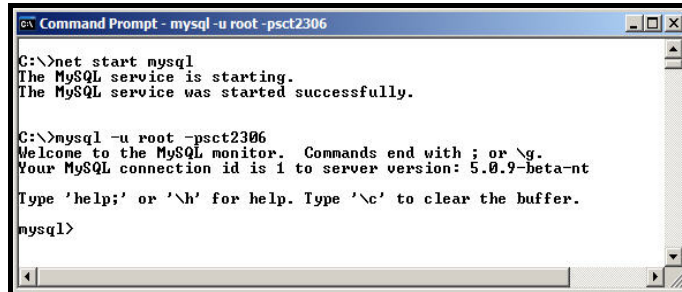
Under the **[mysqld]** section of the my.ini file on the MASTER node (Jasmine):

- ❑ The **log-bin** option enables the binary logging support, which in turn allows actual replication
- ❑ The **server-id=master\_id** option, where **master\_id** must be a positive integer value from **1** to **2^32 - 1**. **server-id** values can be thought of as something similar to an IP address. The server-id uniquely identifies each server instance in the community of replication partners

**27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS**

If these lines are not present add them, start the server and log in as: (Refer to diagram 27.4)

```
<System Prompt> net start mysql
<System Prompt> mysql -u root -psct2306
```



```
Command Prompt - mysql -u root -psct2306
C:\>net start mysql
The MySQL service is starting.
The MySQL service was started successfully.

C:\>mysql -u root -psct2306
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 5.0.9-beta-nt
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>
```

**Diagram 27.4:** MySQL server on the **MASTER** node (Jasmine)

Keep this terminal session of MySQL on the MASTER node (Jasmine) open. This session will be used and required later for additional commands.

## 2. On The SLAVE Nodes:

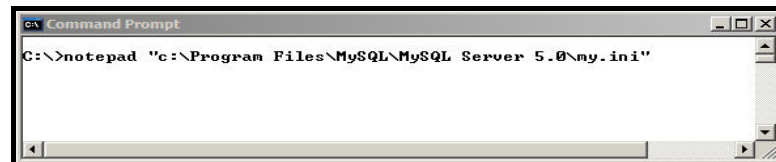
### On The SLAVE Node 1 (Lily):

Stop the mysql service on the SLAVE node 1 (Lily) if already running as:

```
<System Prompt> net stop mysql
```

The file **my.ini** can be opened using **notepad** as follows: (Refer to diagram 27.5)

```
<System Prompt> notepad C:\Program Files\MySQL\MySQL Server 5.0\my.ini
```

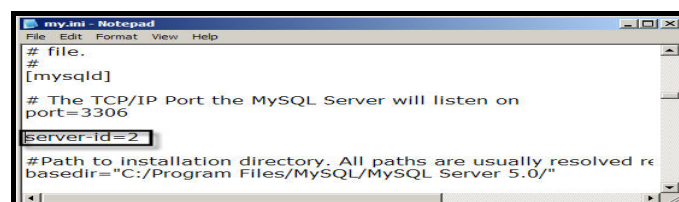


```
Command Prompt
C:\>notepad "c:\Program Files\MySQL\MySQL Server 5.0\my.ini"
```

**Diagram 27.5:** Opening the file **my.ini** on the **SLAVE** node 1 (Lily)

The file **my.ini** on the SLAVE node should contain: (Refer to diagram 27.6)

```
[mysqld]
server-id=2
```



```
my.ini - Notepad
File Edit Format View Help
# file.
#
[mysqld]
# The TCP/IP Port the MySQL Server will listen on
port=3306
server-id=2
#Path to installation directory. All paths are usually resolved re
basedir="C:/Program Files/MySQL/MySQL Server 5.0/"
```

**Diagram 27.6:** Editing the file **my.ini** on the **SLAVE** node 1 (Lily)

This is done, as there is no need to enable binary logging on the **SLAVE** node **1** as this Slave is not going to act as a Master for another computer.

