
Backups and Restoring

Joshua Ferraro <jmf at kados dot org>

2004-12-15

Revision History

2004-12-13

Revision 2.2.0

jmf

Changes made by Joshua Ferraro. Original Document

Table of Contents

1. Using this document	1
1.1. Copyright and License	1
1.2. Disclaimer	2
2. Introduction	2
3. Backing up and Restoring the Koha Database	2
3.1. Backups	2
3.2. Restoration	3
3.3. Role of Replication as Backup	3
4. Backing up the Koha Filesystem	3
4.1. Automating the Filesystem Backup	3
5. Backing up the Local Operating System Configuration	4
6. Storing your Backup Offsite Using scp	4

TO DO:

- Add more script examples for automation
- Discuss replication in more detail

1. Using this document

1.1. Copyright and License

Copyright 2004 Joshua Ferraro <jmf at kados dot org>

This document is related to Koha and is licensed to you under the GNU General Public License version 2 or later (<http://www.gnu.org/licenses/gpl.html>). Koha-related documents may be reproduced and distributed in whole or in part, in any medium physical or electronic, as long as this copyright notice is retained on all copies.

You may create a derivative work and distribute it provided that you:

1. License the derivative work with this same license, or the Linux Documentation Project License (<http://www.tldp.org/COPYRIGHT.html>). Include a copyright notice and at least a pointer to the license used.

2. Give due credit to previous authors and major contributors.

Commercial redistribution is allowed and encouraged; however, the author would like to be notified of any such distributions.

1.2. Disclaimer

No liability for the contents of this document can be accepted. Use the concepts, examples and information at your own risk. There may be errors and inaccuracies, that could be damaging to your system. Proceed with caution, and although this is highly unlikely, the author(s) do not take any responsibility.

All copyrights are held by their by their respective owners, unless specifically noted otherwise. Use of a term in this document should not be regarded as affecting the validity of any trademark or service mark. Naming of particular products or brands should not be seen as endorsements.

2. Introduction

Understanding how to properly backup and restore your Koha installation is a critical step in administering Koha successfully. Not only are these skills needed in an emergency, such as a hard drive failure on your primary production machine, but they are also useful for when you update, upgrade, or anytime you plan to custom modify either the filesystem or the database.

There are three main components of Koha that must be considered when backing up and restoring: the database, the Koha filesystem, and the operating system customizations for the system. This document describes every step necessary for backing up and restoring both the database and the filesystem and points out some important considerations when backing up the operating system customizations -- these customizations often vary widely from system to system.

3. Backing up and Restoring the Koha Database

3.1. Backups

There are many ways to backup MySQL databases, or parts of databases. After reading this section I highly recommend reading the MySQL manual sections that discuss backing up and restoring data. You can find them online at <http://dev.mysql.com/doc/mysql/en/index.html>.

Perhaps the best way to backup your Koha database manually is the **mysqldump** application. To backup the whole database you can invoke it thusly:

```
[root@frodo]# mysqldump --add-drop-table -uroot -pyour password Koha > 2004_11_30_koha.sql
```

Notice that this naming convention uses the date, which makes it easy to organize a number of backups in a directory without being confused about which backup corresponds to which date. It's also clear that it's an SQL backup. Choose your naming conventions wisely and you'll thank yourself when you run into trouble.

Also notice the **--add-drop-table** option. This specifies that when you restore the data you will not append the data to existing data; additionally, if tables in your backup do not exist in the database they will be created.

If you plan to work on a specific table and you'd like to back up just that table you can use **mysqldump**

to do that as well. Suppose you're working on the `issues` table. You can back it up thusly:

```
[root@frodo]# mysqldump --add-drop-table -uroot -pyour password Koha.issues > 2004_11_30_koha.issues.sql
```

For more information on backups using the `mysqldump` utility please see the MySQL manual section 8.8: The `mysqldump` Database Backup Program [<http://dev.mysql.com/doc/mysql/en/mysqldump.html>]. For more ways to backup your database, the MySQL manual is your best reference.

