



Backing up Linux / trixbox Pro with Amazon S3 and JungleDisk

Document v1.1 - Last edited 7/28/2008 by Chris Sherwood, SureTeq, Inc.

That title is a mouthful, but this combination of killer apps will help any trixbox Pro implementer be able to safely and securely back up their configuration files (even though your configs are already backed up by virtue of the Hybrid-Hosted model), system recordings, and voicemails.

First though, what is Amazon S3? You can read about Amazon S3 on their Wikipedia page here: http://en.wikipedia.org/wiki/Amazon_S3, but basically, S3 is an online storage service provided by Amazon Web Services (AWS). Amazon S3 provides unlimited storage through an API, but in order to access it, you will need to either be a great programmer, or use a program such as JungleDisk. Once you're dialed in with Amazon S3, you can now back up unlimited data. The cost? \$0.10 per GB to transfer the data, and \$0.15 per GB per month to store the data. I back up about 6 GB of data from my personal computer meaning that it cost me \$0.60 to transfer the data, and \$0.90/month to keep it stored in the Amazon S3 data center. Not a bad price for piece of mind.

JungleDisk (www.jungledisk.com) is available for Windows, Mac, and Linux (30 day free trial, and \$20 to buy). I use it to back up my documents and pictures to Amazon S3, and I have the backup scheduled to run once a week...so basically, anything I have added or changed gets pushed every Sunday. Worst case scenario for me is losing a week's worth of data. You can also opt to back up monthly, weekly, daily, hourly... whatever you want.

Since trixbox Pro configurations are backed up automatically, but not recordings, I wanted to be able to utilize the power (and cost) of Amazon S3 to back up my system recordings and voicemails. This is a great idea for resellers to implement for their

customers as it provides a value-added service that can be attached to a maintenance contract at very little cost in either time or money to the reseller.

So, how do you set this up on trixbox Pro? Let's do it...

The first step is to sign up for an AWS account with Amazon. You can do that here: <http://aws.amazon.com/s3>. Once you are all signed up, you will receive an Access Key ID and a Secret Access Key. Make note of these (they are always available from your S3 account page) as we will be using them shortly.

Once your Amazon S3 account is in place, it is time to download and configure JungleDisk. Go to the JungleDisk download page (<http://www.jungledisk.com/desktop/download.aspx>) and download the USB Desktop Edition of JungleDisk. This version is made to be a stand-alone executable that can be put on a USB stick for portability, and is meant to be platform independent. The reason we want this version is for ease of setup...the USB Desktop Edition creates our configuration file (jungledisk-settings.xml) for us, and we can use that file with the Linux version to connect to our online storage. Once you have it downloaded, extract the ZIP file to your desired location and then (from Windows) run junglediskmonitor.exe. The JungleDisk monitor will open in your system tray...double click on it to open the interface and then click the 'configure' button. You should see this window:



Click on 'Add bucket...' in the left hand menu, and then click the 'Add Bucket' button on the 'Add a Bucket' screen. Now, you are presented with a screen prompting you for your Amazon S3 information:

Add a Bucket

Amazon Account Information

powered by
amazon
web services™

Your existing Amazon S3 Account keys have been filled in below.
You can change them to add a bucket from a different Amazon S3 account.

💡 Note that a Jungle Disk license is required for each Amazon S3 account.

Enter your Amazon S3 account keys below.
Use the link provided to retrieve your keys from the Amazon site.
Copy and paste the keys from the Amazon site to the boxes below.

[Retrieve your Amazon S3 keys from the Amazon site](#)

Amazon S3 Account Keys

Access Key ID: (required)

Secret Key: (required)

💡 Your keys are stored securely on your computer and never sent to anyone.