**On The SLAVE Node 2 (Daffodil):**

Stop the mysql service on the SLAVE node 2 (Daffodil) if already running as:

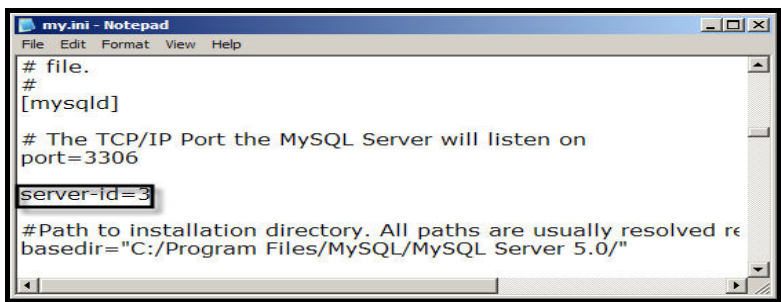
```
<System Prompt> net stop mysql
```

The file **my.ini** can be opened using **notepad** as follows: (Refer to diagram 27.5)

```
<System Prompt> notepad C:\Program Files\MySQL\MySQL Server 5.0\my.ini
```


The file my.ini on the SLAVE node should contain: (Refer to diagram 27.7)

```
[mysqld]
server-id=3
```



**Diagram 27.7:** Editing the file **my.ini** on the **SLAVE** node **2** (Daffodil)

**WARNING**

 **Do not start** the mysql service on the SLAVE nodes (i.e. SLAVE 1 and SLAVE 2) yet.

**3. On The MASTER Node (Jasmine):**

Set up a user account named **rep1** on the MASTER node (Jasmine) via which the SLAVE nodes (Lily and Daffodil) can connect.

This account must be given the **REPLICATION SLAVE** privilege. If the account is used only for replication (which is recommended), do not grant any additional privileges.

Using this account users should be able to access the MASTER node from any host in the domain using a password of choice.

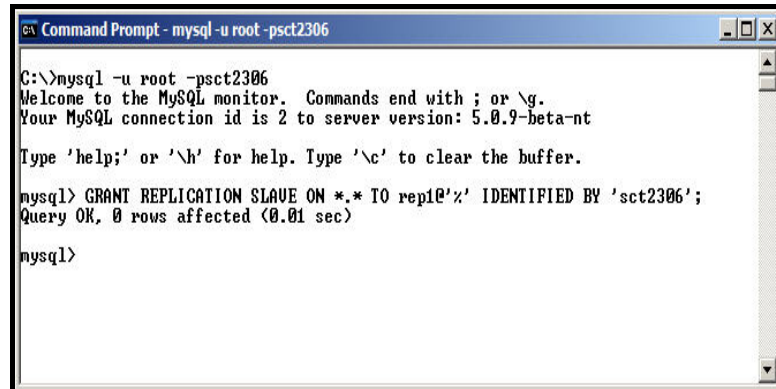
To create the account, use the following **GRANT** statement: (Refer to diagram 27.8)

```
mysql> GRANT REPLICATION SLAVE ON *.* TO rep1@'%'
IDENTIFIED BY 'sct2306';
```

**27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS**

To achieve the same in MySQL versions **older than 4.0.2**, where the REPLICATION SLAVE privilege does not exist. Grant the FILE privilege instead as:

```
mysql> GRANT FILE ON *.* TO repl1@'%' IDENTIFIED BY 'sct2306';
```



```

C:\>mysql -u root -psct2306
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 2 to server version: 5.0.9-beta-nt

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> GRANT REPLICATION SLAVE ON *.* TO repl1@'%' IDENTIFIED BY 'sct2306';
Query OK, 0 rows affected (0.01 sec)

mysql>

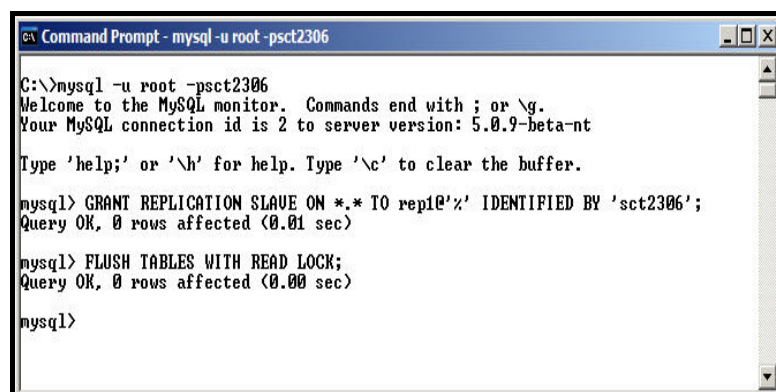
```

