

# Using md5deep for Comparing Directories in Unix

You can compare the contents of two directories by their md5 hashes, which could be useful when you want to make sure that a sync operation went smoothly, for instance. By inspecting the hashes of all the files in the directory and confirming they're identical you can rest assured all data was copied successfully and fully.

You can use `md5sum` to get the md5 sums of all the files in a directory, but comparing like this could be pretty daunting:

```
md5sum dir/*
```

This outputs a list of all files with their md5 sums.

A better way is using `md5deep` instead. If you don't have it you can most likely install it using your package manager very easily (*`sudo apt-get install md5deep`* on Ubuntu).

Then if you run the following you'll get a list of md5 sums of all files in the directory as well as the files of sub-directories:

```
md5deep -r dir/
```

The real solution is in the ability of `md5deep` to compare its own outputs. First you get the md5 sums in a file:

```
md5deep -r -s /dir1> dir1sums
```

And then have `md5deep` read that file and compare the second directory to it:

```
md5deep -r -X dir1sums /dir2
```

If there is no output that means the directories are identical. Otherwise it will display the hashes of files that are different. Thus the comparison has been accomplished.