
GNU MediaGoblin Documentation

Release 0.3.3

Chris Webber, et al

July 30, 2013

CONTENTS

GNU MediaGoblin is a platform for sharing photos, video and other media in an environment that respects our freedom and independence.

This is a Free Software project. It is built by contributors for all to use and enjoy. If you're interested in contributing, see [the wiki](#) which has pages that talk about the ways someone can contribute.

PART 1: SITE ADMINISTRATOR'S GUIDE

This guide covers installing, configuring, deploying and running a GNU MediaGoblin website. It is written for site administrators.

1.1 Foreword

1.1.1 About the Site Administrator's Guide

This is the site administrator manual for GNU MediaGoblin. It covers how to set up and configure MediaGoblin and the kind of information that someone running MediaGoblin would need to know.

We have other documentation at:

- <http://mediagoblin.org/join/> for general “join us” information
- <http://wiki.mediagoblin.org/> for our contributor/developer-focused wiki

1.1.2 Improving the Site Administrator's Guide

There are a few ways—please pick whichever method is convenient for you!

1. Write up a bug report in the bug tracker
2. Tell someone on IRC #mediagoblin on Freenode.
3. Write an email to the devel mailing list.

Information about the bugtracker, IRC and the mailing list is all on the [join page](#).

Patches are the most helpful, but even feedback on what you think could be improved and how to improve it is also helpful.

1.2 About GNU MediaGoblin

Sections

- What is GNU MediaGoblin?
- Why Build GNU MediaGoblin?
- Who Contributes to the Project?
- How Can I Participate?
- How is GNU MediaGoblin licensed?
- Is MediaGoblin an official GNU project? What does that mean?

1.2.1 What is GNU MediaGoblin?

In 2008, a number of free software developers and activists gathered at the FSF to attempt to answer the question “What should software freedom look like on the participatory web?” Their answer, the [Franklin Street Statement](#) has lead to the development of [autonomo.us](#) community, and free software projects including [Identi.ca](#) and [Libre.fm](#).

Identi.ca and Libre.fm address the need for micro-blogging and music sharing services and software that respect users’ freedom and autonomy.

GNU MediaGoblin emerges from this milieu to create a platform for us to share photos, video and other media in an environment that respects our freedom and independence. In the future MediaGoblin will provide tools to facilitate collaboration on media projects.

1.2.2 Why Build GNU MediaGoblin?

The Internet is designed—and works best—as a complex and endlessly resilient network. When key services and media outlets are concentrated in centralized platforms, the network becomes less useful and increasingly fragile. As always, the proprietary nature of these systems, hinders users ability to develop, extend, and understand their software; however, in the case of network services it also means that users must forfeit control of their data to the service providers.

Therefore, we believe that network services must be federated to avoid centralization and that everyone ought to have control over their data. In support of this, we’ve decided to help build the tools to make these kinds of services possible. We hope you’ll join us, both as users and as contributors.

1.2.3 Who Contributes to the Project?

You do!

We are free software activists and folks who have worked on a variety of other projects including: Libre.fm, GNU Social, Status.net, Miro, Miro Community, and OpenHatch among others. We’re programmers, musicians, writers, and painters. We’re friendly and dedicated to software and network freedom.

1.2.4 How Can I Participate?

See [Get Involved](#) on the website. We eagerly look forward to seeing you!

1.2.5 How is GNU MediaGoblin licensed?

GNU MediaGoblin software is released under an AGPLv3 license.

See the `COPYING` file in the root of the source for details.

1.2.6 Is MediaGoblin an official GNU project? What does that mean?

MediaGoblin is an official GNU project! This status means that we meet the GNU Project's rigorous standards for free software. To find out more about what that means, check out the [GNU website](#).

Please feel free to contact us with further questions!

1.3 Deploying MediaGoblin

GNU MediaGoblin is fairly new and so at the time of writing, there aren't easy package-manager-friendly methods to install MediaGoblin. However, doing a basic install isn't too complex in and of itself.

There's an almost infinite way to deploy things... for now, we'll keep it simple with some assumptions and use a setup that combines mediagoblin + virtualenv + fastcgi + nginx on a .deb or .rpm based GNU/Linux distro.

Note: These tools are for site administrators wanting to deploy a fresh install. If instead you want to join in as a contributor, see our [Hacking HOWTO](#) instead.

There are also many ways to install servers... for the sake of simplicity, our instructions below describe installing with nginx. For more recipes, including Apache, see [our wiki](#).

1.3.1 Prepare System

Dependencies

MediaGoblin has the following core dependencies:

- Python 2.6 or 2.7
- [python-lxml](#)
- [git](#)
- SQLite/PostgreSQL
- [Python Imaging Library \(PIL\)](#)
- [virtualenv](#)

On a DEB-based system (e.g. Debian, gNewSense, Trisquel, Ubuntu, and derivatives) issue the following command:

```
sudo apt-get install git-core python python-dev python-lxml \  
python-imaging python-virtualenv
```

On a RPM-based system (e.g. Fedora, RedHat, and derivatives) issue the following command:

```
yum install python-paste-deploy python-paste-script \  
git-core python python-devel python-lxml python-imaging \  
python-virtualenv
```

Configure PostgreSQL

Note: MediaGoblin currently supports PostgreSQL and SQLite. The default is a local SQLite database. This will “just work” for small deployments.

For medium to large deployments we recommend PostgreSQL.

If you don't want/need postgres, skip this section.

These are the packages needed for Debian Wheezy (testing):

```
sudo apt-get install postgresql postgresql-client python-psycopg2
```

The installation process will create a new *system* user named `postgres`, it will have privileges sufficient to manage the database. We will create a new database user with restricted privileges and a new database owned by our restricted database user for our MediaGoblin instance.

In this example, the database user will be `mediagoblin` and the database name will be `mediagoblin` too.

To create our new user, run:

```
sudo -u postgres createuser mediagoblin
```

then answer **NO** to *all* the questions:

```
Shall the new role be a superuser? (y/n) n
Shall the new role be allowed to create databases? (y/n) n
Shall the new role be allowed to create more new roles? (y/n) n
```

then create the database all our MediaGoblin data should be stored in:

```
sudo -u postgres createdb -E UNICODE -O mediagoblin mediagoblin
```

where the first `mediagoblin` is the database owner and the second `mediagoblin` is the database name.

Caution: Where is the password?

These steps enable you to authenticate to the database in a password-less manner via local UNIX authentication provided you run the MediaGoblin application as a user with the same name as the user you created in PostgreSQL. More on this in *Drop Privileges for MediaGoblin*.

Drop Privileges for MediaGoblin

As MediaGoblin does not require special permissions or elevated access, you should run MediaGoblin under an existing non-root user or preferably create a dedicated user for the purpose of running MediaGoblin. Consult your distribution's documentation on how to create "system account" or dedicated service user. Ensure that it is not possible to log in to your system with as this user.

You should create a working directory for MediaGoblin. This document assumes your local git repository will be located at `/srv/mediagoblin.example.org/mediagoblin/` for this documentation. Substitute your preferred local deployment path as needed.

This document assumes that all operations are performed as this user. To drop privileges to this user, run the following command:

```
su - [mediagoblin]
```

Where, "[mediagoblin]" is the username of the system user that will run MediaGoblin.

1.3.2 Install MediaGoblin and Virtualenv

Note: MediaGoblin is still developing rapidly. As a result the following instructions recommend installing from the `master` branch of the git repository. Eventually production deployments will want to transition to running from more consistent releases.

Issue the following commands, to create and change the working directory. Modify these commands to reflect your own environment:

```
mkdir -p /srv/mediagoblin.example.org/  
cd /srv/mediagoblin.example.org/
```

Clone the MediaGoblin repository:

```
git clone git://gitorious.org/mediagoblin/mediagoblin.git
```

And set up the in-package virtualenv:

```
cd mediagoblin  
(virtualenv --system-site-packages . || virtualenv .) && ./bin/python setup.py develop
```

Note: If you have problems here, consider trying to install `virtualenv` with the `--distribute` or `--no-site-packages` options. If your system's default Python is in the 3.x series you may need to run `virtualenv` with the `--python=python2.7` or `--python=python2.6` options.

The above provides an in-package install of `virtualenv`. While this is counter to the conventional `virtualenv` configuration, it is more reliable and considerably easier to configure and illustrate. If you're familiar with Python packaging you may consider deploying with your preferred method.

Assuming you are going to deploy with FastCGI, you should also install `flup`:

```
./bin/easy_install flup
```

This concludes the initial configuration of the development environment. In the future, when you update your code-base, you should also run:

```
./bin/python setup.py develop --upgrade && ./bin/gmg dbupdate
```

Note: If you are running an active site, depending on your server configuration, you may need to stop it first or the `dbupdate` command may hang (and it's certainly a good idea to restart it after the update)

1.3.3 Deploy MediaGoblin Services

Configure MediaGoblin to use the PostgreSQL database

If you are using postgres, edit the `[mediagoblin]` section in your `mediagoblin_local.ini` and put in:

```
sql_engine = postgresql:///mediagoblin
```

if you are running the MediaGoblin application as the same 'user' as the database owner.