**Diagram 27.8:** Setting up a user account for replication on the **MASTER** node (Jasmine)

#### 4. On The MASTER Node (Jasmine):

Flush all the tables and block the write statements on the MASTER node by executing a FLUSH TABLES WITH READ LOCK statement as: (Refer to diagram 27.9)

```
mysql> FLUSH TABLES WITH READ LOCK;
```



```

C:\>mysql -u root -psct2306
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 2 to server version: 5.0.9-beta-nt

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> GRANT REPLICATION SLAVE ON *.* TO repl1@'%' IDENTIFIED BY 'sct2306';
Query OK, 0 rows affected (0.01 sec)

mysql> FLUSH TABLES WITH READ LOCK;
Query OK, 0 rows affected (0.00 sec)

mysql>

```

**Diagram 27.9:** Flushing tables and blocking writes on the **MASTER** node (Jasmine)

Leave this MySQL client terminal on the MASTER node (Jasmine) as shown in Diagram 27.9 running from which the FLUSH TABLES statement is issued. This will enable the read lock to remain effective. If the MySQL client terminal is closed, the lock will be released.

**For any further operations on the MASTER node a different terminal window can be used.**

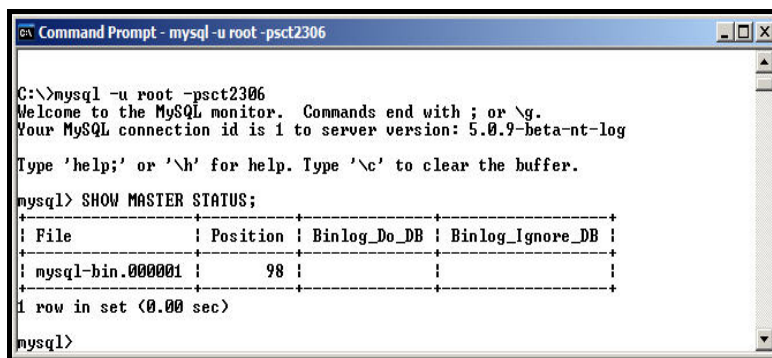
#### 5. On The MASTER Node (Jasmine):

Copy the required databases from the MASTER node (Jasmine)'s data directory i.e. **C:\Program Files\MySQL\MySQL Server 5.0\data** to the SLAVE nodes (Lily and Daffodil)'s data directory i.e. **C:\Program Files\MySQL\MySQL Server 5.0\data**. The databases are usually in form of directories. So this task can simply be done by copying the databases (directories) required from the MASTER node's data directory to the SLAVE node's data directory.

#### 6. On The MASTER Node (Jasmine):

While the read lock placed by FLUSH TABLES WITH READ LOCK is still in effect (Refer to point 4: *Leave this mysql client terminal on the MASTER node running*), read the value of the current binary log name and offset on the MASTER node as: (Refer to diagram 27.12)

```
mysql> SHOW MASTER STATUS;
```



```

C:\>mysql -u root -psct2306
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 5.0.9-beta-nt-log

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| mysql-bin.000001 |      98 |               |                   |
+-----+-----+-----+-----+
1 row in set (0.00 sec)

mysql>

```

**Diagram 27.12:** The MASTER status on the MASTER node (Jasmine)

Note these values down. These values will be required to be entered on to the SLAVE node 1 and Slave node 2 later while setting up the SLAVE nodes.

The **File** column shows the **name** of the log, while **Position** shows the **offset**. In this example, the binary log value is **mysql-bin.000001** and the offset is **98**. This represents the replication coordinates at which the **SLAVE** nodes (i.e. **Lily** and **Daffodil**) should begin processing new updates from the **MASTER** node (**Jasmine**).

