

GmailFS: Using a mail account as a filesystem

MORE THAN MAIL

Why let a 2 Gigabyte email account go to waste? GmailFS lets you use Gmail as a storage medium for a mountable filesystem.

BY FABRIZIO CIACCHI

Google's Gmail webmail system [1] has been heralded as a new vision for Internet mail. A Gmail account comes with a full 2 GB of available file storage. According to the company, a Gmail account means you never have to throw away any messages.

Of course, offering 2 GB of storage to a Linux user only opens the door for more innovation. It wasn't long before the Linux community had a whole new use for this 2 GB storage space that was completely different from Google's original intention. GmailFS, created by Richard Jones [2], uses the 2 GB Gmail storage space as a network-based, mountable Linux filesystem.

What is GmailFS?

GmailFS appears to the user as an ordinary filesystem. According to its creator, "GmailFS supports most file operations, such as read, write, open, close, stat, symlink, link, unlink, truncate, and rename." You can use ordinary Linux command line tools with GmailFS, such

as cp, mv, rm, and grep. Behind the scenes, though, GmailFS saves data in the form of mail messages on the Gmail server. (See the box titled "How GmailFS Works" for more on this highly unusual file storage process.)

GmailFS provides a ready-made network storage system for users who do not want the complication of implementing their own Internet-accessible file server. Using GmailFS, you can access a centrally-stored file from different computers at different locations without ever having to worry about the server side of the connection. For instance, you can use Gmail space to upload some files from work and then download them again at home, or you can make a particular program accessible from any-

where. GmailFS can also serve as a quick and easy tool for informal offsite file backup.

Preparing Your System

GmailFS is built on some essential background components. Before you use GmailFS, you must install:

- FUSE – a tool designed to support Linux filesystems in user space;
- libgmail – a python library that provides access to Gmail.

GmailFS is written in Python, so you also need the Python 2.3 (not 2.4) binary and development files. Many distributions come with Python, but if yours doesn't have it, the box titled "Installing Python" provides a quick summary of how to install Python.

Download version 1.3 of FUSE from

```
root@debbi:/home/fabrizio# cat /etc/fstab
# /etc/fstab: static file system information.
#
#-file system-      <mount point>  <type>  <options>      <dump>  <pass>
proc                /proc          proc    defaults       0        0
/dev/hd1            /              ext3    defaults       1        1
/dev/hdb            none          none    no             0        0
/dev/hdc            /media/cdrom0  iso9660 ro,user,noauto 0        0
/dev/fd0            /media/floppy0 auto    rw,user,noauto 0        0
/dev/fd1            /media/floppy1 auto    rw,user,noauto 0        0
/dev/hda1           /mnt/hda1     vfat    rw,user,noauto 0        0
/usr/local/bin/gmailfs.py /mnt/gmail    gmailfs rw,user,noauto 0        0
root@debbi:/home/fabrizio# df
Filesystem            blocks  bsize  fr  fs  Inodes Disavail  Used  Mounted on
/dev/hd1              3899040  4096  0  ext3  822240  80% /
tmpfs                 47176  0  0  tmpfs  0  0% /dev/shm
/dev/hda1             9249088  4096  0  ext3  9216248  1% /media/hda1
/dev                 3899040  4096  0  ext3  822240  80% /dev
none                  10240  4096  0  none  7468  28% /dev
/usr/local/bin/gmailfs.py 1027568  4096  0  gmailfs  958484  7% /mnt/gmail
root@debbi:/home/fabrizio#
```

Figure 1: /etc/fstab and the output of the df command.

[7], then extract the files and install FUSE. Note: to complete this operation, you need to have compiled your kernel from source so FUSE will install the proper module for your installed kernel. Type in `su` to enter a root password and enter the following:

```
# tar -zxvf fuse-1.3.tar.gz
# cd fuse-1.3/
# sh ./configure
# make
# make install
```

Next you need to install the Python FUSE bindings. These bindings permit you to use FUSE with Python code.

