

Dropbox to Amazon S3 Syncer

Over [on Github](#) I've published a Docker-based system for syncing a Dropbox folder with a S3 bucket. Drop a file in a designated Dropbox subdirectory, and it is automatically uploaded to S3. This is useful when you have a group of non-technical people who want to publish content to the web. For instance, a church might regularly generate podcasts (interesting fact: churches [are prolific podcasters](#)), and the weekly podcast need to be uploaded easily by a volunteer. With a dropbox-to-S3-syncer set up, the process is easy:

- Get each volunteer to install Dropbox
- Designate a folder to be the "public website" folder, and share this
- Anything placed in the designated folder automatically gets synced to S3

Large files may take a while to upload to S3. When `largefile.mp3` is added, a temporary marker file, `largefile.mp3.uploading.txt` appears, giving the uploader feedback that something is happening. When `uploading.txt` disappears (or is replaced with `uploadfailures.txt` in case of errors), the file is uploaded.

This Dropbox-to-S3 system seems to work pretty well when everyone already has Dropbox installed. A few caveats apply:

- The upload script is write-only, to avoid the risk of Dropbox shenanigans trashing public content. This means that removing accidentally added content still requires an S3 browser. In future I might implement a hack whereby the presence of a `filename.delete` file results in `filename` (and `filename.delete`) being deleted.
- The Dropbox folder must be cleaned out manually occasionally, to avoid eating up everyone's Dropbox quotas. This doesn't affect S3-published content, as deletes aren't propagated to S3.
- Of course, this requires a server to host the Dropbox-to-S3 Docker images, and all the sysadmin overhead that implies. The Dropbox-hosting Docker container takes about 318Mb of memory on my Linux host.

Patches welcome, and thanks to Jan B for the [docker-dropbox](#) image this is built on.