#### 7. On The MASTER Node (Jasmine):

After the snapshot has been taken (Refer to point 6) and the log name and offset are recorded the write activity on the MASTER node can be re-enabled as: (Refer to diagram 27.13)

```
mysql> UNLOCK TABLES;
```

## 27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS

```

Command Prompt - mysql -u root -psct2306
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 5.0.9-beta-nt-log

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> SHOW MASTER STATUS;
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| mysql-bin.000001 |      98 |               |                   |
+-----+-----+-----+-----+
1 row in set (0.00 sec)

mysql> UNLOCK TABLES;
Query OK, 0 rows affected (0.01 sec)

mysql>

```

**Diagram 27.13:** Unlocking the tables on the MASTER node (Jasmine)

## 8. On The SLAVE Nodes:

Start the mysql service on the SLAVE nodes (Lily & Daffodil). Ensure that before the mysql service on the SLAVE nodes (Lily & Daffodil) is started the snapshot has been unpacked on both the SLAVE nodes's data directory.

Start the mysql service on the SLAVE node and log into mysql server as: (Refer to diagram 27.14)

```
<System Prompt> net start mysql
```

Execute the following statements via the mysql terminal of the SLAVE nodes, replacing the values within <> with the actual values taken from the MASTER node (i.e. Jasmine's) spec:

```

mysql> CHANGE MASTER TO
-> MASTER_HOST='<master hostname>',
-> MASTER_USER='<replication username>',
-> MASTER_PASSWORD='<replication password>',
-> MASTER_LOG_FILE='<recorded log file name>',
-> MASTER_LOG_POS='<recorded log offset>;

```

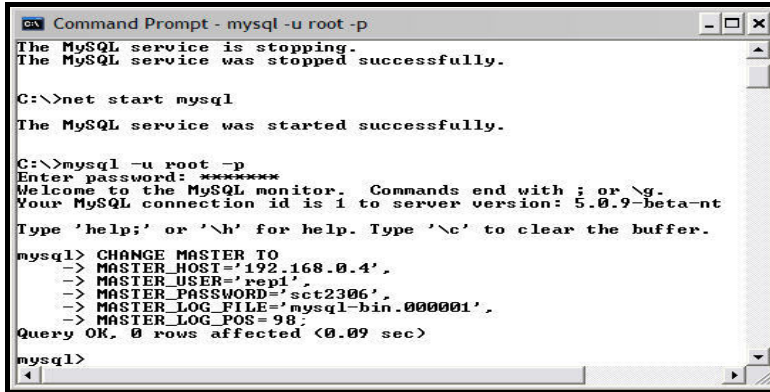
The following table shows the maximum string length for these variables:

MASTER_HOST	60
MASTER_USER	16
MASTER_PASSWORD	32
MASTER_LOG_FILE	255



The above command will look like: (Refer to diagram 27.14)

```
mysql> CHANGE MASTER TO
-> MASTER_HOST='192.168.0.4',
-> MASTER_USER='repl',
-> MASTER_PASSWORD='sct2306',
-> MASTER_LOG_FILE='mysql-bin.000001',
-> MASTER_LOG_POS=98;
```



```
Command Prompt - mysql -u root -p
The MySQL service is stopping.
The MySQL service was stopped successfully.

C:\>net start mysql
The MySQL service was started successfully.

C:\>mysql -u root -p
Enter password: *****
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 5.0.9-beta-nt
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> CHANGE MASTER TO
-> MASTER_HOST='192.168.0.4',
-> MASTER_USER='repl',
-> MASTER_PASSWORD='sct2306',
-> MASTER_LOG_FILE='mysql-bin.000001',
-> MASTER_LOG_POS=98;
Query OK, 0 rows affected (0.09 sec)

mysql>
```

