

el-Studio.com

Tickling both halves of the



JungleDisk for Linux Server Backups without X11

Aug 14

tags: **backup**, **how to**, **linux**

share:

JungleDisk is another great entrant in the online backup business. Unlike Mozy, our favorite for Macs and PCs, JungleDisk also runs on Linux. And, since version 1.30a, it ships a command-line version.

Which means you can **give your Linux servers reliable online backup** — stored on Amazon's S3.

Here's more about **how JungleDisk works**.

Meanwhile, let's get down to business. Our server install has three parts:

- JungleDisk command line program, which makes S3 look like a WebDav server
- DavFS, which makes WebDav look like a regular Linux filesystem
- rsync, the standard Linux power copy program

You've got rsync, and the other two aren't much harder to install.



1. Install DavFS

1. Install davfs2. On Ubuntu Dapper or later this is easy once you **enable the Universe repository**.

```
sudo apt-get install davfs2
```

2. To keep JungleDisk happy, we'll add this line to the bottom of /etc/davfs2/davfs2.conf:

```
kernel_fs fuse
```

That tells davfs to use the newer Fuse backend instead of Coda, which avoids a JungleDisk bug.

3. For other linux distributions, see Chris's **instructions for installing davfs2 from source**

2. Install JungleDisk

1. **Download JungleDisk**, and untar it. Since version 1.30a, the Linux package comes with both GUI and command-line versions.

- `jungledisk` is the command-line version we'll use here

◦ `junglediskmonitor` is the GUI version. This version makes it easy to sign up for S3, and it writes the config file for you automatically — but you’ve got to have X11 installed to run it :(.

2. Sign up for an **Amazon S3 account**. `junglediskmonitor` can help you do this — if you’ve got X-windows installed. The `junglediskmonitor` also creates a settings file.

3. Or if you install JungleDisk on a Mac or Windows PC you can copy the config from that install. (On a Mac, the config can be found here: `~/Library/Preferences/jungledisk-settings.ini`)

4. Or you can make your own `~/.jungledisk/jungledisk-settings.ini`. You’ll need to put your Amazon S3 access and secret keys and your username in place of the

<removed> text below:

```
1 LoginUsername=
2 LoginPassword=
3 AccessKeyID=*<removed>*
4 SecretKey=*<removed>*
5 Bucket=default
6 CacheDirectory=/home/*<removed>*/.jungledisk/cache/
7 ListenPort=2667
8 CacheCheckInterval=120
9 AsyncOperations=1
10 Encrypt=1
11 ProxyServer=
12 EncryptionKey=PROTECTED:
13 DecryptionKeys=PROTECTED:
14 MaxCacheSize=1000
15 MapDrive=
16 UseSSL=0
17 RetryCount=3
18 FastCopy=1
19 WebAccess=0
20 LogDuration=30
21 License=
```

You can **download a sample config file** without the line numbers. A couple of lines will need editing to fit your circumstances.

Lines 3 and 4 are your Amazon S3 credentials. Copy and paste them from **Amazon’s website**. (Click the button labeled “Your Web Services Account”, then choose “AWS Access Identifiers” to see these.)

Line 6: The directory where JungleDisk will write to while its processing. This needs to exist and be writable by the user who will run `jungledisk`. Most folks should just put their Linux username in place of `*<removed>*`.

Lines 1, 2 and 21 can be left blank (as they are here) for your 30-day trial of JungleDisk. Once you buy the software, however, you’ll want to fill them in. Lines 1 and 2 are the email address and password you used to make your JungleDisk purchase.

Line 21 is a license code generated by `junglediskmonitor`. I’m not aware of a way to produce this in the command-line client. The Mac OS X version of JungleDisk keeps the license in the keychain, where **Keychain Access can show it**. If anybody knows where to find this in Windows, please let us know in the comments.

3. Actually Try it Out

Time to try it all out.

1. Start jungledisk

```
./jungledisk/jungledisk
```

2. Mount the jungledisk drive

```
sudo mkdir /mnt/s3
```

```
sudo mount -t davfs -o nolocks,noaskauth http://localhost:2667 /mnt/s3
```

3. Actually copy something over to S3

```
echo "Here's a file to try" > atestfile.txt
```

```
sudo cp atestfile.txt /mnt/s3
```

4. Rsync Backups

If you don't love rsync yet, you're gonna. Can you say "Differential Copy algorithm"? Rsync only copies the changed stuff. And that's what makes over-the-net backups workable.

Here's how to back up the whole enchilada (or at least all of your server's /home directories).

```
rsync -r --inplace --size-only --bwlimit=50 /home /mnt/s3
```

Notice that the rsync command adds a bandwidth limit (`--bwlimit=50`) to the inplace and size-only flags that webDAV and JungleDisk recommend. Without this, you'll find that jungledisk opens more files than the kernel will allow (at least in Ubuntu).

5. Automated Backups

To schedule this stuff automatically, we'll need to add a couple of lines to some of our config files. First we add a line to the end of fstab that lets our backup user mount our jungledrive (use sudo to carefully add this line).

/etc/fstab

```
http://localhost:2667 /mnt/s3 davfs user,nolocks,noaskauth 0 0
```

And here's the shell script we use to actually do the backup. Line 5 here shows that we've got one more choice here — when to start jungledisk. It could be done in line 5 here, but because it's a background process we want to leave running, we could start it on bootup by adding the jungledisk command to the end of /etc/rc.local.

/usr/local/bin/backup-jd

```
1 #!/bin/sh
2 ### Backs up office data to Jungledisk using rsync
3 LOGFILE=/var/log/backup-jd.log
4 ## Start in rc.local or here
5 #/usr/local/bin/jungledisk
6 mount /mnt/s3
7 echo "`date +%F %R`: Start backup-jd" >> $LOGFILE
8 rsync -r --inplace --size-only --bwlimit=50 /home /mnt/s3
9 echo "`date +%F %R`: Finish backup-jd" >> $LOGFILE
10 umount /mnt/s3
```

Updates & Revisions

8/19/07: Add line numbers and a better explanation of the JungleDisk config file, plus a

version for download.

1 Comments for JungleDisk for Linux Server Backups without X11

Thanks, great post. Now have automated backups running on my homeserver, plus a network drive I can access from the office too!

Comment by **Matt** Sep 22, 03:53 PM <#>

© 2000–2007 [el Studio Communications](#), LLC

Subscribe: [RSS](#) or [Atom](#) 