

Daily Backup Script

Create a backup script to run every night and backup:

- **MySQL database**
- `/var/www/BookStack` folder
- **Nginx configuration**

Prerequisites

Prepare directory structure

Before running any script, make sure all the directories exist, otherwise the script will fail (it doesn't check nor creates them).

This is the structure of the backup directory:

```
/var/  
├─ backup  
│   └─ bookstack  
│       ├─ db  
│       └─ files  
└─ nginx
```

This is the structure of the log directory:

```
/var/  
├─ log  
│   └─ bookstack  
│       └─ backup_script
```

To make these paths quickly, just add `-p` to `mkdir`, so it creates all folders along the way.

```
$ sudo mkdir -p /var/backup/bookstack/{db,files,nginx}
```

```
$ sudo mkdir -p /var/log/bookstack/backup_script
```

Adjust permissions

Both directories, especially the backup one will contain sensitive files like the **.env configuration file** etc. For this reason, put **0600 umask to all directories and files**. `-R` will make sure all directories under the one we specified will have their permissions changed. Use `-v` to see what dirs and files were changed exactly.

```
$ sudo chmod -R 600 /var/log/bookstack
```

```
$ sudo chmod -R 600 /var/backup/bookstack/
```

Note that you can create your own directory structure, but make sure to adjust the script accordingly.

Create MySQL configuration file

To backup a MySQL database, you'll utilize `mysqldump` tool. You would normally use it with the `-p` option to specify a password for a user of a database. This however, passes the password in plain text as an option to a command, therefore will be visible to any user who runs `ps aux`. **You want to avoid this, because it's a serious security issue.**

Here's where `my.cnf` comes into play. `my.cnf` is a MariaDB configuration file. MariaDB actually has multiple configuration files in a few directories, `my.cnf` is usually used for user-specific settings. By utilizing `my.cnf`, we can put the password in it, adjust permissions, and then point `mysqldump` to it. This way, the password is hidden in a well-protected file.

The script will run as root and since `my.cnf` will store a password, it should only be readable by root. Pick a directory for the file. I will go for `/root/.my.cnf`. Create it and edit permissions accordingly:

```
$ sudo touch /root/.my.cnf
$ sudo chmod 600 /root/.my.cnf
```

Open it with some editor and add the following to it. Replace `{password}` with a password (without `{ }`) for the user that will be used to access (backup) the database.

```
$ sudo vi /root/.my.cnf
```

```
[mysqldump]
password={password}
```

Create the scripts

Main script

Due to the way I decided to implement logging for this script, I have to divide it in two. It is completely possible to have just one script, so adjust it to your liking. The script will be run as **root**, because it's writing to privileged directories. **You are free to adjust permissions, directories, users etc.**

RUN_bookstack_backup.sh

```
#!/bin/bash
LOG_DIR=/var/log/bookstack/backup_script
CURRENT_DATE=$(date +"%Y-%m-%d")

# Run bookstack_backup_worker.sh with root privileges, pipe it to gawk which puts a timestamp before every
line and writes to file
source ./bookstack_backup_worker.sh | gawk '{ print strftime("[%Y-%m-%d %H:%M:%S %Z]", $0 ) }' >
$LOG_DIR/bookstack-backup_${CURRENT_DATE}.log
```

- `#!/bin/bash` – The so called *shebang*, every bash script should start with one. Read more about it [here](#).
- `LOG_DIR=/var/log/bookstack/backup_script` – Creates a variable `LOG_DIR` specifying the location where to save logs. Simplifies scripts, so you don't have to type out long paths multiple times.
- `CURRENT_DATE=$(date +"%Y-%m-%d")` – Creates a variable `CURRENT_DATE`, which runs a command `date` with options `+"%Y-%m-%d"`. This gives you a basic ISO date in `YYYY-MM-DD` format, separated with dashes. We will append this to the backup filenames.
- `source ./bookstack_backup_worker.sh` – Run the backup script with `source`. Source runs the other script in a sub-shell and that allows us to share variables like `CURRENT_DATE` with the other script. The entire script is piped (`|`) to `gawk`.
- `gawk '{ print strftime("[%Y-%m-%d %H:%M:%S %Z]", $0) }'` – This is my timestamp logging solution. `gawk` takes the output of the entire `bookstack_backup_worker.sh` script and adds a timestamp at the beginning of every line. This timestamping solution is described in more detail in this [TODO](#) guide. After `gawk` adds the timestamp, it sends it to a log file with `>`. The log file's path is specified using the `LOG_DIR` variable and has `CURRENT_DATE` inserted into its name. The entire path and filename then might look like this:
`/var/log/bookstack/backup_script/bookstack-backup_2021-09-04.log`
- To run this script, simply type `sudo ./RUN_bookstack_backup`