Update database data structures

Before you start using the database, you need to run:

```
./bin/gmg dbupdate
```

to populate the database with the MediaGoblin data structures.

Test the Server

At this point MediaGoblin should be properly installed. You can test the deployment with the following command:

```
./lazyserver.sh --server-name=broadcast
```

You should be able to connect to the machine on port 6543 in your browser to confirm that the service is operable.

Connect the Webserver to MediaGoblin with FastCGI

This section describes how to configure MediaGoblin to work via FastCGI. Our configuration example will use nginx, however, you may use any webserver of your choice as long as it supports the FastCGI protocol. If you do not already have a web server, consider nginx, as the configuration files may be more clear than the alternatives.

Create a configuration file at `/srv/mediagoblin.example.org/nginx.conf` and create a symbolic link into a directory that will be included in your nginx configuration (e.g. `/etc/nginx/sites-enabled` or `/etc/nginx/conf.d`) with one of the following commands (as the root user):

```
ln -s /srv/mediagoblin.example.org/nginx.conf /etc/nginx/conf.d/
ln -s /srv/mediagoblin.example.org/nginx.conf /etc/nginx/sites-enabled/
```

Modify these commands and locations depending on your preferences and the existing configuration of your nginx instance. The contents of this `nginx.conf` file should be modeled on the following:

```
server {
#####
# Stock useful config options, but ignore them :)
#####
include /etc/nginx/mime.types;

autoindex off;
default_type application/octet-stream;
sendfile on;

# Gzip
gzip on;
gzip_min_length 1024;
gzip_buffers 4 32k;
gzip_types text/plain text/html application/x-javascript text/javascript text/xml text/css;

#####
# Mounting MediaGoblin stuff
# This is the section you should read
#####

# Change this to update the upload size limit for your users
client_max_body_size 8m;

# prevent attacks (someone uploading a .txt file that the browser
# interprets as an HTML file, etc.)
add_header X-Content-Type-Options nosniff;

server_name mediagoblin.example.org www.mediagoblin.example.org;
access_log /var/log/nginx/mediagoblin.example.access.log;
error_log /var/log/nginx/mediagoblin.example.error.log;

# MediaGoblin's stock static files: CSS, JS, etc.
location /mgoblin_static/ {
```

```
    alias /srv/mediagoblin.example.org/mediagoblin/mediagoblin/static/;
}

# Instance specific media:
location /mgoblin_media/ {
    alias /srv/mediagoblin.example.org/mediagoblin/user_dev/media/public/;
}

# Theme static files (usually symlinked in)
location /theme_static/ {
    alias /srv/mediagoblin.example.org/mediagoblin/user_dev/theme_static/;
}

# Mounting MediaGoblin itself via FastCGI.
location / {
    fastcgi_pass 127.0.0.1:26543;
    include /etc/nginx/fastcgi_params;

    # our understanding vs nginx's handling of script_name vs
    # path_info don't match :)
    fastcgi_param PATH_INFO $fastcgi_script_name;
    fastcgi_param SCRIPT_NAME "";
}
}
```

Now, nginx instance is configured to serve the MediaGoblin application. Perform a quick test to ensure that this configuration works. Restart nginx so it picks up your changes, with a command that resembles one of the following (as the root user):

```
sudo /etc/init.d/nginx restart
sudo /etc/rc.d/nginx restart
```

Now start MediaGoblin. Use the following command sequence as an example:

```
cd /srv/mediagoblin.example.org/mediagoblin/
./lazyserver.sh --server-name=fcgi fcgi_host=127.0.0.1 fcgi_port=26543
```

Visit the site you've set up in your browser by visiting <http://mediagoblin.example.org>. You should see MediaGoblin!

Note: The configuration described above is sufficient for development and smaller deployments. However, for larger production deployments with larger processing requirements, see the “*Considerations for Production Deployments*” documentation.

1.4 Considerations for Production Deployments

This document contains a number of suggestions for deploying MediaGoblin in actual production environments. Consider “*Deploying MediaGoblin*” for a basic overview of how to deploy MediaGoblin.

1.4.1 Deploy with Paste

The instance configured with `./lazyserver.sh` is not ideal for a production MediaGoblin deployment. Ideally, you should be able to use an “init” or “control” script to launch and restart the MediaGoblin process.

Use the following command as the basis for such a script:

```
CELERY_ALWAYS_EAGER=true \  
/srv/mediagoblin.example.org/mediagoblin/bin/paster serve \  
/srv/mediagoblin.example.org/mediagoblin/paste.ini \  
--pid-file=/var/run/mediagoblin.pid \  
--server-name=fcgi fcgi_host=127.0.0.1 fcgi_port=26543
```

The above configuration places MediaGoblin in “always eager” mode with Celery, this means that submissions of content will be processed synchronously, and the user will advance to the next page only after processing is complete. If we take Celery out of “always eager mode,” the user will be able to immediately return to the MediaGoblin site while processing is ongoing. In these cases, use the following command as the basis for your script:

```
CELERY_ALWAYS_EAGER=false \  
/srv/mediagoblin.example.org/mediagoblin/bin/paster serve \  
/srv/mediagoblin.example.org/mediagoblin/paste.ini \  
--pid-file=/var/run/mediagoblin.pid \  
--server-name=fcgi fcgi_host=127.0.0.1 fcgi_port=26543
```

1.4.2 Separate Celery

While the `./lazyserver.sh` configuration provides an efficient way to start using a MediaGoblin instance, it is not suitable for production deployments for several reasons:

In nearly every scenario, work on the Celery queue will need to balance with the demands of other processes, and cannot proceed synchronously. This is a particularly relevant problem if you use MediaGoblin to host video content. Processing with Celery ought to be operationally separate from the MediaGoblin application itself, this simplifies management and support better workload distribution.

Basically, if you’re doing anything beyond a trivial workload, such as image hosting for a small set of users, or have limited media types such as “ASCII art” or icon sharing, you will need to run `celeryd` as a separate process.

Build an *init script* around the following command:

```
CELERY_CONFIG_MODULE=mediagoblin.init.celery.from_celery ./bin/celeryd
```

Modify your existing MediaGoblin and application init scripts, if necessary, to prevent them from starting their own `celeryd` processes.

1.4.3 Set up sentry to monitor exceptions

We have a plugin for [raven](#) integration, see the “*raven plugin*” documentation.

1.4.4 Use an Init Script

Look in your system’s `/etc/init.d/` or `/etc/rc.d/` directory for examples of how to build scripts that will start, stop, and restart MediaGoblin and Celery. These scripts will vary by distribution/operating system.

These are scripts provided by the MediaGoblin community:

Debian

- [GNU MediaGoblin init scripts](#) by Joar Wandborg

Arch Linux

- [MediaGoblin - ArchLinux rc.d scripts](#) by Jeremy Pope

- [Mediagoblin init script on Archlinux](#) by Chimo

1.5 Configuring MediaGoblin

So! You've got MediaGoblin up and running, but you need to adjust some configuration parameters. Well you've come to the right place!

1.5.1 MediaGoblin's config files

When configuring MediaGoblin, there are two files you might want to make local modified versions of, and one extra file that might be helpful to look at. Let's examine these.

mediagoblin.ini This is the config file for MediaGoblin, the application. If you want to tweak settings for MediaGoblin, you'll usually tweak them here.

paste.ini This is primarily a server configuration file, on the Python side (specifically, on the WSGI side, via [paste deploy / paste script](#)). It also sets up some middleware that you can mostly ignore, except to configure sessions... more on that later. If you are adding a different Python server other than fastcgi / plain HTTP, you might configure it here. You probably won't need to change this file very much.

There's one more file that you certainly won't change unless you're making coding contributions to mediagoblin, but which can be useful to read and reference:

mediagoblin/config_spec.ini This file is actually a specification for mediagoblin.ini itself, as a config file! It defines types and defaults. Sometimes it's a good place to look for documentation... or to find that hidden option that we didn't tell you about. :)

1.5.2 Making local copies

Let's assume you're doing the virtualenv setup described elsewhere in this manual, and you need to make local tweaks to the config files. How do you do that? Let's see.

To make changes to mediagoblin.ini

```
cp mediagoblin.ini mediagoblin_local.ini
```

To make changes to paste.ini

```
cp paste.ini paste_local.ini
```

From here you should be able to make direct adjustments to the files, and most of the commands described elsewhere in this manual will "notice" your local config files and use those instead of the non-local version.

Note: Note that all commands provide a way to pass in a specific config file also, usually by a `-cf` flag.

1.5.3 Common changes

Enabling email notifications

You'll almost certainly want to enable sending email. By default, MediaGoblin doesn't really do this... for the sake of developer convenience, it runs in "email debug mode".

To make MediaGoblin send email, you need a mailer daemon.

Change this in your `mediagoblin.ini` file:

```
email_debug_mode = false
```

You should also change the “from” email address by setting `email_sender_address`. For example:

```
email_sender_address = "foo@example.com"
```

If you have more custom SMTP settings, you also have the following options at your disposal, which are all optional, and do exactly what they sound like.