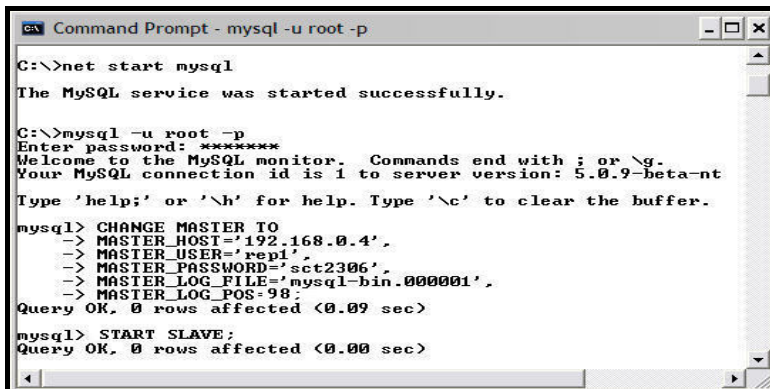
Diagram 27.14: Starting MySQL service on SLAVE nodes (Lily & Daffodil)

## 9. On The SLAVE Nodes:

Once the SLAVE nodes Lily and Daffodil are configured to point to the MASTER node Jasmine using the replication account username and password, the actual replication can be commenced, by starting the SLAVE threads on the SLAVE nodes as:

Issue the following command at the mysql prompt: (Refer to diagram 27.15)

```
mysql> START SLAVE;
```



```
Command Prompt - mysql -u root -p
C:\>net start mysql
The MySQL service was started successfully.

C:\>mysql -u root -p
Enter password: *****
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 1 to server version: 5.0.9-beta-nt
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> CHANGE MASTER TO
-> MASTER_HOST='192.168.0.4',
-> MASTER_USER='repl',
-> MASTER_PASSWORD='sct2306',
-> MASTER_LOG_FILE='mysql-bin.000001',
-> MASTER_LOG_POS=98;
Query OK, 0 rows affected (0.09 sec)

mysql> START SLAVE;
Query OK, 0 rows affected (0.00 sec)

mysql>
```

Diagram 27.15: Starting the SLAVE thread on the SLAVE nodes (Lily & Daffodil)

## 27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS

Once the slaves are replicated, two files named **master.info** and **relay-log.info** will be automatically created on the SLAVE nodes Lily and Daffodil by mysql in their data directories i.e. **C:\Program Files\MySQL\MySQL Server 5.0\data**. The SLAVE nodes use these two files to keep track of how much of the MASTER node (Jasmine)'s binary log it has processed.

These are status files and as disk files they survive a SLAVE server's shutdown. The next time the SLAVE node(s) starts up, it reads these files to determine how far it has proceeded in reading binary logs from the MASTER and in processing its own relay logs.

### **WARNING**



Do not remove or edit these files, unless one really knows what one is doing and what will be the implications. Even in that case, always prefer to use the CHANGE MASTER TO statement.

### **REMINDER**



Once a snapshot is created, the same snapshot can be used to set up other slaves by following the **SLAVE PORTION** of the procedure just described.

Another snapshot of the MASTER need not be taken. The same snapshot can be used for each additional SLAVE.

## **10. On The MASTER Node:**

Once the replication process as explained till now is followed and setup on all the three nodes i.e. 1 MASTER and 2 SLAVE nodes, the SHOW PROCESSLIST command can be used to verify and view replication status.

The SHOW PROCESSLIST statement provides information about what is happening on the MASTER and on the SLAVE(s) regarding replication.

Execute the following command on the MASTER node (Jasmine):

```
mysql> SHOW PROCESSLIST \G;
```

### **Output:**

```
***** 1. row *****
  Id: 1
  User: root
  Host: localhost:1048
  db: NULL
  Command: Query
  Time: 0
  State: NULL
  Info: show processlist
```

**Output:**

```

***** 2. row *****
  Id: 2
  User: repl
  Host: LILY:1150
  db: NULL
Command: Binlog Dump
  Time: 124
  State: Has sent all binlog to slave; waiting for binlog to be
updated
  Info: NULL
***** 3. row *****
  Id: 3
  User: repl
  Host: DAFFODIL:1060
  db: NULL
Command: Binlog Dump
  Time: 36
  State: Has sent all binlog to slave; waiting for binlog to be
updated
  Info: NULL
3 rows in set (0.02 sec)

```

Here, thread 1 is the SQL thread that is processing the updates stored in the relay logs, thread 2 and 3 are replication threads for connected SLAVE nodes. The information indicates that all outstanding updates have been sent to the SLAVE nodes and that the MASTER is waiting for more updates to occur.