Slave script

bookstack_backup_worker.sh

```
#!/bin/bash
BOOKSTACK_DIR=/var/www/BookStack
NGINX_DIR=/etc/nginx
BACKUP_DIR=/var/backup/bookstack
DB_BACKUP_DIR=$BACKUP_DIR/db
WEBROOT_BACKUP_DIR=$BACKUP_DIR/files
NGINX_BACKUP_DIR=$BACKUP_DIR/nginx

exec 2>&1

# MYSQL DATABASE BACKUP
echo "Starting BACKUP SCRIPT..."
echo "Starting MySQL backup..."
echo "Backing up to $DB_BACKUP_DIR..."
mysqldump --defaults-extra-file=/root/.my.cnf -v -u user database | gzip -vc > $DB_BACKUP_DIR/bookstackdb-
backup_$(CURRENT_DATE).sql.gz
echo "Done..."

# WEBSERVER BACKUP
# Archive and compress BookStack webroot folder and save it to backup location with current date
echo "Backing up BookStack webroot directory to $WEBROOT_BACKUP_DIR..."
tar -czvf $WEBROOT_BACKUP_DIR/bookstack-backup_$(CURRENT_DATE).tar.gz $BOOKSTACK_DIR
echo "Done..."

# NGINX CONFIG BACKUP
# Archive and compress Nginx config folder and save it to backup location with current date
echo "Backing up Nginx to $NGINX_BACKUP_DIR..."
tar -czvf $NGINX_BACKUP_DIR/nginx-backup_$(CURRENT_DATE).tar.gz $NGINX_DIR
echo "Done..."
echo "Finished..."
```

VARIABLE PREPARATION

- `#!/bin/bash` – See above.
- `BOOKSTACK_DIR=/var/www/BookStack` – Creates variable `BOOKSTACK_DIR` pointing to the directory where BookStack is stored.
- `NGINX_DIR=/etc/nginx` – Points to the Nginx configuration folder.
- `BACKUP_DIR=/var/backup/bookstack` – Points to the directory where all BookStack backups will be saved.
- `DB_BACKUP_DIR=$BACKUP_DIR/db` – Points to the database backup directory inside of the main backup directory. This variable uses the previously created `BACKUP_DIR` to make it

shorter.

- `WEBROOT_BACKUP_DIR=$BACKUP_DIR/files` - Backup dir for the BookStack files.
- `NGINX_BACKUP_DIR=$BACKUP_DIR/nginx` - Backup dir for Nginx configuration files.

`exec 2>&1` - Redirect `stderr (2)` to `stdout (1)`. This means all errors and normal messages will be redirected to standard output (stdout), which is what we would normally see in a terminal. All messages generated by this script will then go from stdout to the master script, which then redirects it to `gawk` using pipe `|` (described in the master script).

MYSQL DATABASE BACKUP

- `echo` - Prints to the stdout
- `mysqldump` - Utility used to backup a MySQL/MariaDB database
 - `--defaults-extra-file=/root/.my.cnf` - Points to the configuration file explained above
 - `-v` - Enable verbose output
 - `-u user` - User that will access (backup) the database (needs appropriate privileges)
 - `database` - Name of the database to backup
- `| gzip -vc > $DB_BACKUP_DIR/bookstackdb-backup_$CURRENT_DATE.sql.gz`
 - `|` - pipe output of `mysqldump` (the database backup file) to `gzip`
 - `gzip` - Reduce the size of the backup by compressing it
 - `-v` - Enable verbose output
 - `-c > $DB_BACKUP_DIR/bookstackdb-backup_$CURRENT_DATE.sql.gz` - Use of variables shortens the string and inserts current date into the filename. `.sql.gz` explains that it's a `.sql` file compressed with `gzip`.