You can grab the Python FUSE bindings from FUSE's CVS site [8] or from the website of GmailFS author Richard Jones [9]; after you have downloaded the file, `su` to root and enter the following commands:

```
# tar -zxvf fuse-python.tar.gz
# cd fuse-python/
# python setup.py build
# python setup.py install
```

The hardest part is now done. The next step is to obtain the CVS version of libgmail [4]; you'll need the `constant.py` and `libgmail.py` files.

To obtain these files, you can download them from CVS with your browser [10]. After you have downloaded the lat-

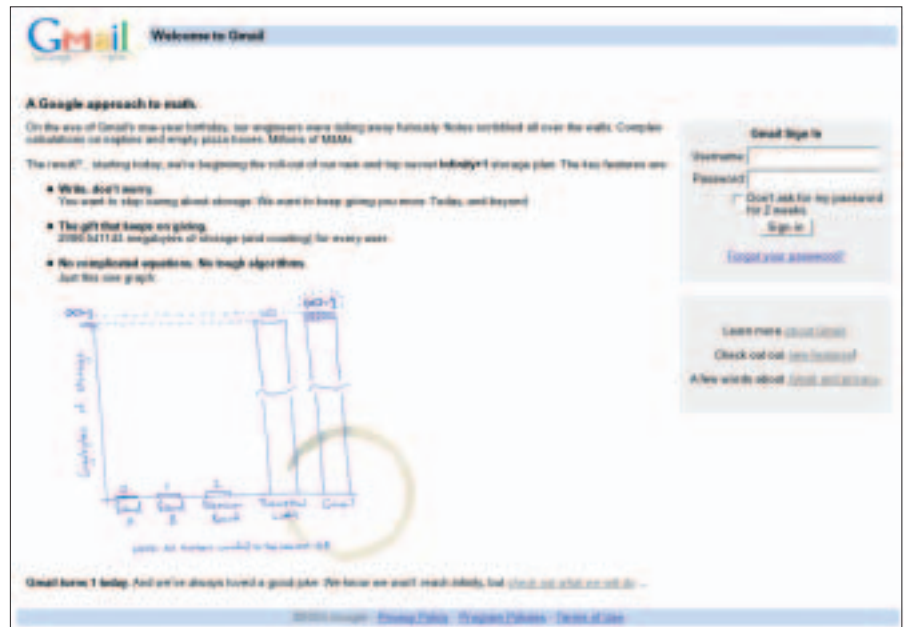


Figure 2: Gmail was launched last year on April 1st with 1 GB of space. On this year's April 1st, space was extended to 2GB and is still growing to "Infinity+1".

est versions of these files, enter the necessary commands (as root) to move the libgmail files to a location on your system where Python can read them:

```
# cd libgmailCVS/
# mv libgmail.py constants.py /usr/local/lib/python2.3/site-packages/
```

Installing GmailFS

Now you are ready to install GmailFS. Open your browser, go to Richard Jones's website [2], and download the

latest release of GmailFS [11]. After completing the download, decompress the file and copy some files to the correct locations. Type in `su` to enter your root password and then:

```
# tar -zxvf gmailfs-0.3.tar.gz
# cd gmailfs-0.3/
# mv gmailfs.py /usr/local/lib
# mv mount.gmailfs /sbin
# mv gmailfs.conf /etc
```

The next step is to edit the `/etc/gmailfs.conf` file. This is the file in which you set parameters such as the Gmail account username and password, and the filesystem name. The filesystem name is very important because GmailFS will use the filesystem name to associate the messages on the server with the filesystem. Choose a hard-to-guess name. (We have used the default `linux_fs_3` in this example.) Because GmailFS groups messages according to the filesystem name, it is actually possible to create multiple filesystems for a single Gmail account.