**11. On The SLAVE Nodes:****On The SLAVE Node 1 (Lily):**

Execute the following command on the SLAVE node 1 (Lily):

```
mysql> SHOW PROCESSLIST \G;
```

**Output:**

```

***** 1. row *****
  Id: 1
  User: root
  Host: localhost:1142
  db: NULL
Command: Query
  Time: 0
  State: NULL
  Info: SHOW PROCESSLIST

```

**27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS**

**Output:**

```

***** 2. row *****
  Id: 2
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 209
  State: Waiting for master to send event
  Info: NULL
***** 3. row *****
  Id: 3
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 194
  State: Has read all relay log; waiting for the slave I/O
thread to update it
  Info: NULL
3 rows in set (0.00 sec)

```

**On The SLAVE Node 2 (Daffodil):**

Execute the following command on the SLAVE node 2 (Daffodil):

```
mysql> SHOW PROCESSLIST \G;
```

**Output:**

```

***** 1. row *****
  Id: 1
  User: root
  Host: localhost:1055
  db: NULL
Command: Query
  Time: 0
  State: NULL
  Info: Show Processlist
***** 2. row *****
  Id: 2
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 118
  State: Waiting for master to send event

```

**Output:** (Continued)

```

Info: NULL
***** 3. row *****
  Id: 3
  User: system user
  Host:
  db: NULL
Command: Connect
  Time: 103
  State: Has read all relay log; waiting for the slave I/O
thread to update it
  Info: NULL
3 rows in set (0.03 sec)
    
```

This information indicates that thread 1 is the SQL thread that is processing the updates stored in the relay logs, thread 2 is the I/O thread that is communicating with the MASTER server, and thread 3 is the SQL thread that is processing the updates stored in the relay logs. Currently, both threads are idle, waiting for further updates.

**HINT**



Note that the value in the **Time** column tells how late the SLAVE is compared to the MASTER.

**Testing The Replication Setup**

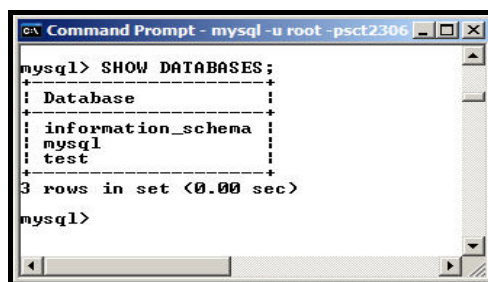
To test whether the data from the MASTER node (Jasmine) is actually replicated on to the SLAVE node(s) (Lily and Daffodil) follow the steps:

**1. On The MASTER Node (Jasmine):**

Issue the following command: (Refer to diagram 27.16)

```
mysql> SHOW DATABASES;
```

This will list the databases available as of now on the MASTER node (Jasmine).



**Diagram 27.16:** Databases on the MASTER node (Jasmine)

**27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS**

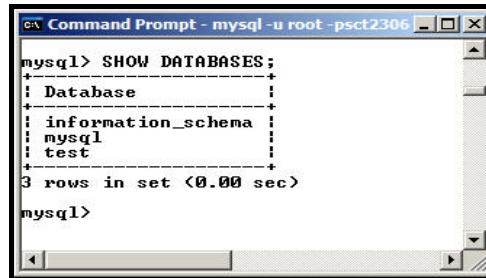
## 2. On The SLAVE Nodes:

### On The SLAVE Node 1 (Lily):

Issue the following commands: (Refer to diagram 27.17)

```
mysql> SHOW DATABASES;
```

This will list the databases available as of now on the SLAVE node 1 (Lily).



```
mysql> SHOW DATABASES;
+-----+
| Database |
+-----+
| information_schema |
| mysql        |
| test        |
+-----+
3 rows in set <0.00 sec>

mysql>
```