< Back Next > Cancel

Enter in your Access Key ID and Secret Key (available from your Amazon S3 logon page) and click 'Next.' You will then be prompted to use an existing bucket or create a new bucket. What the heck is a bucket? A bucket is Amazon's term for a collection of S3 files...using buckets, you can logically separate sets of files backed up to Amazon S3. If you have multiple customers, you would want to create a separate bucket for each customer. Select 'Create a new bucket' and click 'Next.' You are prompted for the name of the new bucket...enter it in and click 'Next' followed by 'Finish.' A good naming convention for buckets is the server ID of the trixbox, or perhaps tb[server ID].

Now that your bucket has been created, it shows up in the left-hand window of the main JungleDisk screen. There are various options and configuration settings (including encryption for anyone who is super paranoid), but for the purposes of this documentation, I'm going to take all defaults.

Now that JungleDisk knows how to connect to our S3 account, and it knows the name of our bucket, the configuration file (jungledisk-settings.xml) has been populated with all of the appropriate information. Now, it's time to get into the Linux CLI.

Connect to your trixbox using Putty or another SSH program such as SecureCRT.

Go to the `/usr/local` directory...this is where I'm putting the jungledisk application, but you can actually put it anywhere. I will also create a jungledisk directory:

```
cd /usr/local
mkdir jungledisk
```

With the USB Desktop Edition of JungleDisk, we already have all of the files we need. We now want to get the appropriate files over to the trixbox Pro. I prefer to use WinSCP to copy files from Windows to Linux. An explanation of how to use WinSCP is beyond the scope of this document, but the bottom line is that you need to get two files (jungledisk and jungledisk-settings.xml) moved over to your trixbox Pro and placed in the `/usr/local/jungledisk` directory.

The jungledisk executable file copies over in non-executable mode, so let's make sure we can run it on Linux:

```
chmod +x /usr/local/jungledisk/jungledisk
```

JungleDisk works by utilizing FUSE (File System in Userspace). I don't even really know what this is...but we need it for JungleDisk to function, so let's install it onto our system:

```
cd /usr/src
wget ftp://rpmfind.net/linux/dag/redhat/el4/en/i386/dag/RPMS/dkms-2.0.19-2.el4.rf.noarch.rpm
wget ftp://rpmfind.net/linux/dag/redhat/el4/en/i386/dag/RPMS/dkms-fuse-2.7.3-1.nodist.rf.noarch.rpm
rpm -ivh dkms*
modprobe fuse
```

Ok, so now that FUSE has been installed, let's connect to our online storage:

```
/usr/local/jungledisk/jungledisk /mnt/jungledisk -o config=/usr/local/jungledisk/jungledisk-settings.xml
```

This command mounts our online S3 bucket (the one configured in jungledisk-settings.xml) on `/mnt/jungledisk` using the settings in the `/usr/local/jungledisk/jungledisk-settings.xml`. You can now go to the `/mnt/jungledisk` directory and copy files to the Amazon S3 bucket you created.

We need to perform some other one-time miscellaneous tasks to make this solution work...these are very important, so don't skip them!

Edit the startup file (`rc.local`) so that FUSE and JungleDisk start automatically upon system boot:

```
nano -w /etc/rc.local
```

Type in these two lines at the bottom of the file:

```
/sbin/modprobe fuse
/usr/local/jungledisk/jungledisk /mnt/jungledisk -o config=/usr/local/jungledisk/jungledisk-settings.xml
```

CTRL+X to exit, 'Y' when prompted to save.

Create a symbolic link for the jungledisk executable so that it can be run from anywhere:

```
ln -s /usr/local/jungledisk/jungledisk /usr/bin/jungledisk
```

Ok, so now we can access the bucket...so what? Well...now, we need to create a backup script to back up our files to the bucket...after that, we will set the script to run automatically on a nightly basis.

The best way to accomplish this is going to be with rsync. Rsync is a built-in Linux application that syncs files between two locations. The reason rsync is preferable for this backup is because it has the ability to only back up new and changed files...very nice since we're paying 10 cents per GB worth of transfer (not much money...but hey...every little bit counts).

