

Lazy Backup

If there's one thing good UI designers know, it's that the best UI is not to have one at all. Applications should just save, security should just work, and computers should just backup.

Apparently that last task is a harder than it appears, since I still haven't found decent backup software for Unix (OS X and GNU/Linux).

Here is how the software should work:

1. I install it.
2. I point it at some storage server (ideally Amazon EC2 and S3, but if that's too hard then a GNU/Linux server with a large drive).
3. I give it a maximum space limit (e.g. store no more than 200GB).
4. I give it a maximum up-bandwidth limit (e.g. use no more than 5K/s).
5. I tell it to run.

From then on, it should just work. In the background, it will upload my files to the server using only 5K/s of bandwidth. If I get disconnected from the Internet or reboot my computer, when I get back on it will pick up where it left off. If a file changes it will only send the diff and store that as well. When I run out of disk space it will delete the old diffs.

It will preserve all the Unix ACLs and permissions and weird Mac OS X resource forks and stuff so that if my drive ever dies I can make a full bootable restore from the backup.

Does this software exist?

The closest I've seen is [rdiff-backup](#), which is very nice but fails to automate some key steps.

If it doesn't exist, let me know if you're interested in writing it (a wrapper around `rdiff-backup` to do it shouldn't be too hard, I would think). I'd be willing to offer a bounty.

You should follow me on twitter [here](#).

November 29, 2006

[Read comments.](#)