

# Salem Radio Labs

## APPLICATION NOTE 003

### Configuring Rivendell Hot Standby Hosts on SuSE

#### SCOPE

This application note details the procedures for configuring one or more host systems to act as a 'hot standby' Rivendell server. Such a 'hot standby' system can be used to quickly restore a Rivendell network in the event of failure in the primary server.

#### PREREQUISITES

This procedure assumes that Rivendell is already properly installed on both the master and any standby host(s). This includes proper mySQL installation and configuration. For further information on basic Rivendell installation, see SRL Application Note #1.

Ideally, the versions of mySQL being employed on the master and any standby system(s) should be identical, as data replication between different versions is not supported for all combinations. For more information regarding cross-version compatibility, consult the mySQL documentation.

#### OVERVIEW

It is possible to configure Rivendell so as to have two alternative configurations, called 'Master' and 'Standby'. Each configuration can use a different mySQL server and audio storage location (either local or connected via the network), thus allowing a backup configuration that can be quickly activated in the event of a problem with the Master setup. It is also possible to configure a standby system so that it is automatically kept up to date with any changes (configuration changes, new recordings, etc) made to the Master system. These capabilities are collectively referred to a 'hot standby' capability.

Configuring hot standby for one or more Rivendell hosts consists of three basic steps:

- 1) Enabling database replication
- 2) Enabling audio replication
- 3) Setting up the master and standby configurations.

#### Package Requirements

In addition to the standard Rivendell dependencies, both the Master Server and any standby hosts will need the 'rsync' package installed. The package is included in SuSE, or can be downloaded from:

<http://rsync.samba.org/>

#### CONFIGURING DATABASE REPLICATION

##### Preparing the Master Server

On the master server, copy the 'my.cnf-master' file (from the Rivendell distribution) to '/etc/my.cnf'. This file can be found in the 'conf/' directory of the Rivendell sources, or in the package documentation directory ('/usr/share/doc/packages/rivendell') if an RPM distribution is being used.

Next, create a mySQL user account called 'repl' with rights to connect from each hot standby host. This account should have RELOAD, REPLICATION SLAVE, REPLICATION CLIENT and SUPER rights enabled, and no others.

Finally, shutdown and then restart the mySQL server by doing (as root):

```
/etc/init.d/mysql stop
/etc/init.d/mysql start
```

### Generating The Initial Data Snapshot

Before starting replication, it is necessary to generate an initial “snapshot” of data on the master server for importation into the slave. (IMPORTANT NOTE: this procedure will lock the Rivendell database while it is being performed, effectively taking it off-line. Plan accordingly!) To generate the snapshot, proceed as follows:

On the master mySQL server , use the mySQL monitor program to log into mySQL with an administrator account, using a command such as:

```
mysql -u root -p
```

Once logged in, enter the following command:

```
FLUSH TABLES WITH READ LOCK;
```

Next, open up a root shell and make a copy of all the files in the '/var/lib/mysql' directory, by doing:

```
cd /var/lib
tar -zcf snapshot.tar.gz mysql
```

You should now have a file called 'snapshot.tar.gz', containing a snapshot of the server data.

Next, enter the following command in the mySQL monitor:

```
SHOW MASTER STATUS;
```

You should see output similar to the following:

```
+-----+-----+-----+-----+
| File           | Position | Binlog_Do_DB | Binlog_Ignore_DB |
+-----+-----+-----+-----+
| rdserv-bin.000001 | 269414518 |               |                   |
+-----+-----+-----+-----+
```

Note down the values of the 'File' and 'Position' fields. You will need these when starting replication on the slave system(s).

Finally, enter the following in the mySQL monitor:

```
UNLOCK TABLES;
```

Enter 'quit' to exit the monitor, and you're done creating the snapshot.

### Preparing the Standby Host

(IMPORTANT NOTE: this procedure will completely **destroy** any previously existing mySQL data on the slave system!). On the standby host, copy the 'my.cnf-standby' file (from the Rivendell distribution) to '/etc/my.cnf'. This file can be found in the 'conf/' directory of the Rivendell sources, or in the package documentation directory ('/usr/share/doc/packages/rivendell') if an RPM distribution is being used. Open the '/etc/my.cnf' in a text editor and set the following values appropriately for your site:

```
server-id = <id-num>
```

where <id-num> is a site unique number that identifies this MySQL server. By convention, the master Rivendell server gets a server-id of '1', and standby hosts are sequentially numbered, starting with '100'.

```
master-host = <hostname>
```

where <hostname> is the hostname of the master Rivendell server.

```
master-user = <username>
```

where <username> is the name of the account on the master server to use for replication. Normally, this is the `repl` account.

```
master-password = <password>
```

where <password> is the password for accessing the replication account on the master server.