To mount the GmailFS, you can use these commands:

```
# mkdir GmailMount
```

How GmailFS Works

GmailFS is a Python application that depends upon a pair of important background components. The first of these components is FUSE [3], a component that runs in kernel space to support the implementation of a filesystem in userspace. The other necessary ingredient is libgmail [4], a Python library that supports programmatic access to Gmail. GmailFS results from the combination of these components. FUSE allows GmailFS to operate in userspace, which greatly simplifies its design and implementation, and libgmail lets GmailFS

invisibly access a Gmail account to store and retrieve data. The result is that the user perceives GmailFS as an ordinary file system, and Gmail perceives the GmailFS file system as a collection of email messages.

The meta-data associated with the file is stored in the subject line of the mail message. GmailFS stores the actual file data in the form of an email attachment. Because Gmail limits the size of an email attachment, a large file may be stored in multiple attachments.

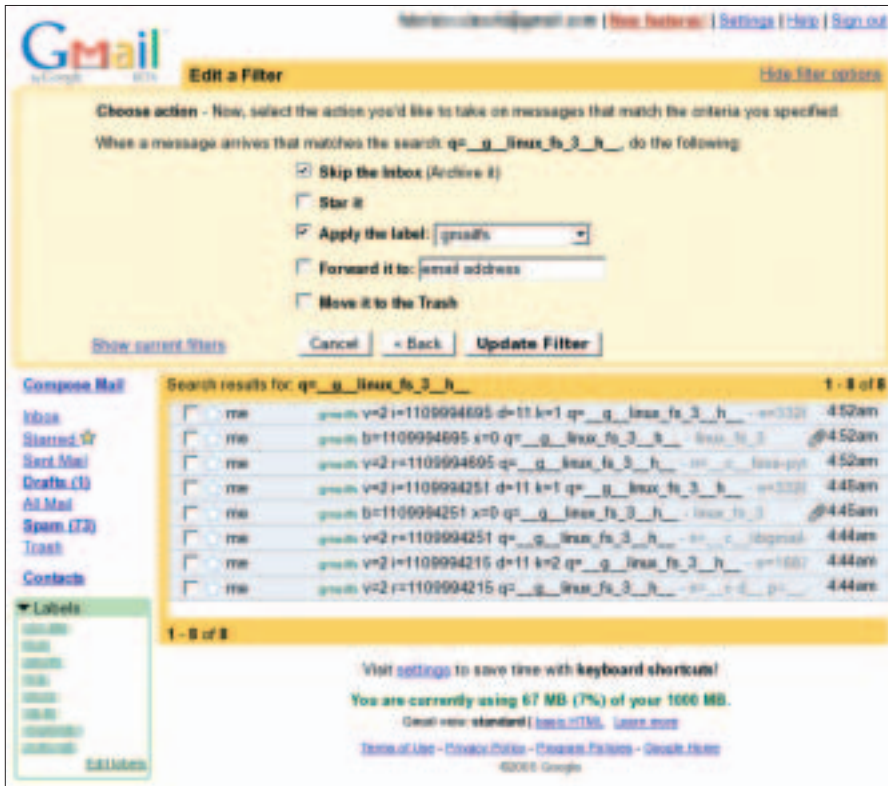


Figure 3: If you set up filtering correctly, the filesystem messages will appear in the summary under the filter box.

```
# mount -t gmailfs \
/usr/local/bin/gmailfs.py \
GmailMount/ -o username=[
gmailuser];password=[
gmailpass],
fsname=[fs_name]
```

where `[gmailuser]` and `[gmailpassword]` are your Gmail username and password, and `[fs_name]` is the filesystem name (Note: If you set the parameters in the `/etc/gmailfs.conf` file, you can omit the

bold code.). To automatically mount the Gmail file system, you need to modify your `/etc/fstab` file. In the example, we added the line that acts as the mount command:

```
# nano /etc/fstab
- add the following line \
at the end of the file-
/usr/local/bin/gmailfs.py \
/path/of/mount/point \
gmailfs rw,user,noauto, \
```

Installing Python 2.3

The commands for installing Python 2.3 packages are as follows:

In Debian:

```
# apt-get install python2.3
python2.3-dev
```

Slackware

```
# swaket --install
python-2.3.4-i486-1
```

Gentoo

```
# emerge python
```

Mandrake

```
# urpmi libpython2.3
libpython2.3-devel
```

Fedora and Redhat (with Apt4Rpm [5])

```
# apt-get install
python-devel-2.3.4-13.1
(Redhat rpm and Yum packages are also
available.)
```

If you want to install from source, you need to download the Python 2.3.5 source file [6]. The following commands assume you have the development components necessary to compile. Type in `su` to a root password, and then:

```
# tar -zxvf Python-2.3.5.tgz
# cd Python-2.3.5/
# ./configure
# make
# make install
```

```
username=[gmailuser], \
password=[gmailpass], \
fsname=[fs_name]
- press CTRL+O to save \
and CTRL+X to exit -
```

(Note: If you set the parameters in the `/etc/gmailfs.conf` file, you can omit the bold code.)

Tuning Gmail

When you complete all these steps and reboot your computer, you will see that GmailFS is mounted in the proper directory. Figure 1 shows the `/etc/fstab` syntax and the output of the `df` command.

Now that you have a fully operative GmailFS file system, let's take a look at some fine-tuning. First of all, you need to create a label and a filter, so that your GmailFS messages will not be stored with your actual email messages.

Enter your Gmail account, and click on *Edit Labels* in the green box at the left side, then insert something like `gmailfs` in the text box and click Create.

Click on *Inbox* and select all the emails with *Me* as sender. Then click on *More Actions* and select *Apply Label - gmailfs*. This step is necessary because, in order to create the filter, you need to see the common characteristics of the emails that compose the file system.

Click on *Settings* then *Filter* and, at the end of the page, select *Create a new*

Compiling the kernel

To compile and install the kernel from source, you need to go to the Linux Kernel Archives [12] and download the latest stable version of the 2.6.x kernel series [13]. After you have downloaded the file in `.tar.gz` format, `su` to root and enter the following:

```
# mv linux-2.6.x.tar.gz /usr/src
# cd /usr/src
# tar -zxvf linux-2.6.x.tar.gz
# ln -s linux-2.6.x linux
# cd linux/
# make menuconfig
# make
# make install
# make modules_install
```

If you have some problems compiling the kernel, you can follow the Kernel-Build-HOWTO [14].

Advertisement

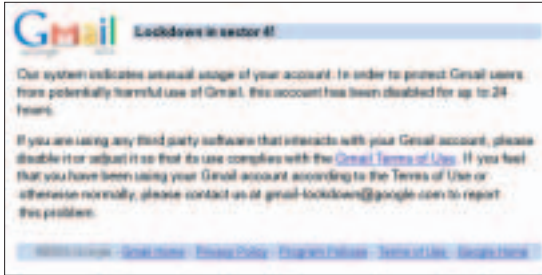


Figure 4: Using Gmail File System with a graphical environment like Gnome may result in your account being disabled for a while.

filter. Use the filesystem name we entered earlier to create the filter. The filter has the syntax `q = __g__[fs_name] __h__`. Click on Next Step.

If all is correct, you will see the file system emails in the summary under the filter box (Figure 2). Now check the options labeled *Skip the Inbox (Archive it)* and *Apply the label – gmailfs*. Click on Update Filter.

You should now be able to use your GmailFS filesystem without any problems. If you wish, you can use GmailFS with a graphical environment like GNOME, mounting it on request like an external USB device. We have met with some problems using GmailFS with a graphical environment; in some cases, the high frequency of access to the file system makes Gmail think it is experiencing a hacker attack.




Figure 5: GmailFS appears to the user as an ordinary filesystem.

Conclusion

One gigabyte of free network file storage accessible from anywhere on the planet? You may be saying, “This is too good to be true.” But before you rebuild your whole work environment around GmailFS, it is worth considering its limitations.

First, as you might guess, the whole concept of using email to store and retrieve files is inherently less efficient

than conventional file storage. Don't expect speedy performance from GmailFS. It is also worth remembering that GmailFS is something of an experiment. It isn't really intended for high-reliability production settings. Richard Jones, the creator of GmailFS, does not recommend storing the only copy of any important file on a GmailFS filesystem.

