

Setting up and running rsnapshot for backups

Bradley Tinder, Systems Integrator, SPK and Associates

Installation

First, you'll need to download rsnapshot from the official website <http://rsnapshot.org/>. There are packages for various Linux distributions as well as a generic Unix tarball, which will let you set it up manually. For this demonstration, I downloaded the latest version 1.3.1 in RPM format and used rpm to install it on my RedHat machine. After getting the program installed, you'll need to set it up with the configuration file.

Configuration

Rsnapshot will read /etc/rsnapshot.conf as its default configuration file, however you can create your own configuration file and pass the full path to it on the command line. I'll give an example of that a bit later in this document. If you're installing by package, you should have the default rsnapshot.conf file installed in /etc. Load that file into an editor and we'll go through some of the key things you need to setup for your rsnapshot to work properly.

Before we start, let's go over some quick prerequisites:

1. You'll need to first determine WHERE you plan to save your snapshots. At one of our clients, we use a local NAS device to store snapshots for quick retrieval of files.
2. You'll need to determine a retention period and frequency of snapshots. You can get as granular as hourly snapshots or roll out to longer periods like months and years.
3. Depending on how much data you have, snapshots can take a while to process, so sometimes it's best to run them during off hours.

After loading the configuration file into your favorite text editor, the first parameter you'll need to edit is the following:

IMPORTANT! Make sure to use tabs in between fields, not spaces!

snapshot_root

This needs to be set to the directory where you want all of your snapshots for this configuration file to be held.

Going further, the following parameter must be uncommented if you plan to do snapshots of remote filesystems:

cmd_ssh

Make sure it points to your ssh binary, which is typically in /usr/bin/ssh.

Most of the other parameters past this one are very self explanatory. If you have special file locations for various utilities, make sure you point to the correct location.

The next critical parameters are the **retain** parameters. These relate to how many snapshots you want to keep at various intervals. They must be unique and in ascending order starting with hourly.



www.spkaa.com
Ph: 888-310-4540

SPK and Associates
900 E Hamilton Ave, Ste. 100
Campbell, CA 95008

Here's an example that I setup for a client, where they keep 7 daily snapshots, 4 weekly and 3 monthly snapshots:

```
retain    daily    7
retain    weekly   4
retain    monthly  3
```

You can customize this to however fits your snapshot and backup methods, but I think it's a good start if you have the disk space.

The next parameter that you might want to change is the **exclude_file** parameter. This parameter just points to a file which contains a list of files/directories that you want to exclude from snapshotting. This gives you the power to exclude directories which you don't want to backup, and thus saving you disk space. Each line of the file contains a file pattern to exclude when rsnapshot starts. For example, if you didn't want to backup any files that ended with .tmp, you could put a line in that file that reads:

```
*.tmp
```

And all files that end in .tmp will be excluded.

Finally, the most important parameter of all, the **backup** parameter! This one controls what directories you want to snapshot. The syntax of the backup parameter is as follows:

backup<tab>source directory<tab>destination directory under snapshot_root<tab>options

For example, I want to backup my /etc directory to a directory named foemachine under the snapshot_root directory. My backup parameter would be:

```
backup    /etc/    foemachine/
```

If I wanted to backup the /etc directory on a remote server named boof:

```
backup    snapshotuser@boof:/etc/    boof/
```

This will create a directory named boof under the snapshot_root in the daily subdirectory. Under there will be the contents of the /etc directory from boof.

There are a lot more examples and ideas in the default rsnapshot.conf file, feel free to check those out, or if you have further questions, feel free to ask me!

Last, we need to check that the configuration file is good to use, and then schedule it to run via cron! To check that the configuration is valid, run this command as root:

rsnapshot configtest

If everything looks okay, you'll get a confirmation of that. Otherwise, it will list errors that you need to correct. Go ahead and fix them, and re-run the configtest.

If everything checks out, schedule to run your daily, weekly and monthly snapshots. One important thing to note is that you cannot run two instances of the rsnapshot command at the same time, or the second instance will not run. I have



www.spkaa.com
Ph: 888-310-4540

SPK and Associates
900 E Hamilton Ave, Ste. 100
Campbell, CA 95008

written a BASH script that executes these for me so no two commands are running at the same time. Look for this script and some more advanced examples of setting up a solid snapshot solution for your infrastructure in a future guide!