

# Linux Fu: Fusing Hackaday

[Originalartikel](#)

[Backup](#)

<html> <p>Unix and, by extension, Linux, has a mantra to make everything possible look like a file. Files, of course, look like files. But also devices, network sockets, and even system information show up as things that appear to be files. There are plenty of advantages to doing that since you can use all the nice tools like

grep

and

find

to work with files. However, making your own programs expose a filesystem can be hard. Filesystem code traditionally works at the kernel module level, where mistakes can wipe out lots of things and debugging is difficult. However, there is FUSE &#8212; the file system in user space library &#8212; that allows you to write more or less ordinary code and expose anything you want as a file system. You&#8217;ve probably seen FUSE used to mount, say, remote drives via ssh or Dropbox.

We&#8217;ve even looked at FUSE before, even <a href=„<https://hackaday.com/2021/08/31/linux-fu-user-space-file-systems-now-for-windows-too/>“>for Windows</a>.</p><p>What&#8217;s missing, naturally, is the Hackaday RSS feed, mountable as a normal file. And that&#8217;s what we&#8217;re building today.</p><p><a href=„<https://hackaday.com/wp-content/uploads/2022/02/mount.png>“ target=„\_blank“><img data-attachment-id=„519698“ data-permalink=„<https://hackaday.com/2022/02/16/linux-fu-fusing-hackaday/mount-2/>“ data-original-file=„<https://hackaday.com/wp-content/uploads/2022/02/mount.png>“ data-original-size=„982,702“ data-comments-opened=„1“ data-image-meta=„{&quot;aperture&quot;:&quot;0&quot;,&quot;credit&quot;:&quot;,&quot;camera&quot;:&quot;,&quot;caption&quot;:&quot;,&quot;created\_timestamp&quot;:&quot;0&quot;,&quot;copyright&quot;:&quot;,&quot;focal\_length&quot;:&quot;0&quot;,&quot;iso&quot;:&quot;0&quot;,&quot;shutter\_speed&quot;:&quot;0&quot;,&quot;title&quot;:&quot;,&quot;orientation&quot;:&quot;0&quot;};“ data-image-title=„mount“ data-image-description=„“ data-image-caption=„“ data-medium-file=„<https://hackaday.com/wp-content/uploads/2022/02/mount.png?w=400>“ data-large-file=„<https://hackaday.com/wp-content/uploads/2022/02/mount.png?w=800>“ class=„alignright wp-image-519698 size-medium“ src=„<https://hackaday.com/wp-content/uploads/2022/02/mount.png?w=400>“ alt=„“ width=„400“ height=„286“ srcset=„<https://hackaday.com/wp-content/uploads/2022/02/mount.png> 982w, <https://hackaday.com/wp-content/uploads/2022/02/mount.png?resize=250,179> 250w, <https://hackaday.com/wp-content/uploads/2022/02/mount.png?resize=400,286> 400w, <https://hackaday.com/wp-content/uploads/2022/02/mount.png?resize=800,572> 800w“ referrerpolicy=„no-referrer“ /></a>Writing a FUSE filesystem isn&#8217;t that hard, but there are a lot of tedious jobs. You essentially have to provide callbacks that FUSE uses to do things when the operating system asks for them. Open a file, read a file, list a directory, etc. The problem is that for some simple projects, you don&#8217;t care about half of these things, but you still have to provide them.</p><p>Luckily, there are libraries that can make it a lot easier. I&#8217;m going to show you a simple C++ program that can mount your favorite RSS feed (assuming your favorite one is

Hackaday, of course) as a file system. Granted, that's not amazing, but it is kind of neat to be able to grep through the front page stories from the command line or view the last few articles using Dolphin.

**Pick a Library**

There are plenty of libraries and wrappers around FUSE. I picked one by [jachappell] over on [GitHub](https://github.com/jachappell/Fusepp). It was pretty simple and hides just enough of FUSE to be handy, but not so much as to be a problem. All the code is hidden around in

**Fuse.h**

. One thing to note is that the library assumes you are using

**libfuse**

3.0. If you don't already have it, you'll have to install the

**libfuse**

3.0 development package from your package manager. There are other choices of libraries, of course, and you could just write to the underlying

**libfuse**

implementation, but a good library can make it much simpler to get started.

Just to keep things simple, I [forked the original project on GitHub](https://github.com/wd5gnr/Fusepp) and added a [fusehad](https://github.com/wd5gnr/Fusepp/tree/master/fusehad) directory.

**Constraints**

To keep things simple, I decided not to worry about performance too much. Since the data is traveling over the network, I do attempt to cache it, and I don't refresh data later. Of course, you can't write to the filesystem at all. This is purely for reading Hackaday.

These constraints make things easier. Of course, if you were writing your own filesystem, you might relax some of these, but it still helps to get something as simple as possible working first.

**Making it Work First**

Speaking of which, the first order of business is to be able to read the Hackaday RSS feed and pull out the parts we need. Again, not worrying about performance, I decided to do that with a pipe and calling out to

**curl**

. Yes, that's cheating, but it works just fine, and that's why we have tools in the first place.