- `email_smtp_host`
- `email_smtp_port`
- `email_smtp_user`
- `email_smtp_pass`

All other configuration changes

To be perfectly honest, there are quite a few options and we haven’t had time to document them all.

So here’s a cop-out section saying that if you get into trouble, hop onto IRC and we’ll help you out. Details for the IRC channel is on the [join page](#) of the website.

1.5.4 Celery

FIXME: List Celery configuration here.

1.6 Media Types

In the future, there will be all sorts of media types you can enable, but in the meanwhile there are three additional media types: video, audio and ascii art.

First, you should probably read “*Configuring MediaGoblin*” to make sure you know how to modify the `mediagoblin` config file.

1.6.1 Enabling Media Types

Media types are enabled in your `mediagoblin` configuration file, typically it is created by copying `mediagoblin.ini` to `mediagoblin_local.ini` and then applying your changes to `mediagoblin_local.ini`. If you don’t already have a `mediagoblin_local.ini`, create one in the way described.

Most media types have additional dependencies that you will have to install. You will find descriptions on how to satisfy the requirements of each media type on this page.

To enable a media type, edit the `media_types` list in your `mediagoblin_local.ini`. For example, if your system supported image and video media types, then the list would look like this:

```
media_types = mediagoblin.media_types.image, mediagoblin.media_types.video
```

Note that after enabling new media types, you must run `dbupdate` like so:

```
./bin/gmg dbupdate
```

If you are running an active site, depending on your server configuration, you may need to stop it first (and it's certainly a good idea to restart it after the update).

1.6.2 How does MediaGoblin decide which media type to use for a file?

MediaGoblin has two methods for finding the right media type for an uploaded file. One is based on the file extension of the uploaded file; every media type maintains a list of supported file extensions. The second is based on a sniffing handler, where every media type may inspect the uploaded file and tell if it will accept it.

The file-extension-based approach is used before the sniffing-based approach, if the file-extension-based approach finds a match, the sniffing-based approach will be skipped as it uses far more processing power.

1.6.3 Video

To enable video, first install gstreamer and the python-gstreamer bindings (as well as whatever gstreamer extensions you want, good/bad/ugly). On Debianoid systems

```
sudo apt-get install python-gst0.10 \  
    gstreamer0.10-plugins-base \  
    gstreamer0.10-plugins-bad \  
    gstreamer0.10-plugins-good \  
    gstreamer0.10-plugins-ugly \  
    gstreamer0.10-ffmpeg
```

Add `mediagoblin.media_types.video` to the `media_types` list in your `mediagoblin_local.ini` and restart MediaGoblin.

Run

```
./bin/gmg dbupdate
```

Now you should be able to submit videos, and mediagoblin should transcode them.

Note: You almost certainly want to separate Celery from the normal paste process or your users will probably find that their connections time out as the video transcodes. To set that up, check out the “*Considerations for Production Deployments*” section of this manual.

1.6.4 Audio

To enable audio, install the gstreamer and python-gstreamer bindings (as well as whatever gstreamer plugins you want, good/bad/ugly), `scipy` and `numpy` are also needed for the audio spectrograms. To install these on Debianoid systems, run:

```
sudo apt-get install python-gst0.10 gstreamer0.10-plugins-{base,bad,good,ugly} \  
    gstreamer0.10-ffmpeg python-numpy python-scipy
```

The `scikits.audiolab` package you will install in the next step depends on the `libsndfile1-dev` package, so we should install it. On Debianoid systems, run

```
sudo apt-get install libsndfile1-dev
```

Note: `scikits.audiolab` will display a warning every time it's imported if you do not compile it with `alsa` support. `Alsa` support is not necessary for the GNU MediaGoblin application but if you do not wish the `alsa` warnings from `audiolab` you should also install `libasound2-dev` before installing `scikits.audiolab`.

Then install `scikits.audiolab` for the spectrograms:

```
./bin/pip install scikits.audiolab
```

Add `mediagoblin.media_types.audio` to the `media_types` list in your `mediagoblin_local.ini` and restart MediaGoblin.

Run

```
./bin/gmg dbupdate
```

You should now be able to upload and listen to audio files!

1.6.5 Ascii art

To enable `ascii art` support, first install the `chardet` library, which is necessary for creating thumbnails of `ascii art`

```
./bin/easy_install chardet
```

Next, modify (and possibly copy over from `mediagoblin.ini`) your `mediagoblin_local.ini`. In the `[mediagoblin]` section, add `mediagoblin.media_types.ascii` to the `media_types` list.

For example, if your system supported `image` and `ascii art` media types, then the list would look like this:

```
media_types = mediagoblin.media_types.image, mediagoblin.media_types.ascii
```

Run

```
./bin/gmg dbupdate
```

Now any `.txt` file you uploaded will be processed as `ascii art`!

1.6.6 STL / 3d model support

To enable the “STL” 3d model support plugin, first make sure you have a recentish `Blender` installed and available on your execution path. This feature has been tested with `Blender 2.63`. It may work on some earlier versions, but that is not guaranteed (and is surely not to work prior to `Blender 2.5X`).

Add `mediagoblin.media_types.stl` to the `media_types` list in your `mediagoblin_local.ini` and restart MediaGoblin.

Run

```
./bin/gmg dbupdate
```

You should now be able to upload `.obj` and `.stl` files and MediaGoblin will be able to present them to your wide audience of admirers!

1.7 How to Get Help with MediaGoblin

There are a couple of ways to get help with problems with MediaGoblin:

1. ask for help on IRC
2. ask for help on the devel mailing list

Details for both IRC and the mailing list are on the [join page](#) of the website.

1.8 Release Notes

This chapter has important information for releases in it. If you're upgrading from a previous release, please read it carefully, or at least skim over it.

1.8.1 0.3.3

Do this to upgrade

1. Make sure to run `bin/gmg dbupdate` after upgrading.
2. OpenStreetMap is now a plugin, so if you want to use it, add the following to your config file:

```
[plugins] [[mediagoblin.plugins.geolocation]]
```

If you have your own theme, you may need to make some adjustments to it as some theme related things may have changed in this release. If you run into problems, don't hesitate to [contact us](#) (IRC is often best).

New features

- New dropdown menu for accessing various features.
- Significantly improved URL generation. Now mediagoblin won't give up on making a slug if it looks like there will be a duplicate; it'll try extra hard to generate a meaningful one instead.

Similarly, linking to an id no longer can possibly conflict with linking to a slug; `/u/username/m/id:35/` is the kind of reference we now use to linking to entries with ids. However, old links with entries that linked to ids should work just fine with our migration. The only urls that might break in this release are ones using colons or equal signs.
- New template hooks for plugin authoring.
- As a demonstration of new template hooks for plugin authoring, openstreetmap support now moved to a plugin!
- Method to add media to collections switched from icon of paperclip to button with "add to collection" text.
- Bug where videos often failed to produce a proper thumbnail fixed!
- Copying around files in mediagoblin now much more efficient, doesn't waste gobs of memory.
- Video transcoding now optional for videos that meet certain criteria. By default, MediaGoblin will now transcode webm videos that are smaller in resolution than the mediagoblin defaults, and mediagoblin can also be configured to allow theora files to not be transcoded as well.
- Per-user license preference option; always want your uploads to be BY-SA and tired of changing that field? You can now set your license preference in your user settings.
- Video player now responsive; better for mobile!
- You can now delete your account from the user preferences page if you so wish.

Other changes

- Plugin writers: Internal restructuring led to `mediagoblin.db.sql*` be `mediagoblin.db.*` starting from 0.3.3
- Dependency list has been reduced not requiring the "webob" package anymore.

- And many small fixes/improvements, too numerous to list!

1.8.2 0.3.2

This will be the last release that is capable of converting from an earlier MongoDB-based MediaGoblin instance to the newer SQL-based system.

Do this to upgrade

1. Make sure to run `bin/gmg dbupdate` after upgrading.

New features

- **3d model support!**

You can now upload STL and OBJ files and display them in MediaGoblin. Requires a recent-ish Blender; for details see: *Deploying MediaGoblin*

- **trim_whitespace**

We bundle the optional plugin `trim_whitespace` which reduces the size of the delivered html output by reducing redundant whitespace.

See *Part 2: Core plugin documentation* for plugin documentation

- **A new API!**

It isn't well documented yet but we do have an API. There is an [android application in progress](#) which makes use of it, and there are some demo applications between [automgtic](#), an automatic media uploader for your desktop and [OMGMG](#), an example of a web application hooking up to the API.

This is a plugin, so you have to enable it in your mediagoblin config file by adding a section under `[plugins]` like:

```
[plugins]
[[mediagoblin.plugins.api]]
```

Note that the API works but is not nailed down... the way it is called may change in future releases.

- **OAuth login support**

For applications that use OAuth to connect to the API.

This is a plugin, so you have to enable it in your mediagoblin config file by adding a section under `[plugins]` like:

```
[plugins]
[[mediagoblin.plugins.oauth]]
```

- **Collections**

We now have user-curated collections support. These are arbitrary galleries that are customizable by users. You can add media to these by clicking on the paperclip icon when logged in and looking at a media entry.