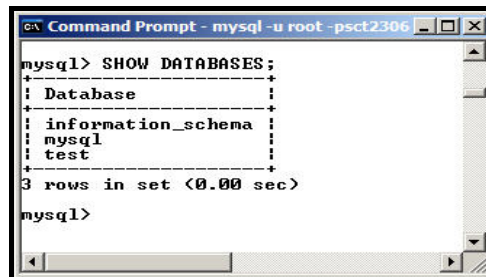
Diagram 27.17: Databases on the SLAVE node 1 (Lily)

### On The SLAVE Node 2 (Daffodil):

Issue the following commands: (Refer to diagram 27.18)

```
mysql> SHOW DATABASES;
```

This will list the databases available as of now on the SLAVE node 2 (Daffodil).



```
mysql> SHOW DATABASES;
+-----+
| Database |
+-----+
| information_schema |
| mysql        |
| test        |
+-----+
3 rows in set <0.00 sec>

mysql>
```

Diagram 27.18: Databases on the SLAVE node 2 (Daffodil)

## 3. On The MASTER Node (Jasmine):

Issue the following command: (Refer to diagram 27.19)

```
mysql> CREATE DATABASE sct;
```

This will create a database named **sct** on the MASTER node (Jasmine).

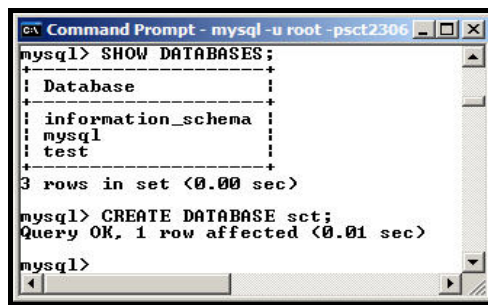


Diagram 27.19: Create Database on the MASTER node (Jasmine)

**4. On The SLAVE Nodes:**

**On The SLAVE Node 1 (Lily):**

Issue the following commands: (Refer to diagram 27.20)

```
mysql> SHOW DATABASES;
```

This will list the databases available on the SLAVE node 1 (Lily) along with the new database named **sct** created via the MASTER node (Jasmine).




Diagram 27.20: Databases on the SLAVE node 1 (Lily)

This shows that the data replication has taken place between the MASTER (Jasmine) and the SLAVE node 1 (Lily).

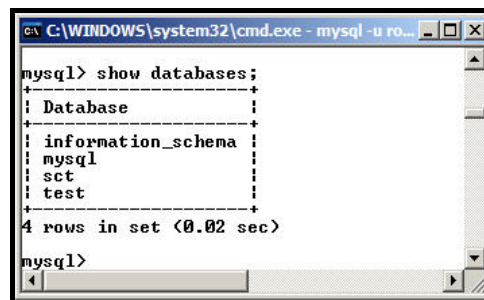
**On The SLAVE Node 2 (Daffodil):**

Issue the following commands: (Refer to diagram 27.21)

```
mysql> SHOW DATABASES;
```

This will list the databases available on the SLAVE node 2 (Daffodil) along with the new database named **sct** created via the MASTER node (Jasmine).

**27. ONE-WAY DATA REPLICATION IN MySQL ON WINDOWS**



```

C:\WINDOWS\system32\cmd.exe - mysql -u ro...
mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| mysql          |
| sct            |
| test           |
+-----+
4 rows in set (0.02 sec)

mysql>

```

**Diagram 27.21:** Databases on the SLAVE node 2 (Daffodil)

This shows that the ONE-WAY data replication has taken place between the MASTER (Jasmine) and the SLAVE node 2 (Daffodil).

### **REMINDER**



MySQL supports only one-way replication. MySQL replication currently does not support any locking protocol between MASTER and SLAVE to guarantee the atomicity of a distributed (CROSS-SERVER) update.

It is not advisable to co-chain TWO servers in a TWO-WAY replication relationship.

However, if desired a two-way replication mechanism can be setup and the load can be distributed programmatically by reversing the steps followed for implementing one-way replication.

The procedure of implementing a two-way replication mechanism on Windows is available as an appendix on the Book's accompanying CDROM.

Use this technique with CAUTION.