The

**HaDFS.cpp**

file has a few functions related to FUSE and some helper functions, too. However, I wanted to focus on getting the RSS feed working so I put the related code into a function I made up called

**userinit**

. I found out the hard way that naming it

## init

would conflict with the library. The normal FUSE system processes your command line arguments; a good thing, as you'll see soon. So the main in

## HaD.cpp

is really simple: 

```
#include <stdio.h>#include „HaDFS.h“int main(int argc, char *argv[]){ HaDFS fs; if (fs.userinit()) { fprintf(stderr, „Can't fetch feed\n“); return 99; } int status; status= fs.run(argc, argv); return status;}
```

 However, for now, I simply commented out the line that calls

## fs.run

. That left me with a simple program that just calls

## userinit

. Reading the feed isn't that hard since I'm conscripting

## curl

. Each topic is in a structure and there is an array of these structures. If you try to load too many stories, the code just quietly discards the excess (see

## MAXTOPIC

). The

## topics

global variable tells how many stories we've actually loaded. 

```
 The popen function runs a command line and gives us the stdout stream as a bunch of lines. Processing the lines is just brute force looking for <title> and <link> to identify the data we need. I filtered curl through grep to make sure I didn't get a lot of extra lines, by the way, and I assumed lowercase, but a -i option could easily fix that. The redirect is to prevent curl from polluting stderr, although normally FUSE will disconnect the output streams so it doesn't really matter. Note that I add an HTML extension to each fake file name so opening one is more likely to get to the browser.
```

By putting a `printf` in the code I was able to make sure the feed fetching was working the way I expected. Note that I don't fetch the actual pages until later in the process. For now, I just

want the titles and the URL links.

**The Four Functions**

There are four functions we need to create in a subclass to get a minimal read-only filesystem going: `getattr`, `readdir`, `open`, and `read`. These functions pretty much do what you expect. The `getattr` call will return 755 for our root (and only) directory and 444 for any other file that exists. The `readdir` outputs entries for . and .. along with our files. `Open` and `read` do just what you think they do.

There are some other functions, but those are ones I added to help myself:

- `userinit` – Called to kick off the file system data
- `trimrss` – Trim an RSS line to make it easier to parse
- `pathfind` – Turn a file name into a descriptor (an index into the array of topics)
- `readurl` – Return a string with the contents of a URL (uses curl)

There's not much to it. You'll see in the code that there are a few things to look out for like catching someone trying to write to a file since that isn't allowed.

**Debugging and Command Line Options**

Of course, it doesn't matter how simple it is, it isn't going to work the first time is it? Of course, first, you have to remember to put the call to `fs.run` back in the main function. But, of course, things won't work like you expect for any of a number of reasons. There are a few things to remember as you go about running and debugging.

When you build your executable, you simply run it and provide a command line argument to specify the mount point which, of course, should exist. I have a habit of using `/tmp/mnt` while debugging, but it can be anywhere you have permissions.

Under normal operation, FUSE detaches your program so you can just kill it. You'll need to use `umount` command (`fusermount -u`) with the mount point as an argument. Even if your program dies with a segment fault, you'll need to use the `umount` command or you will probably get the dreaded "disconnected endpoint" error message.

Being detached leads to a problem. If you put `printf` statements in your code, they will never show up after detachment. For this reason, FUSE understands the `-f` flag which tells the system to keep your filesystem running in the foreground. Then you can see messages and a clean exit, like a `Control+C`, will cleanly unmount the filesystem. You can also use `-d` which enables some built-in debugging and implies `-f`. The `-s` flag turns off threading which can make debugging easier, or harder if you are dealing with a thread-related problem.

You can use `gdb`, and there are some [good articles](https://blog.jeffli.me/blog/2014/08/30/use-gdb-to-understand-fuse-file-system/) about that. But for such a simple piece of code, it isn't really necessary.

**What's Next?**

The documentation for the library is almost nothing. However, the library closely mirrors the `libfuse` API so the documentation for that (mostly in [the source code](http://libfuse.github.io/doxygen/fuse-3_810_84_2include_2fuse_8h.html)) will help you go further. If you want to graduate from FUSE to a real file system, you have a long road. The video below gives some background on Linux VFS, but that's just the start down that path.

**Cloud Storage**

Maybe stick to FUSE for a while. If you prefer Python, [no problem](https://hackaday.com/2013/11/06/writing-a-fuse-filesystem-in-python/). FUSE is very popular for mapping [cloud storage into your](https://hackaday.com/2020/11/10/linux-fu-send-in-the-cloud-clones/)

*filesystem</a>, but with your own coding, you could just as easily expose your Arduino or anything else your computer can communicate with.</p> </html>*

From:

<https://schnipsl.qgelm.de/> - **Qgelm**



Permanent link:

[https://schnipsl.qgelm.de/doku.php?id=wallabag:wb2linux-fu\\_-fusing-hackaday](https://schnipsl.qgelm.de/doku.php?id=wallabag:wb2linux-fu_-fusing-hackaday)

Last update: **2025/06/27 11:17**