- **OpenStreetMap licensing display improvements**

More accurate display of OSM licensing, and less disruptive: you click to “expand” the display of said licensing.

Geolocation is also now on by default.

- **Miscellaneous visual improvements**

We've made a number of small visual improvements including newer and nicer looking thumbnails and improved checkbox placement.

1.8.3 0.3.1

Do this to upgrade

1. Make sure to run `bin/gmg dbuptdate` after upgrading.
2. If you set up your server config with an older version of mediagoblin and the mediagoblin docs, it's possible you don't have the "theme static files" alias, so double check to make sure that section is there if you are having problems.

New features

- **theming support**

MediaGoblin now also includes theming support, which you can read about in the section *Theming MediaGoblin*.

- **flatpages**

MediaGoblin has a flatpages plugin allowing you to add pages that are aren't media-related like "About this site...", "Terms of service...", etc.

See *Part 2: Core plugin documentation* for plugin documentation

1.8.4 0.3.0

This release has one important change. You need to act when upgrading from a previous version!

This release changes the database system from MongoDB to SQL(alchemy). If you want to setup a fresh instance, just follow the instructions in the deployment chapter. If on the other hand you want to continue to use one instance, read on.

To convert your data from MongoDB to SQL(alchemy), you need to follow these steps:

1. Make sure your MongoDB is still running and has your data, it's needed for the conversion.
2. Configure the `sql_engine` URI in the config to represent your target database (see: *Deploying MediaGoblin*)
3. You need an empty database.
4. Then run the following command:

```
bin/gmg [-cf mediagoblin_config.ini] convert_mongo_to_sql
```

5. Start your server and investigate.
6. That's it.

1.9 Theming MediaGoblin

We try to provide a nice theme for MediaGoblin by default, but of course, you might want something different! Maybe you want something more light and colorful, or maybe you want something specifically tailored to your organization. Have no fear—MediaGoblin has theming support! This guide should walk you through installing and making themes.

1.9.1 Installing a theme

Installing the archive

Say you have a theme archive such as `goblincities.tar.gz` and you want to install this theme! Don't worry, it's fairly painless.

1. `cd ./user_dev/themes/`
2. Move the theme archive into this directory
3. `tar -xzvf <tar-archive>`
4. Open your configuration file (probably named `mediagoblin_local.ini`) and set the theme name:

```
[mediagoblin]
# ...
theme = goblincities
```

5. Link the assets so that they can be served by your web server:

```
$ ./bin/gmg theme assetlink
```

Note: If you ever change the current theme in your config file, you should re-run the above command!

(In the near future this should be even easier ;))

Set up your webserver to serve theme assets

If you followed the nginx setup example in *Connect the Webserver to MediaGoblin with FastCGI* you should already have theme asset setup. However, if you set up your server config with an older version of mediagoblin and the mediagoblin docs, it's possible you don't have the "theme static files" alias, so double check to make sure that section is there if you are having problems.

If you are simply using this for local development and serving the whole thing via paste/lazyserver, assuming you don't have a `paste_local.ini`, the asset serving should be done for you.

Configuring where things go

By default, MediaGoblin's install directory for themes is `./user_dev/themes/` (relative to the MediaGoblin checkout or base config file.) However, you can change this location easily with the `theme_install_dir` setting in the `[mediagoblin]` section.

For example:

```
[mediagoblin]
# ... other parameters go here ...
theme_install_dir = /path/to/themes/
```

Other variables you may consider setting:

theme_web_path When theme-specific assets are specified, this is where MediaGoblin will set the urls. By default this is `"/theme_static/"` so in the case that your theme was trying to access its file `"images/shiny_button.png"` MediaGoblin would link to `/theme_static/images/shiny_button.png`.

theme_linked_assets_dir Your web server needs to serve the theme files out of some directory, and MediaGoblin will symlink the current theme’s assets here. See the “Link the assets” step in *Installing the archive*.

1.9.2 Making a theme

Okay, so a theme layout is pretty simple. Let’s assume we’re making a theme for an instance about hedgehogs! We’ll call this the “hedgehogified” theme.

Change to where your `theme_install_dir` is set to (by default, `./user_dev/themes/` ... make those directories or otherwise adjust if necessary):

```
hedgehogified/
|- theme.cfg           # configuration file for this theme
|- templates/         # override templates
|  '- mediagoblin/
|     |- base.html    # overriding mediagoblin/base.html
|     '- root.html    # overriding mediagoblin/root.html
'- assets/
|  '- images/
|     | |- im_a_hedgehog.png # hedgehog-containing image used by theme
|     | '- custom_logo.png # your theme’s custom logo
|     '- css/
|         '- hedgehog.css    # your site’s hedgehog-specific css
|- README.txt         # Optionally, a readme file (not required)
|- AGPLv3.txt         # AGPL license file for your theme. (good practice)
'- CC0_1.0.txt        # CC0 1.0 legalcode for the assets [if appropriate!]
```

The top level directory of your theme should be the symbolic name for your theme. This is the name that users will use to refer to activate your theme.

Note: It’s important to note that templates based on MediaGoblin’s code should be released as AGPLv3 (or later), like MediaGoblin’s code itself. However, all the rest of your assets are up to you. In this case, we are waiving our copyright for images and CSS into the public domain via CC0 (as MediaGoblin does) but do what’s appropriate to you.

1.9.3 The config file

The config file is not presently strictly required, though it is nice to have. Only a few things need to go in here:

```
[theme]
name = Hedgehog-ification
description = For hedgehog lovers ONLY
licensing = AGPLv3 or later templates; assets (images/css) waived under CC0 1.0
```

The name and description fields here are to give users an idea of what your theme is about. For the moment, we don’t have any listing directories or admin interface, so this probably isn’t useful, but feel free to set it in anticipation of a more glorious future.

Licensing field is likewise a textual description of the stuff here; it’s recommended that you preserve the “AGPLv3 or later templates” and specify whatever is appropriate to your assets.

Templates

Your template directory is where you can put any override and custom templates for MediaGoblin.

These follow the general MediaGoblin theming layout, which means that the MediaGoblin core templates are all kept under the `./mediagoblin/` prefix directory.

You can copy files right out of MediaGoblin core and modify them in this matter if you wish.

To fit with best licensing form, you should either preserve the MediaGoblin copyright header borrowing from a MediaGoblin template, or you may include one like the following:

```
{#
# [YOUR THEME], a MediaGoblin theme
# Copyright (C) [YEAR] [YOUR NAME]
#
# This program is free software: you can redistribute it and/or modify
# it under the terms of the GNU Affero General Public License as published by
# the Free Software Foundation, either version 3 of the License, or
# (at your option) any later version.
#
# This program is distributed in the hope that it will be useful,
# but WITHOUT ANY WARRANTY; without even the implied warranty of
# MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
# GNU Affero General Public License for more details.
#
# You should have received a copy of the GNU Affero General Public License
# along with this program. If not, see <http://www.gnu.org/licenses/>.
#}
```

Assets

Put any files, such as images, CSS, etc, that are specific to your theme in here.

You can reference these in your templates like so:

```

```

This will tell MediaGoblin to reference this image from the current theme.

Licensing file(s)

You should include `AGPLv3.txt` with your theme as this is required for the assets. You can copy this from `mediagoblin/licenses/`.

In the above example, we also use `CC0` to waive our copyrights to images and css, so we also included `CC0_1.0.txt`

A README.txt file

A `README` file is not strictly required, but probably a good idea. You can put whatever in here, but restating the license choice clearly is probably a good idea.

Simple theming by adding CSS

Many themes won't require anything other than the ability to override some of MediaGoblin's core css. Thankfully, doing so is easy if you combine the above steps!

In your theme, do the following (make sure you make the necessary directories and `cd` to your theme's directory first):

```
$ cp /path/to/mediagoblin/mediagoblin/templates/mediagoblin/extra_head.html templates/mediagoblin/
```

Great, now open that file and add something like this at the end:

```
<link rel="stylesheet" type="text/css"
      href="{ request.staticdirect('/css/theme.css', 'theme') }"/>
```

You can name the css file whatever you like. Now make the directory for `assets/css/` and add the file `assets/css/theme.css`.

You can now put custom CSS files in here and any CSS you add will override default MediaGoblin CSS.

Packaging it up!

Packaging a theme is really easy. It's just a matter of making an archive!

Change to the installed themes directory and run the following:

```
tar -cvfz yourtheme.tar.gz yourtheme
```

Where “yourtheme” is replaced with your theme name.

That's it!

1.10 Plugins

GNU MediaGoblin supports plugins that allow you to augment MediaGoblin's behavior.

This chapter covers discovering, installing, configuring and removing plugins.

1.10.1 Discovering plugins

MediaGoblin comes with core plugins. Core plugins are located in the `mediagoblin.plugins` module of the MediaGoblin code. Because they come with MediaGoblin, you don't have to install them, but you do have to add them to your config file if you're interested in using them.

