

# Re: MySQL Clustering

---

*Source:* <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/isp/2006-03/msg00032.html>

---

- *From:* "L. Jason Godsey" <lannygodsey@xxxxxxxxxx>
  - *Date:* Wed, 22 Mar 2006 04:35:41 -0800 (PST)
- 

The following solves replication breaking over the same auto int race condition. This can happen because of fast inserts from multiple hosts connected to multiple back end servers (8 in my case), or because of disconnected operation.

The way I solved this in the past was using composite keys. Instead of just the auto integer field, I also used a server\_id field.

Prior to inserting anything into this database, issue the query  
set @server\_id=@@server\_id;, reuse the db connection for the insert.

Sometimes you can use something like:

```
$query = "set @server_id=@@server_id; insert into.....";
```

Don't skip the seemingly redundant, @@server\_id=@@ bit and go with insert into VALUES (@@server\_id)... then each server doing the replication will use it's own server\_id which isn't what you want.

```
create table users (  
id int auto_increment,  
server_id int,  
login varchar(32),  
pass varchar(32),  
primary key (id,server_id)  
);  
create unique index _i_users_login on users (login);
```

```
mysql> set @server_id=@@server_id;  
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> insert into users (server_id, login, pass) values (@server_id,  
'lanny', 'examples');  
Query OK, 1 row affected (0.00 sec)
```

```
mysql> select * from users;  
+-----+-----+-----+-----+  
| id | server_id | login | pass |
```

## Re: MySQL Clustering

```
+-----+-----+-----+-----+
| 1 | 2 | lanny | examples |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

```
mysql> set @server_id=32; insert into users (server_id, login, pass)
values (@server_id, 'lanny3', 'examples');
Query OK, 0 rows affected (0.00 sec)
```

Query OK, 1 row affected (0.00 sec)

```
mysql> select * from users;
+-----+-----+-----+-----+
| id | server_id | login | pass |
+-----+-----+-----+-----+
| 1 | 2 | lanny | examples |
| 2 | 2 | lanny2 | examples |
| 3 | 32 | lanny3 | examples |
+-----+-----+-----+-----+
```

--- Cody Baker <cody@xxxxxxxxxxxxxx> wrote:

If your database is simply read only then the load balanced situation should work fine. If your database is read/write then your load balancer could cause you problems under certain circumstances. If your updates are not time critical then it should be fine. If, however, your databases are used for a time critical updates (ex. session data for a web page) then it's important that your users always hit the same database server because the replication can become delayed at times. If the user has a random chance of hitting any of your database servers then for example your user may add something to their cart, and refresh the page connecting to another server, and find the item missing from their cart. Other no-nos for replicated databases especially on mysql < 5, are auto-incrementing fields in tables. 5.x has a solution for this (

<http://dev.mysql.com/doc/refman/5.0/en/replication-auto-increment.html>

).

Re: MySQL Clustering

Thank You,

Cody Baker  
cody@xxxxxxxxxxxxxx  
<http://www.wilshire.net>  
Jon Simola wrote:

On 3/15/06, James Ryan <james@xxxxxxxxxxxxxxxxxxxxxxxx> wrote:

Not sure if this is the right list for this (I apologize if its  
not),

but has anybody ran a MySQL 2+ node cluster under  
FreeBSD 5.x

behind a

load balancer; and if so, could you offer any tips or  
warnings?

If you're talking about the actual MySQL clustering server setup,  
I've

never had a chance to try it as our dataset is too large.

I've run a pair of MySQL servers in a round-robin master setup (A  
slaves from B, B slaves from A) and that worked rather well,  
replication was impressively quick.

--  
Jon Simola  
Systems Administrator  
ABC Communications

---

freebsd-isp@xxxxxxxxxxxxx mailing list  
<http://lists.freebsd.org/mailman/listinfo/freebsd-isp>  
To unsubscribe, send any mail to

"freebsd-isp-unsubscribe@xxxxxxxxxxxxx"

---

freebsd-isp@xxxxxxxxxxxxx mailing list  
<http://lists.freebsd.org/mailman/listinfo/freebsd-isp>

Re: MySQL Clustering

To unsubscribe, send any mail to  
"freebsd-isp-unsubscribe@xxxxxxxxxxx"

---

freebsd-isp@xxxxxxxxxxx mailing list  
<http://lists.freebsd.org/mailman/listinfo/freebsd-isp>  
To unsubscribe, send any mail to "freebsd-isp-unsubscribe@xxxxxxxxxxx"