Next, populate the snapshot data from the master system. Copy the snapshot file you created on the master to the standby system, then open a root shell on the standby system and do the following:

```
cd /var/lib  
tar -zxvf /path/to/snapshot/file/snapshot.tar.gz
```

Finally, restart the MySQL server by doing:

```
/etc/init.d/mysql start
```

### Starting Replication

To actually start replication, log into a MySQL account with administrative rights on the standby host (using the `mysql -u root -p` command) and issue the following commands:

```
CHANGE MASTER TO  
MASTER_HOST='<hostname>',  
MASTER_USER='repl',  
MASTER_PASSWORD='<password>',  
MASTER_LOG_FILE='<filename>',  
MASTER_LOG_POS=<position>;
```

where <hostname> and <password> are the hostname and MySQL password of the master system and <filename> and <position> are the values of the `File` and `Position` fields reported by the `SHOW_MASTER_STATUS` command on the master.

Next, enter the following command:

```
START SLAVE;
```

Each of these commands should return 'Query OK'. If errors occur, consult the MySQL log (at `/var/lib/mysql/mysql.log`) for clues as to what could be going wrong. Replication should now be operational.

## CONFIGURING AUDIO REPLICATION

### Preparing the Master Server

Configure an `rsync` module for the Rivendell audio archive by inserting the following into the `/etc/rsyncd.conf` file:

```
[rivendell]
  path = /var/snd
  comment = Rivendell Audio Archive
  read only = true
```

In addition, make sure that the `hosts allow=` directive in `/etc/rsyncd.conf` is set so as to allow the standby host(s) to connect. See the `rsyncd.conf(5)` man page for more details.

Next, open `Yast2`, click on `System` and then click on `Runlevel Editor`. Select the `rsyncd` service and the click `Enable`, then `Finish`.

Next, create the replication check file by doing the following as the Rivendell user:

```
touch /var/snd/repl.chk
```

The audio replication code will check to see that this file exists each time before replicating audio. This is to prevent catastrophes with mirror deletes in the event of a failure of the storage device containing the primary archive.

### Preparing the Standby Host

Create a `/var/snd.standby` directory with enough free space to be able to contain all of the audio data from the master server, and make sure that it is owned as well as readable and writable by the Rivendell user and group (normally, user `'rd'` and group `'users'`).

Run the initial audio synchronization by entering the following command as the Rivendell user:

```
rsync <hostname>::rivendell/*.wav /var/snd.standby
```

where `<hostname>` is the name of the master server. Since this is the initial synchronization, this command may take a significant amount of time to complete, as the entire audio archive must be transferred.

### Enabling Audio Replication

Add the following cron job for the Rivendell user (using the `crontab -e` command):

```
0 * * * * /usr/local/bin/rd_audio_sync <hostname>
```

where `<hostname>` is the name of the master server. This will cause the archive to be synchronized once per day. Some sites may want synchronization to occur more frequently. This can be accomplished by adjusting the parameters of the cron job – see the `crontab(5)` man page for details.

## SETTING UP CONFIGURATIONS

### Stop Daemons

Before proceeding with the rest of this procedure, stop the Rivendell daemons on the standby host by doing (as root):

```
/etc/init.d/rivendell stop
```

## Database Configurations

Prepare two 'rd.conf' files, one each for your 'Master' and 'Standby' setup. The master file should be saved as '/etc/rc.conf.master', while the standby file should be saved as '/etc/rd.conf.standby'. Be sure to back up anything in the original '/etc/rd.conf' file, as it's contents will be overwritten.

## Audio Directories

If you haven't already, create two directories, called '/var/snd.master' and '/var/snd.standby'. They should both be owned by the Rivendell user, as well as readable and writable by that user. Move the contents of the original '/var/snd' directory to '/var/snd.master', and then delete '/var/snd'.

## Start It Up

To generate the initial master configuration, enter the following (as root):

```
rd_config master
```

This will generate the appropriate symbolic links and then start up the Rivendell daemons.

## SWITCHING BETWEEN CONFIGURATIONS

To switch between the Master and Standby configurations, all you need to do is:

Close any open Rivendell programs.

Do the following command (as root):

```
rd_config <config-name>
```

where <config-name> is either master or standby. This command will shut down the Rivendell daemons, regenerate the appropriate sym links for the selected configuration and then restart the daemons. **IMPORTANT NOTE:** running this command *will cause any active audio operations (playout or recording) to be interrupted.*

## REVISION HISTORY

Thu 19 Oct 2006 – Fred Gleason <[fredg@salemradiolabs.com](mailto:fredg@salemradiolabs.com)>  
Fixed missing '-x' switch provided to tar(1) command  
(reported by Andrew Struiksma <[andy@kpiradio.com](mailto:andy@kpiradio.com)>).

Sun 03 Sep 2006 – Fred Gleason <[fredg@salemradiolabs.com](mailto:fredg@salemradiolabs.com)>  
Added sections on generating and installing initial data snapshot.

Fri 28 Jan 2005 – Fred Gleason <[fredg@salemradiolabs.com](mailto:fredg@salemradiolabs.com)>  
Added directions on starting the rsyncd daemon.

Thu 04 Nov 2004 – Fred Gleason <[fredg@salemradiolabs.com](mailto:fredg@salemradiolabs.com)>  
Initial document creation.