You can also write your own plugins and additionally find plugins elsewhere on the Internet. Once you find a plugin you like, you need to first install it, then add it to your configuration.

1.10.2 Installing plugins

Core plugins

MediaGoblin core plugins don't need to be installed because they come with MediaGoblin. Further, when you upgrade MediaGoblin, you will also get updates to the core plugins.

Other plugins

If the plugin is available on the [Python Package Index](#), then you can install the plugin with pip:

```
pip install <plugin-name>
```

For example, if we wanted to install the plugin named “mediagoblin-licenses” (which allows you to customize the licenses you offer for your media), we would do:

```
pip install mediagoblin-licenses
```

Note: If you’re using a virtual environment, make sure to activate the virtual environment before installing with pip. Otherwise the plugin may get installed in a different environment than the one MediaGoblin is installed in. Also make sure, you use e.g. pip-2.7 if your default python (and thus pip) is python 3 (e.g. in Ubuntu).

Once you’ve installed the plugin software, you need to tell MediaGoblin that this is a plugin you want MediaGoblin to use. To do that, you edit the `mediagoblin.ini` file and add the plugin as a subsection of the plugin section.

For example, say the “mediagoblin-licenses” plugin has the Python package path `mediagoblin_licenses`, then you would add `mediagoblin_licenses` to the `plugins` section as a subsection:

```
[plugins]

[[mediagoblin_licenses]]
license_01=abbrev1, name1, http://url1
license_02=abbrev2, name1, http://url2
```

1.10.3 Configuring plugins

Configuration for a plugin goes in the subsection for that plugin. Core plugins are documented in the administration guide. Other plugins should come with documentation that tells you how to configure them.

Example 1: Core MediaGoblin plugin

If you wanted to use the core MediaGoblin `flatpages` plugin, the module for that is `mediagoblin.plugins.flatpagesfile` and you would add that to your `.ini` file like this:

```
[plugins]

[[mediagoblin.plugins.flatpagesfile]]
# configuration for flatpagesfile plugin here!
about-view = '/about', about.html
terms-view = '/terms', terms.html
```

(Want to know more about the `flatpagesfile` plugin? See [flatpagesfile plugin](#))

Example 2: Plugin that is not a core MediaGoblin plugin

If you installed a hypothetical restrictive plugin which is in the module `restrictfive`, your `.ini` file might look like this (with comments making the bits clearer):

```
[plugins]

[[restrictfive]]
# configuration for restrictfive here!
```

Check the plugin’s documentation for what configuration options are available.

1.10.4 Removing plugins

To remove a plugin, use `pip uninstall`. For example:

```
pip uninstall mediagoblin-licenses
```

Note: If you're using a virtual environment, make sure to activate the virtual environment before uninstalling with pip. Otherwise the plugin may get installed in a different environment.

1.10.5 Upgrading plugins

Core plugins

Core plugins get upgraded automatically when you upgrade MediaGoblin because they come with MediaGoblin.

Other plugins

For plugins that you install with pip, you can upgrade them with pip:

```
pip install -U <plugin-name>
```

The `-U` tells pip to upgrade the package.

1.10.6 Troubleshooting plugins

Sometimes plugins just don't work right. When you're having problems with plugins, think about the following:

1. Check the log files.

Some plugins will log errors to the log files and you can use that to diagnose the problem.

2. Try running MediaGoblin without that plugin.

It's easy to disable a plugin from MediaGoblin. Add a `-` to the name in your config file.

For example, change:

```
[[mediagoblin.plugins.flatpagesfile]]
```

to:

```
[[-mediagoblin.plugins.flatpagesfile]]
```

That'll prevent the `mediagoblin.plugins.flatpagesfile` plugin from loading.

3. If it's a core plugin that comes with MediaGoblin, ask us for help!

If it's a plugin you got from somewhere else, ask them for help!

PART 2: CORE PLUGIN DOCUMENTATION

2.1 flatpagesfile plugin

This is the flatpages file plugin. It allows you to add pages to your MediaGoblin instance which are not generated from user content. For example, this is useful for these pages:

- About this site
- Terms of service
- Privacy policy
- How to get an account here
- ...

2.1.1 How to configure

Add the following to your MediaGoblin .ini file in the [plugins] section:

```
[[mediagoblin.plugins.flatpagesfile]]
```

This tells MediaGoblin to load the flatpagesfile plugin. This is the subsection that you'll do all flatpagesfile plugin configuration in.

2.1.2 How to add pages

To add a new page to your site, you need to do two things:

1. add a route to the MediaGoblin .ini file in the flatpagesfile subsection
2. write a template that will get served when that route is requested

Routes

First, let's talk about the route.

A route is a key/value in your configuration file.

The key for the route is the route name. You can use this with `url()` in templates to have MediaGoblin automatically build the urls for you. It's very handy.

It should be “unique” and it should be alphanumeric characters and hyphens. I wouldn't put spaces in there.

Examples: `flatpages-about`, `about-view`, `contact-view`, ...

The value has two parts separated by commas:

1. **route path:** This is the url that this route matches.

Examples: `/about`, `/contact`, `/pages/about`, ...

You can do anything with this that you can do with the `route_path` parameter of `routes.Route`. For more details, see [the routes documentation](#).

Example: `/siteadmin/{adminname:\w+}`

Note: If you're doing something fancy, enclose the route in single quotes.

For example: `'/siteadmin/{adminname:\w+}'`

2. **template:** The template to use for this url. The template is in the `flatpagesfile` template directory, so you just need to specify the file name.

Like with other templates, if it's an HTML file, it's good to use the `.html` extensions.

Examples: `index.html`, `about.html`, `contact.html`, ...

Here's an example configuration that adds two flat pages: one for an “About this site” page and one for a “Terms of service” page:

```
[[mediagoblin.plugins.flatpagesfile]]
about-view = '/about', about.html
terms-view = '/terms', terms.html
```

Note: The order in which you define the routes in the config file is the order in which they're checked for incoming requests.

Templates

To add pages, you must edit template files on the file system in your `local_templates` directory.

The directory structure looks kind of like this:

```
local_templates
|- flatpagesfile
  |- flatpage1.html
  |- flatpage2.html
  |- ...
```

The `.html` file contains the content of your page. It's just a template like all the other templates you have.

Here's an example that extends the `flatpagesfile/base.html` template:

```
{% extends "flatpagesfile/base.html" %}
{% block mediagoblin_content %}
<h1>About this site</h1>
<p>
  This site is a MediaGoblin instance set up to host media for
```

```
    me, my family and my friends.  
</p>  
{% endblock %}
```

Note: If you have a bunch of flatpages that kind of look like one another, take advantage of Jinja2 template extending and create a base template that the others extend.

2.1.3 Recipes

Url variables

You can handle urls like `/about/{name}` and access the name that's passed in in the template.

Sample route:

```
about-page = '/about/{name}', about.html
```

Sample template:

```
{% extends "flatpagesfile/base.html" %}  
{% block mediagoblin_content %}  
  
<h1>About page for {{ request.matchdict['name'] }}</h1>  
  
{% endblock %}
```

See the [the routes documentation](#) for syntax details for the route. Values will end up in the `request.matchdict` dict.

2.2 sampleplugin

This is a sample plugin. It does nothing interesting other than show one way to structure a MediaGoblin plugin.

The code for this plugin is in `mediagoblin/plugins/sampleplugin/`.

2.3 OAuth plugin

Warning: In its current state. This plugin has received no security audit. Development has been entirely focused on Making It Work(TM). Use this plugin with caution. Additionally, this and the API may break... consider it pre-alpha. There's also a known issue that the OAuth client doesn't do refresh tokens so this might result in issues for users.

The OAuth plugin enables third party web applications to authenticate as one or more GNU MediaGoblin users in a safe way in order retrieve, create and update content stored on the GNU MediaGoblin instance.

The OAuth plugin is based on the [oauth v2.25 draft](#) and is pointing by using the `oauthlib.oauth2.draft25.WebApplicationClient` from [oauthlib](#) to a mediagoblin instance and building the OAuth 2 provider logic around the client.

There are surely some aspects of the OAuth v2.25 draft that haven't made it into this plugin due to the technique used to develop it.

2.3.1 Set up the OAuth plugin

1. Add the following to your MediaGoblin .ini file in the [plugins] section:

```
[[mediagoblin.plugins.oauth]]
```

2. Run:

```
gmg dbupdate
```

in order to create and apply migrations to any database tables that the plugin requires.

Note: This only enables the OAuth plugin. To be able to let clients fetch data from the MediaGoblin instance you should also enable the API plugin or some other plugin that supports authenticating with OAuth credentials.

2.3.2 Authenticate against GNU MediaGoblin

Note: As mentioned in [capabilities](#) GNU MediaGoblin currently only supports the [Authorization Code Grant](#) procedure for obtaining an OAuth access token.