Because GmailFS does not support cryptography, you cannot really count on any privacy. Even if Google's security structures are sufficient to keep outside intruders from accessing your files, there is nothing to prevent Google itself from scanning your files for consumer and demographic information. (This is an issue with all free web-based mail systems – not just Gmail.)

If you can see your way around these limitations, however, you may find that you do indeed have uses for the very original GmailFS. Linux users have done something brilliant that could eventually lead to a conceptual revolution. Think about using the same files at home and at work without the need for a portable storage device such as a floppy, CD or USB memory stick.

Until the security and performance issues are addressed, I will still use my USB memory stick to move files around, but GmailFS may be a good alternative for many users. If you need a Gmail account, I will be glad to send you an invite. Send an email to fabrizio@ciacchi.it (I have 50 invites). ■

Gmail Tricks

For very busy people, it is possible to check the Gmail account with a mobile device [18]. If you are using the Mozilla Firefox browser you can also change the stylesheet [19] of the Gmail homepage, or use an extension to check for new mail [20] of your Gmail account without having a tab open for it. For Windows users, there is a program called Gmail Drive [21] that acts like GmailFS, and a program called Gmail notifier [22] that is launched at startup and puts itself in the tray.

INFO

- [1] Gmail homepage: <http://www.gmail.com>
- [2] Richard Jones's Website: <http://ltsv-141.richard.jones.name>
- [3] FUSE homepage: <http://fuse.sourceforge.net>
- [4] Libgmail homepage: <http://sourceforge.net/projects/libgmail/>
- [5] Apt4Rpm homepage: <http://apt.freshrpms.net>
- [6] Python 2.3.5 Download: <http://www.python.org/ftp/python/2.3.5/Python-2.3.5.tgz>
- [7] FUSE file list: http://sourceforge.net/project/showfiles.php?group_id=121684
- [8] FUSE's CVS website: <http://cvs.sourceforge.net/viewcvs.py/fuse/python/>
- [9] Pythons FUSE bindings: <http://richard.jones.name/google-hacks/gmail-filesystem/fuse-python.tar.gz>
- [10] Libgmail CVS: <http://cvs.sourceforge.net/viewcvs.py/libgmail/libgmail/>
- [11] GmailFS download: <http://ltsv-141.richard.jones.name/google-hacks/gmail-filesystem/gmailfs-0.3.tar.gz>
- [12] Linux Kernel Archives: <http://www.kernel.org>
- [13] Download Linux Kernel 2.6.x: <http://www.kernel.org/pub/linux/kernel/v2.6/>
- [14] Kernel-Build-HOWTO: <http://www.digitalhermit.com/linux/Kernel-Build-HOWTO.html>
- [15] Gallina homepage: <http://ion.gluch.org.mx/files/Hacks/gallina>
- [16] Goollery homepage: <http://www.wirzm.ch/goollery>
- [17] Gmail-lite homepage: <http://gmail-lite.sourceforge.net>
- [18] Pocket Gmail: <http://gmail.wiggster.com>
- [19] Skins for Gmail: <http://persistent.info/archives/2004/10/05/gmail-skinning>
- [20] Gmail Notifier Extension: <https://addons.update.mozilla.org/extensions/moreinfo.php?id=173>
- [21] Gmail Drive: <http://www.viksoe.dk/code/gmail.htm>
- [22] Gmail Notifier: <http://www.google.com/downloads>

THE AUTHOR

Fabrizio Ciacchi <http://fabrizio.ciacchi.it> is an Italian Student of Computer Science at the University of Pisa. He also works as a consultant and writes articles on Linux.

Other things to do with Gmail

Some other projects also make innovative use of Gmail. The most famous is Gallina [15], a Blog system that uses the Gmail messages as blog posts; there is also Goollery [16] a photo gallery application. These two programs are written in PHP and use the Gmail-lite library [17].