Now you can use the `gzip` program to compress the sql file and reduce it's overall size for easier transport:

```
[root@frodo]# gzip 2004_11_30_koha.sql
```

3.2. Restoration

To unzip your file (in case you need to restore your database) you simply type:

```
[root@frodo]# gzip -d filename.tgz
```

An sql file can be used to restore data using the MySQL command-line tool. Here's an example of how to restore the Koha database using an sql file called `koha.sql`:

```
[root@frodo]# mysql -uroot -pyour password Koha < koha.sql
```

To restore a single table such as `issues` from a file called `koha.issues.sql` you would type:

```
[root@frodo]# mysql -uroot -pyour password Koha.issues < koha.issues.sql
```

3.3. Role of Replication as Backup

In addition to being useful in load balancing, replication can serve as a realtime backup for your Koha system. You will need to manually promote a slave to serve as master, and you will likely need to adjust your local configurations depending on how you have your systems setup, but having replication in place will probably ensure the fastest route to recovery from a complete system crash. See the replication section of the MySQL manual for more details.

4. Backing up the Koha Filesystem

Even if you don't plan to use the latest CVS updates for your installation, it might be worth your time to centralize your Koha installation filesystem to ease the backup and restore process. For more info on how to achieve that see the *Updating Koha* section of the Koha documentation.

To backup the Koha installation directory tree you need to ensure that all your Koha files are being considered. Most of them are located by default in `/usr/local/koha`. The `koha-httpd.conf` and the `koha.conf` files are located in the `/etc/` directory.

4.1. Automating the Filesystem Backup

Here's a short Perl script that I run with `cron` every evening. It backs up the filesystem using `tar` and `gzip`. You'll need to change some of the Configurable Variables for the script to work in your environment. Other than that, it's pretty much ready to go -- setup `cron` to run it nightly.

```
#!/usr/bin/perl -w
#This script assumes that you have set up a CVS repository
#that is symlinked to your installation directories. For
#info on setting that up see the Koha manual "Updating Koha"
#Configurable Variables
my $koha_base = "/build/websites/openils.com/koha"; #location of Koha cvs repo
my $bak_loc = "/build/backups"; #where to put backups
my ($sec,$min,$hour,$mday,$mon,$year,$wday,$yday,$isdst)
    = localtime(time);
$year += "1900";
my $date = "$year-$mon-$mday";
#Backup the filesystem
print "Backing up the filesystem\n";
print "Please wait ...";
`rm -rf $koha_base.bak`; #remove the previous backup
`cp $koha_base $koha_base.bak`; #make a new backup
`tar -cvzf $bak_loc/$date-kohafilesystem.tgz $koha_base.bak`; #targzip it
```

5. Backing up the Local Operating System Configuration

It's best to write a script to perform all the local operating system configurations as you would manually perform them to ensure that in the event of an emergency, you can simply run the script to set everything the way you had it before.

The list of possible custom configuration files for your system is unlimited, so I can just list a few common examples that you might want to consider including as part of your regular backup scheme:

- `httpd.conf`
- `my.cnf`
- any custom **cron** jobs related to Koha

6. Storing your Backup Offsite Using `scp`

You may want to keep a copy of your database, filesystem and operating system customizations on your production server to ease the process of restoring should the need arise. However, what if your hard drive on that machine crashes? It's a very good idea to store your backups on a different machine (or two or three) and it's probably a good ideal to store them in a different building too.

Perhaps the best way to avoid the problems inherent in physical backup medium such as a tape (that you may never really be sure is going to work anyway) is to use **ssh** and in particular **scp** as a part of your backup process. Why mess with tapes when you can securely copy a file from one system to another? You can invoke **scp** thusly:

```
[root@frodo]# scp koha.bak jmf@sam.openils.com:/build/backups/
```

The `:` signifies that this is a network address and also allows a specific directory to be specified on the receiving machine -- in this case, `sam.openils.com` has a `/build/backups` directory where I store backups.