Authorization Code Grant

Note: As mentioned in [incapabilities](#) GNU MediaGoblin currently does not support [client registration](#)

The [authorization code grant](#) works in the following way:

Definitions

Authorization server The GNU MediaGoblin instance

Resource server Also the GNU MediaGoblin instance ;)

Client The web application intended to use the data

Redirect uri An URI pointing to a page controlled by the *client*

Resource owner The GNU MediaGoblin user who's resources the client requests access to

User agent Commonly the GNU MediaGoblin user's web browser

Authorization code An intermediate token that is exchanged for an *access token*

Access token A secret token that the *client* uses to authenticate itself against the *resource server* as a specific *resource owner*.

Brief description of the procedure

1. The *client* requests an *authorization code* from the *authorization server* by redirecting the *user agent* to the [Authorization Endpoint](#). Which parameters should be included in the redirect are covered later in this document.
2. The *authorization server* authenticates the *resource owner* and redirects the *user agent* back to the *redirect uri* (covered later in this document).
3. The *client* receives the request from the *user agent*, attached is the *authorization code*.

4. The *client* requests an *access token* from the *authorization server*
5. ?????
6. Profit!

Detailed description of the procedure

TBD, in the meantime here is a proof-of-concept GNU MediaGoblin client:

<https://github.com/jwandborg/omgmg/>

and here are some detailed descriptions from other OAuth 2 providers:

- <https://developers.google.com/accounts/docs/OAuth2WebServer>
- <https://developers.facebook.com/docs/authentication/server-side/>

and if you're unsure about anything, there's the [OAuth v2.25 draft](#), the [OAuth plugin source code](#) and the [#mediagoblin IRC channel](#).

2.3.3 Capabilities

- [Authorization endpoint](#) - Located at `/oauth/authorize`
- [Token endpoint](#) - Located at `/oauth/access_token`
- [Authorization Code Grant](#)
- [Client Registration](#)

2.3.4 Incapabilities

- Only bearer tokens are issued.
- [Implicit Grant](#)
- [Force TLS for token endpoint](#) - This one is up the the siteadmin
- [Authorization scope](#) and *state*
- ...

2.4 Trim whitespace plugin

Mediagoblin templates are written with 80 char limit for better readability. However that means that the html output is very verbose containing LOTS of whitespace. This plugin inserts a Middleware that filters out whitespace from the returned HTML in the `Response()` objects.

Simply enable this plugin by putting it somewhere where python can reach it and put it's path into the `[plugins]` section of your `mediagoblin.ini` or `mediagoblin_local.ini` like for example this:

```
[plugins] [[mediagoblin.plugins.trim_whitespace]]
```

There is no further configuration required. If this plugin is enabled, all text/html documents should not have lots of whitespace in between elements, although it does a very naive filtering right now (just keep the first whitespace and delete all subsequent ones).

Nonetheless, it is a useful plugin that might serve as inspiration for other plugin writers.

It was originally conceived by Sebastian Spaeth. It is licensed under the GNU AGPL v3 (or any later version) license.

2.5 raven plugin

2.5.1 Set up the raven plugin

1. Add the following to your MediaGoblin .ini file in the [plugins] section:

```
[[mediagoblin.plugins.raven]]
sentry_dsn = <YOUR SENTRY DSN>
# Logging is very high-volume, set to 0 if you want to turn off logging
setup_logging = 1
```

PART 3: PLUGIN WRITER'S GUIDE

This guide covers writing new GNU MediaGoblin plugins.

3.1 Foreword

3.1.1 About the Plugin Writer's Guide

This guide covers writing plugins for GNU MediaGoblin. It's very much a work in progress partially because we just started writing it and partially because the plugin API is currently in flux.

3.1.2 Improving the Plugin Writer's Guide

There are a few ways—please pick whichever method is convenient for you!

1. Write up a bug report in the bug tracker
2. Tell someone on IRC #mediagoblin on Freenode.
3. Write an email to the devel mailing list.

Information about the bugtracker, IRC and the mailing list is all on the [join page](#).

Patches are the most helpful, but even feedback on what you think could be improved and how to improve it is also helpful.

3.2 Quick Start

This is a quick start. It's not comprehensive, but it walks through writing a basic plugin called “sampleplugin” which logs “I've been started!” when `setup_plugin()` has been called.

3.2.1 Step 1: Files and directories

GNU MediaGoblin plugins are Python projects at heart. As such, you should use a standard Python project directory tree:

```
sampleplugin/  
|- README  
|- LICENSE  
|- setup.py  
|- sampleplugin/  
    |- __init__.py
```

The outer `sampleplugin` directory holds all the project files.

The `README` should cover what your plugin does, how to install it, how to configure it, and all the sorts of things a `README` should cover.

The `LICENSE` should have the license under which you're distributing your plugin.

The inner `sampleplugin` directory is the Python package that holds your plugin's code.

The `__init__.py` denotes that this is a Python package. It also holds the plugin code and the `hooks` dict that specifies which hooks the `sampleplugin` uses.

3.2.2 Step 2: README

Here's a rough `README`. Generally, you want more information because this is the file that most people open when they want to learn more about your project.

```
README  
=====
```

```
This is a sample plugin. It logs a line when ``setup__plugin()`` is  
run.
```

3.2.3 Step 3: LICENSE

GNU MediaGoblin plugins must be licensed under the AGPLv3 or later. So the `LICENSE` file should be the AGPLv3 text which you can find at <http://www.gnu.org/licenses/agpl-3.0.html>

3.2.4 Step 4: setup.py

This file is used for packaging and distributing your plugin.

We'll use a basic one:

```
from setuptools import setup, find_packages  
  
setup(  
    name='sampleplugin',  
    version='1.0',  
    packages=find_packages(),  
    include_package_data=True,  
    install_requires=[],  
    license='AGPLv3',  
)
```

See <http://docs.python.org/distutils/index.html#distutils-index> for more details.

3.2.5 Step 5: the code

The code for `__init__.py` looks like this:

```

1  import logging
2  from mediagoblin.tools.pluginapi import Plugin, get_config
3
4
5  # This creates a logger that you can use to log information to
6  # the console or a log file.
7  _log = logging.getLogger(__name__)
8
9
10 # This is the function that gets called when the setup
11 # hook fires.
12 def setup_plugin():
13     _log.info("I've been started!")
14     config = get_config('sampleplugin')
15     if config:
16         _log.info('%r' % config)
17     else:
18         _log.info('There is no configuration set.')
19
20
21 # This is a dict that specifies which hooks this plugin uses.
22 # This one only uses one hook: setup.
23 hooks = {
24     'setup': setup_plugin
25 }
```

Line 12 defines the `setup_plugin` function.

Line 23 defines `hooks`. When MediaGoblin loads this file, it sees `hooks` and registers all the callables with their respective hooks.

3.2.6 Step 6: Installation and configuration

To install the plugin for development, you need to make sure it's available to the Python interpreter that's running MediaGoblin.

There are a couple of ways to do this, but we're going to pick the easy one.

Use `python` from your MediaGoblin virtual environment and do:

```
python setup.py develop
```

Any changes you make to your plugin will be available in your MediaGoblin virtual environment.

Then adjust your `mediagoblin.ini` file to load the plugin:

```
[plugins]

[[sampleplugin]]
```

3.2.7 Step 7: That's it!

When you launch MediaGoblin, it'll load the plugin and you'll see evidence of that in the log file.

That's it for the quick start!

3.2.8 Where to go from here

See the documentation on the plugin API for code samples and other things you can use when building your plugin.

See [Hitchhiker's Guide to Packaging](#) for more information on packaging your plugin.

3.3 Database

3.3.1 Accessing Existing Data

If your plugin wants to access existing data, this is quite straight forward. Just import the appropriate models and use the full power of SQLAlchemy. Take a look at the (upcoming) database section in the Developer's Chapter.

3.3.2 Creating new Tables

If your plugin needs some new space to store data, you should create a new table. Please do not modify core tables. Not doing so might seem inefficient and possibly is. It will help keep things sane and easier to upgrade versions later.

So if you create a new plugin and need new tables, create a file named `models.py` in your plugin directory. You might take a look at the core's `db.models` for some ideas. Here's a simple one:

```
from mediagoblin.db.base import Base
from sqlalchemy import Column, Integer, Unicode, ForeignKey

class MediaSecurity(Base):
    __tablename__ = "yourplugin__media_security"

    # The primary key *and* reference to the main media_entry
    media_entry = Column(Integer, ForeignKey('core__media_entries.id'),
        primary_key=True)
    get_media_entry = relationship("MediaEntry",
        backref=backref("security_rating", cascade="all, delete-orphan"))

    rating = Column(Unicode)

MODELS = [MediaSecurity]
```

That's it.

Some notes:

- Make sure all your `__tablename__` start with your plugin's name so the tables of various plugins can't conflict in the database. (Conflicts in python naming are much easier to fix later).
- Try to get your database design as good as possible in the first attempt. Changing the database design later, when people already have data using the old design, is possible (see next chapter), but it's not easy.

3.3.3 Changing the Database Schema Later

If your plugin is in use and instances use it to store some data, changing the database design is a tricky thing.

1. Make up your mind how the new schema should look like.

2. Change `models.py` to contain the new schema. Keep a copy of the old version around for your personal reference later.
3. Now make up your mind (possibly using your old and new `models.py`) what steps in SQL are needed to convert the old schema to the new one. This is called a “migration”.
4. Create a file `migrations.py` that will contain all your migrations and add your new migration.