From the Linux CLI, do the following:

```
cd /usr/local/jungledisk
nano -w backup
```

Ok...we are now editing a new file called 'backup.' Copy and paste everything between the lines below into the backup script (exclude the lines):

```
#!/bin/bash

#####
# #
# Backup to Amazon S3 #
# #
# Created by Chris Sherwood - Fonality #
# #
# Version 1.0 - 07/25/2008 #
# #
#####

# This document assumes you are using /mnt/jungledisk as your S3
mount point

# *** Rsync your files ***
#
# Add or remove these as necessary for your organization

# Asterisk configuration files
rsync -atv /etc/asterisk/* /mnt/jungledisk/asteriskconfig

# System recordings (and all Asterisk sounds)
rsync -atv /var/lib/asterisk/* /mnt/jungledisk/asterisk

# System voicemails
```

```
rsync -atv /var/spool/asterisk/voicemail/* /mnt/jungledisk/  
voicemail
```

```
# ACD and FindMe recordings
```

```
rsync -atv /var/spool/asterisk/monitor/* /mnt/jungledisk/  
acdrecordings
```

```
# CDR CSV files
```

```
rsync -atv /var/log/asterisk/cdr-csv/* /mnt/jungledisk/cdr
```

```
# Just to know we're done
```

```
echo "All done!"
```

```
#
```

```
# OPTIONS:
```

```
#
```

```
# To mail the results of any of the rsyncs to you, append:
```

```
#
```

```
# | mail email@domain.com
```

```
#
```

```
# to the end of any of the rsync commands. Note, that is a pipe  
( | ) before mail.
```

```
#
```

CTRL+X to exit and then 'Y' when asked if you want to save.

Now, we have our script ready...you can edit the script to remove anything you don't want to back up, but by default, the script will back up Asterisk config files, system recordings, voicemails, ACD recordings, and CDR logs.

Now we need to make our backup file executable:

```
chmod +x /usr/local/jungledisk/backup
```

That's it! Assuming your JungleDisk is all connected and ready to rock 'n roll, you can now run backup (./backup) and you'll see your files sync to the S3 bucket. If you run it a second time, it should go much faster, and you'll see fewer files sync'd. This means everything is working properly.

Ok, so how do we automate this script? We do that with crontab.

```
nano -w /etc/crontab
```

Crontab looks confusing, but it's actually fairly simple...the format of a crontab command is as follows:

(minute 0-59) (hour 0-23) (day of month 1-31) (month 1-12) (day of week 0-7, where both 0 and 7 are Sunday) (user running command) (command)

So if we wanted our backup to run once a week on Sunday morning at 2:00am, it would look like this:

```
0 2 * * 0 root /usr/local/jungledisk/backup
```

If we wanted to run it every day at 12:05am, it would look like this:

```
05 0 * * * root /usr/local/jungledisk/backup
```

If we wanted to run it every 5 minutes, it would look like this:

```
0-59/5 * * * * root /usr/local/jungledisk/backup
```

So...basically, figure out how often you want it to run, and at what time, and add a line to the bottom of your crontab file accordingly.

Ok, so there's the whole solution...here's an additional tip for resellers: You can pretty much standardize this process for all of your customers. All you have to do is create a new bucket for each customer, and then edit one line in the jungledisk-settings.xml file to point to the new bucket...the rest of the instructions are the same.

A note from the author:

Hey everyone!

If you have found this document useful and plan on purchasing some VoIP equipment, I would greatly appreciate if you use the banner link for Voipsupply.com below. I have purchased equipment from Voipsupply.com many times, and I have always found their pricing to be great and their shipping fast, so I definitely recommend them. By using the link below, you won't be raising YOUR purchase price at all...you would simply be donating a small percentage of your purchase to me, and I would definitely appreciate it!

Thanks,

-Chris Sherwood

My Voipsupply.com banner link:

**LARGEST SELECTION OF
DIGIUM PRODUCTS...**