WEBSERVER BACKUP

- `tar -czvf $WEBROOT_BACKUP_DIR/bookstack-backup_$CURRENT_DATE.tar.gz $BOOKSTACK_DIR`
 - `tar` - Command used to create archives and compress them
 - `-c` - Create an archive
 - `-z` - Compress with `gzip`
 - `-v` - Enable verbose output (print files and directories being backed up)
 - `-f` - Has to be always the last option, after which goes the name of the archive to be created
 - `$WEBROOT_BACKUP_DIR` - Location of the web backup directory
 - `/bookstack-backup_$CURRENT_DATE.tar.gz` - Filename with current date in `.tar.gz` (compressed archive) format. `/` at the beginning is appended after `$WEBROOT_BACKUP_DIR` to form a full path to the file, e.g. - `/var/backup/bookstack/files/bookstack-backup_2021-09-24.tar.gz`
 - `$BOOKSTACK_DIR` - Tells `tar` which directory to backup

NGINX CONFIG BACKUP

- `tar -czvf $NGINX_BACKUP_DIR/nginx-backup_$CURRENT_DATE.tar.gz $NGINX_DIR`
 - Same as above, just with different paths and variables

You can now run manual backups. Logs will be located in `/var/log/bookstack/backup_script` . Use cron to automate it.

Add to cron

In order to run the script daily, setup `cron` to do it for you.

Save scripts to a location

These scripts will run as `root`, save them to a safe location and give them appropriate privileges (700).

Create the directory:

```
$ sudo su
(root)$ mkdir -p /root/scripts/bookstack_backup
```

Copy the scripts from an old directory:

```
(root)$ cd /root/scripts/bookstack_backup
(root)$ cp /home/user/bookstack_backup_worker.sh /home/user/RUN_bookstack_backup.sh .
```

Adjust permissions:

```
(root)$ chmod 700 bookstack_backup_worker.sh RUN_bookstack_backup.sh
```

Modify the script

Take a look at this script that we are about to add to crontab. Do you see anything wrong?

```
#!/bin/bash
LOG_DIR=/var/log/bookstack/backup_script
CURRENT_DATE=$(date +"%Y-%m-%d")

# Run bookstack_backup_worker.sh with root privileges, pipe it to gawk which puts timestamp before every line
and writes to file
source ./bookstack_backup_worker.sh | gawk '{ print strftime("[%Y-%m-%d %H:%M:%S %Z]", $0 )' >
$LOG_DIR/bookstack-backup_${CURRENT_DATE}.log
```

Don't worry if you don't, I also hadn't seen anything when I looked at it before.

Remember how we ran this script previously? We had them both in the same directory and typed `./RUN_bookstack_backup.sh`. Inside of this script we called `source` with `./` again. This works standalone, but try adding it to cron and you will sooner or later realize the script is **not working** and the reason for it is `./`. There might be other reasons why some scripts work standalone, but don't in cron, but this is what helped in my case. Generally, cron doesn't like relative paths and user specific things, especially if you run the script as another user. Replacing `./bookstack_backup_worker.sh` with a proper full path to the other script solved my problems:

```
#!/bin/bash
LOG_DIR=/var/log/bookstack/backup_script
CURRENT_DATE=$(date +"%Y-%m-%d")

# Run bookstack_backup_worker.sh with root privileges, pipe it to gawk which puts timestamp before every line
and writes to file
source /root/scripts/bookstack_backup/bookstack_backup_worker.sh | gawk '{ print strftime("[%Y-%m-%d
%H:%M:%S %Z]", $0 )' > $LOG_DIR/bookstack-backup_$CURRENT_DATE.log
```

Tip: Use `*****` in `/etc/crontab` to run the script every minute to troubleshoot quicker.

Edit crontab

If you aren't too afraid of messing up and you want to run the scripts as root, you can open `/etc/crontab` directly

```
(root)$ vi /etc/crontab
```

Explaining what options you have when setting a cron job is out of the scope of this guide. To run the script every day at 4AM, put the following line at the end of the file:

```
0 4 *** root /root/scripts/bookstack_backup/RUN_bookstack_backup.sh
```

Use <https://crontab.guru> to convert your desired time to cronjob command.

Revision #15

Created 22 September 2021 13:41:07 by Marek

Updated 27 September 2021 00:26:50 by Marek