Take a look at the core’s `db/migrations.py` for some good examples on what you might be able to do. Here’s a simple one to add one column:

```
from mediagoblin.db.migration_tools import RegisterMigration, inspect_table
from sqlalchemy import MetaData, Column, Integer

MIGRATIONS = {}

@RegisterMigration(1, MIGRATIONS)
def add_license_preference(db):
    metadata = MetaData(bind=db.bind)

    security_table = inspect_table(metadata, 'yourplugin__media_security')

    col = Column('security_level', Integer)
    col.create(security_table)
    db.commit()
```

3.4 Plugin API

3.4.1 pluginapi Module

This module implements the plugin api bits.

Two things about things in this module:

1. they should be excessively well documented because we should pull from this file for the docs
2. they should be well tested

How do plugins work?

Plugins are structured like any Python project. You create a Python package. In that package, you define a high-level `__init__.py` module that has a `hooks` dict that maps hooks to callables that implement those hooks.

Additionally, you want a `LICENSE` file that specifies the license and a `setup.py` that specifies the metadata for packaging your plugin. A rough file structure could look like this:

```
myplugin/
|- setup.py           # plugin project packaging metadata
|- README            # holds plugin project information
|- LICENSE           # holds license information
|- myplugin/         # plugin package directory
    |- __init__.py   # has hooks dict and code
```

Lifecycle

1. All the modules listed as subsections of the `plugins` section in the config file are imported. MediaGoblin registers any hooks in the `hooks` dict of those modules.
2. After all plugin modules are imported, the `setup` hook is called allowing plugins to do any set up they need to do.

`mediagoblin.tools.pluginapi.get_config(key)`
Retrieves the configuration for a specified plugin by key

Example:

```
>>> get_config('mediagoblin.plugins.sampleplugin')
{'foo': 'bar'}
>>> get_config('myplugin')
{}
>>> get_config('flatpages')
{'directory': '/srv/mediagoblin/pages', 'nesting': 1}}
```

`mediagoblin.tools.pluginapi.register_routes(routes)`
Registers one or more routes

If your plugin handles requests, then you need to call this with the routes your plugin handles.

A “route” is a `routes.Route` object. See [the routes.Route documentation](#) for more details.

Example passing in a single route:

```
>>> register_routes(('about-view', '/about',
...                 'mediagoblin.views:about_view_handler'))
```

Example passing in a list of routes:

```
>>> register_routes([
...     ('contact-view', '/contact', 'mediagoblin.views:contact_handler'),
...     ('about-view', '/about', 'mediagoblin.views:about_handler')
... ])
```

Note: Be careful when designing your route urls. If they clash with core urls, then it could result in DISASTER!

`mediagoblin.tools.pluginapi.register_template_path(path)`
Registers a path for template loading

If your plugin has templates, then you need to call this with the absolute path of the root of templates directory.

Example:

```
>>> my_plugin_dir = os.path.dirname(__file__)
>>> template_dir = os.path.join(my_plugin_dir, 'templates')
>>> register_template_path(template_dir)
```

Note: You can only do this in `setup_plugins()`. Doing this after that will have no effect on template loading.

`mediagoblin.tools.pluginapi.register_template_hooks(template_hooks)`
Register a dict of template hooks.

Takes `template_hooks` as an argument, which is a dictionary of template hook names/keys to the templates they should provide. (The value can either be a single template path or an iterable of paths.)

Example:

```
{ "media_sidebar": "/plugin/sidemess/mess_up_the_side.html",  
  "media_descriptionbox": [ "/plugin/sidemess/even_more_mess.html",  
                           "/plugin/sidemess/so_much_mess.html" ] }
```

`mediagoblin.tools.pluginapi.get_hook_templates` (*hook_name*)

Get a list of hook templates for this `hook_name`.

Note: for the most part, you access this via a template tag, not this method directly, like so:

```
{% template_hook "media_sidebar" %}
```

... which will include all templates for you, partly using this method.

However, this method is exposed to templates, and if you wish, you can iterate over templates in a template hook manually like so:

```
{% for template_path in get_hook_templates("media_sidebar") %}  
  <div class="extra_structure">  
    {% include template_path %}  
  </div>  
{% endfor %}
```

Returns: A list of strings representing template paths.

PART 4: DEVELOPER'S ZONE

This chapter contains various information for developers.

4.1 Codebase Documentation

Sections

- [Software Stack](#)
- [What's where](#)

This chapter covers the libraries that GNU MediaGoblin uses as well as various recipes for getting things done.

Note: This chapter is in flux. Clearly there are things here that aren't documented. If there's something you have questions about, please ask!

See the [join page](#) on the [website](#) for where we hang out.

For more information on how to get started hacking on GNU MediaGoblin, see [the wiki](#).

4.1.1 Software Stack

- Project infrastructure
 - [Python](#): the language we're using to write this
 - [Nose](#): for unit tests
 - [virtualenv](#): for setting up an isolated environment to keep mediagoblin and related packages (potentially not required if MediaGoblin is packaged for your distro)
- Data storage
 - [SQLAlchemy](#): SQL ORM and database interaction library for Python. Currently we support sqlite and postgres as backends.
- Web application
 - [Paste Deploy](#) and [Paste Script](#): we'll use this for configuring and launching the application
 - [werkzeug](#): nice abstraction layer from HTTP requests, responses and WSGI bits
 - [Beaker](#): for handling sessions and caching

- `Jinja2`: the templating engine
- `WTForms`: for handling, validation, and abstraction from HTML forms
- `Celery`: for task queuing (resizing images, encoding video, ...)
- `Babel`: Used to extract and compile translations.
- `Markdown (for python)`: implementation of `Markdown` text-to-html tool to make it easy for people to write richtext comments, descriptions, and etc.
- `lxml`: nice xml and html processing for python.
- Media processing libraries
 - `Python Imaging Library`: used to resize and otherwise convert images for display.
 - `GStreamer`: (Optional, for video hosting sites only) Used to transcode video, and in the future, probably audio too.
 - `chardet`: (Optional, for ascii art hosting sites only) Used to make ascii art thumbnails.
- Front end
 - `JQuery`: for groovy JavaScript things

4.1.2 What's where

After you've run checked out `mediagoblin` and followed the `virtualenv` instantiation instructions, you're faced with the following directory tree:

```
mediagoblin/
|- mediagoblin/           # source code
| | |- tests/
| | |- templates/
| | |- auth/
| | \- submit/
|- docs/                 # documentation
|- devtools/             # some scripts for developer convenience
|
| # the below directories are installed into your virtualenv checkout
|
|- bin/                  # scripts
|- develop-eggs/
|- lib/                  # python libraries installed into your virtualenv
|- include/
|- mediagoblin.egg-info/
|- parts/
|- user_dev/             # sessions, etc
```

As you can see, all the code for GNU MediaGoblin is in the `mediagoblin` directory.

Here are some interesting files and what they do:

`routing.py` maps url paths to views

`views.py` views handle http requests

`models.py` holds the sqlalchemy schemas—these are the data structures we're working with

You'll notice that there are several sub-directories: `tests`, `templates`, `auth`, `submit`, ...

`tests` holds the unit test code.

`templates` holds all the templates for the output.

`auth` and `submit` are modules that encapsulate authentication and media item submission. If you look in these directories, you'll see they have their own `routing.py`, `view.py`, and `models.py` in addition to some other code.

4.2 Original Design Decisions

Sections

- Why GNU MediaGoblin?
- Why Python
- Why WSGI Minimalism
- Why MongoDB
- Why Sphinx for documentation
- Why AGPLv3 and CC0?
- Why (non-mandatory) copyright assignment?

This chapter talks a bit about design decisions.

4.2.1 Why GNU MediaGoblin?

Chris and Will on “Why GNU MediaGoblin”:

Chris came up with the name MediaGoblin. The name is pretty fun. It merges the idea that this is a Media hosting project with Goblin which sort of sounds like gobbling. Here's a piece of software that gobbles up your media for all to see.

According to Wikipedia, a goblin is:

a legendary evil or mischievous illiterate creature, described as grotesquely evil or evil-like phantom

So are we evil? No. Are we mischievous or illiterate? Not really. So what kind of goblin are we thinking about? We're thinking about these goblins:

Those are pretty cute goblins. Those are the kinds of goblins we're thinking about.

Chris started doing work on the project after thinking about it for a year. Then, after talking with Matt and Rob, it became an official GNU project. Thus we now call it GNU MediaGoblin.

That's a lot of letters, though, so in the interest of brevity and facilitating easier casual conversation and balancing that with what's important to us, we have the following rules:

1. “GNU MediaGoblin” is the name we're going to use in all official capacities: web site, documentation, press releases, ...
2. In casual conversation, it's ok to use more casual names.
3. If you're writing about the project, we ask that you call it GNU MediaGoblin.
4. If you don't like the name, we kindly ask you to take a deep breath, think a happy thought about cute little goblins playing on a playground and taking cute pictures of themselves, and let it go. (Will added this one.)



Figure 4.1: *Figure 1: Cute goblin with a beret. Illustrated by Chris Webber*



Figure 4.2: *Figure 2: Snuggly goblin. Illustrated by Karen Rustad*

4.2.2 Why Python

Chris Webber on “Why Python”:

Because I know Python, love Python, am capable of actually making this thing happen in Python (I’ve worked on a lot of large free software web applications before in Python, including [Miro Community](#), the [Miro Guide](#), a large portion of [Creative Commons](#), and a whole bunch of things while working at [Imaginary Landscape](#)). Me starting a project like this makes sense if it’s done in Python.

You might say that PHP is way more deployable, that Rails has way more cool developers riding around on fixie bikes—and all of those things are true. But I know Python, like Python, and think that Python is pretty great. I do think that deployment in Python is not as good as with PHP, but I think the days of shared hosting are (thankfully) coming to an end, and will probably be replaced by cheap virtual machines spun up on the fly for people who want that sort of stuff, and Python will be a huge part of that future, maybe even more than PHP will. The deployment tools are getting better. Maybe we can use something like Silver Lining. Maybe we can just distribute as `.debs` or `.rpms`. We’ll figure it out when we get there.

Regardless, if I’m starting this project, which I am, it’s gonna be in Python.

4.2.3 Why WSGI Minimalism

Chris Webber on “Why WSGI Minimalism”:

If you notice in the technology list I list a lot of components that are very “django-like”, but not actually [Django](#) components. What can I say, I really like a lot of the ideas in Django! Which leads to the question: why not just use Django?

While I really like Django’s ideas and a lot of its components, I also feel that most of the best ideas in Django I want have been implemented as good or even better outside of Django. I could just use Django and replace the templating system with Jinja2, and the form system with wtforms, and the database with MongoDB and MongoKit, but at that point, how much of Django is really left?

I also am sometimes saddened and irritated by how coupled all of Django’s components are. Loosely coupled yes, but still coupled. WSGI has done a good job of providing a base layer for running applications on and if you know how to do it yourself ¹, it’s not hard or many lines of code at all to bind them together without any framework at all (not even say [Pylons](#), [Pyramid](#) or [Flask](#) which I think are still great projects, especially for people who want this sort of thing but have no idea how to get started). And even at this already really early stage of writing MediaGoblin, that glue work is mostly done.

Not to say I don’t think Django isn’t great for a lot of things. For a lot of stuff, it’s still the best, but not for MediaGoblin, I think.

One thing that Django does super well though is documentation. It still has some faults, but even with those considered I can hardly think of any other project in Python that has as nice of documentation as Django. It may be worth learning some lessons on documentation from Django ², on that note.

I’d really like to have a good, thorough hacking-howto and deployment-howto, especially in the former making some notes on how to make it easier for Django hackers to get started.

4.2.4 Why MongoDB

Chris Webber on “Why MongoDB”:

¹ <http://pythonpaste.org/webob/do-it-yourself.html>

² <http://pycon.blip.tv/file/4881071/>

In case you were wondering, I am not a NOSQL fanboy, I do not go around telling people that MongoDB is web scale. Actually my choice for MongoDB isn't scalability, though scaling up really nicely is a pretty good feature and sets us up well in case large volume sites eventually do use MediaGoblin. But there's another side of scalability, and that's scaling down, which is important for federation, maybe even more important than scaling up in an ideal universe where everyone ran servers out of their own housing. As a memory-mapped database, MongoDB is pretty hungry, so actually I spent a lot of time debating whether the inability to scale down as nicely as something like SQL has with sqlite meant that it was out.

But I decided in the end that I really want MongoDB, not for scalability, but for flexibility. Schema evolution pains in SQL are almost enough reason for me to want MongoDB, but not quite. The real reason is because I want the ability to eventually handle multiple media types through MediaGoblin, and also allow for plugins, without the rigidity of tables making that difficult. In other words, something like:

```
{ "title": "Me talking until you are bored",
  "description": "blah blah blah",
  "media_type": "audio",
  "media_data": {
    "length": "2:30",
    "codec": "OGG Vorbis"},
  "plugin_data": {
    "licensing": {
      "license": "http://creativecommons.org/licenses/by-sa/3.0/"}}
```

Being able to just dump media-specific information in a media_data hashtable is pretty great, and even better is having a plugin system where you can just let plugins have their own entire key-value space cleanly inside the document that doesn't interfere with anyone else's stuff. If we were to let plugins to deposit their own information inside the database, either we'd let plugins create their own tables which makes SQL migrations even harder than they already are, or we'd probably end up creating a table with a column for key, a column for value, and a column for type in one huge table called "plugin_data" or something similar. (Yo dawg, I heard you liked plugins, so I put a database in your database so you can query while you query.) Gross.

I also don't want things to be too loose so that we forget or lose the structure of things, and that's one reason why I want to use MongoKit, because we can cleanly define a much structure as we want and verify that documents match that structure generally without adding too much bloat or overhead (MongoKit is a pretty lightweight wrapper and doesn't inject extra MongoKit-specific stuff into the database, which is nice and nicer than many other ORMs in that way).

4.2.5 Why Sphinx for documentation

Will Kahn-Greene on "Why Sphinx":

[Sphinx](#) is a fantastic tool for organizing documentation for a Python-based project that makes it pretty easy to write docs that are readable in source form and can be "compiled" into HTML, LaTeX and other formats.

There are other doc systems out there, but given that GNU MediaGoblin is being written in Python and I've done a ton of documentation using Sphinx, it makes sense to use Sphinx for now.

4.2.6 Why AGPLv3 and CC0?

Chris, Brett, Will, Rob, Matt, et al curated into a story where everyone is the hero by Will on "Why AGPLv3 and CC0":

The [AGPL v3](#) preserves the freedoms guaranteed by the GPL v3 in the context of software as a service. Using this license ensures that users of the service have the ability to examine the source, deploy their own

instance, and implement their own version. This is really important to us and a core mission component of this project. Thus we decided that the software parts should be under this license.

However, the project is made up of more than just software: there's CSS, images, and other output-related things. We wanted the templates/images/css side of the project all permissive and permissive in the same absolutely permissive way. We're waiving our copyrights to non-software things under the CC0 waiver.

That brings us to the templates where there's some code and some output. The template engine we're using is called Jinja2. It mixes HTML markup with Python code to render the output of the software. We decided the templates are part of the output of the software and not the software itself. We wanted the output of the software to be licensed in a hassle-free way so that when someone deploys their own GNU MediaGoblin instance with their own templates, they don't have to deal with the copyleft aspects of the AGPLv3 and we'd be fine with that because the changes they're making are identity-related. So at first we decided to waive our copyrights to the templates with a CC0 waiver and then add an exception to the AGPLv3 for the software such that the templates can make calls into the software and yet be a separately licensed work. However, Brett brought up the question of whether this allows some unscrupulous person to make changes to the software through the templates in such a way that they're not bound by the AGPLv3: i.e. a loophole. We thought about this loophole and between this and the extra legalese involved in the exception to the AGPLv3, we decided that it's just way simpler if the templates were also licensed under the AGPLv3.

Then we have the licensing for the documentation. Given that the documentation is tied to the software content-wise, we don't feel like we have to worry about ensuring freedom of the documentation or worry about attribution concerns. Thus we're waiving our copyrights to the documentation under CC0 as well.

Lastly, we have branding. This covers logos and other things that are distinctive to GNU MediaGoblin that we feel represents this project. Since we don't currently have any branding, this is an open issue, but we're thinking we'll go with a CC BY-SA license.

By licensing in this way, we make sure that users of the software receive the freedoms that the AGPLv3 ensures regardless of what fate befalls this project.

So to summarize:

- software (Python, JavaScript, HTML templates): licensed under AGPLv3
- non-software things (CSS, images, video): copyrights waived under CC0 because this is output of the software
- documentation: copyrights waived under CC0 because it's not part of the software
- branding assets: we're kicking this can down the road, but probably CC BY-SA

This is all codified in the `COPYING` file.

4.2.7 Why (non-mandatory) copyright assignment?

Chris Webber on "Why copyright assignment?":

GNU MediaGoblin is a GNU project with non-mandatory but heavily encouraged copyright assignment to the FSF. Most, if not all, of the core contributors to GNU MediaGoblin will have done a copyright assignment, but unlike some other GNU projects, it isn't required here. We think this is the best choice for GNU MediaGoblin: it ensures that the Free Software Foundation may protect the software by enforcing the AGPL if the FSF sees fit, but it also means that we can immediately merge in changes from a new contributor. It also means that some significant non-FSF contributors might also be able to enforce the AGPL if seen fit.

Again, assignment is not mandatory, but it is heavily encouraged, even incentivized: significant contributors who do a copyright assignment to the FSF are eligible to have a unique goblin drawing produced for them by the project's main founder, Christopher Allan Webber. See [the wiki](#) for details.

INDICES AND TABLES

- *search*
- *genindex*

This guide was built on July 30, 2013.

PYTHON MODULE INDEX

m

mediagoblin.tools.pluginapi, ??