

Getting Started with Engines

January 13, 2015

In this guide you will learn about engines and how they can be used to provide additional functionality to their host applications through a clean and very easy-to-use interface.

After reading this guide, you will know:

- What makes an engine.
- How to generate an engine.
- Building features for the engine.
- Hooking the engine into an application.
- Overriding engine functionality in the application.

1 What are engines?

Engines can be considered miniature applications that provide functionality to their host applications. A Rails application is actually just a “supercharged” engine, with the `Rails::Application` class inheriting a lot of its behavior from `Rails::Engine`.

Therefore, engines and applications can be thought of almost the same thing, just with subtle differences, as you’ll see throughout this guide. Engines and applications also share a common structure.

Engines are also closely related to plugins. The two share a common `lib` directory structure, and are both generated using the `rails plugin new` generator. The difference is that an engine is considered a “full plugin” by Rails (as indicated by the `--full` option that’s passed to the generator command). We’ll actually be using the `--mountable` option here, which includes all the features of `--full`, and then some. This guide will refer to these “full plugins” simply as “engines” throughout. An engine **can** be a plugin, and a plugin **can** be an engine.

The engine that will be created in this guide will be called “blorgh”. This engine will provide blogging functionality to its host applications, allowing for new articles and comments to be created. At the beginning of this guide, you will be working solely within the engine itself, but in later sections you’ll see how to hook it into an application.

Engines can also be isolated from their host applications. This means that an application is able to have a path provided by a routing helper such as `articles_path` and use an engine also that provides a path also called `articles_path`, and the two would not clash. Along with this, controllers, models and table names are also namespaced. You’ll see how to do this later in this guide.

It’s important to keep in mind at all times that the application should **always** take precedence over its engines. An application is the object that has final say in what goes on in its environment. The engine should only be enhancing it, rather than changing it drastically.

To see demonstrations of other engines, check out Devise, an engine that provides authentication for its parent applications, or Forem, an engine that provides forum functionality. There's also Spree which provides an e-commerce platform, and RefineryCMS, a CMS engine.

Finally, engines would not have been possible without the work of James Adam, Piotr Sarnacki, the Rails Core Team, and a number of other people. If you ever meet them, don't forget to say thanks!

2 Generating an engine

To generate an engine, you will need to run the plugin generator and pass it options as appropriate to the need. For the “blorgh” example, you will need to create a “mountable” engine, running this command in a terminal:

```
$ rails plugin new blorgh --mountable
```

The full list of options for the plugin generator may be seen by typing:

```
$ rails plugin --help
```

The `--mountable` option tells the generator that you want to create a “mountable” and namespace-isolated engine. This generator will provide the same skeleton structure as would the `--full` option. The `--full` option tells the generator that you want to create an engine, including a skeleton structure that provides the following:

- An app directory tree
- A `config/routes.rb` file:

```
Rails.application.routes.draw do
  end
```

- A file at `lib/blorgh/engine.rb`, which is identical in function to a standard Rails application's `config/application.rb` file:

```
module Blorgh
  class Engine < ::Rails::Engine
    end
  end
end
```

The `--mountable` option will add to the `--full` option:

- Asset manifest files (`application.js` and `application.css`)
- A namespaced `ApplicationController` stub
- A namespaced `ApplicationHelper` stub
- A layout view template for the engine
- Namespace isolation to `config/routes.rb`:

```
Blorgh::Engine.routes.draw do
end
```

- Namespace isolation to `lib/blorgh/engine.rb`:

```
module Blorgh
  class Engine < ::Rails::Engine
    isolate_namespace Blorgh
  end
end
```

Additionally, the `--mountable` option tells the generator to mount the engine inside the dummy testing application located at `test/dummy` by adding the following to the dummy application's routes file at `test/dummy/config/routes.rb`:

```
mount Blorgh::Engine => "/blorgh"
```

2.1 Inside an Engine

2.1.1 Critical Files At the root of this brand new engine's directory lives a `blorgh.gemspec` file. When you include the engine into an application later on, you will do so with this line in the Rails application's `Gemfile`:

```
gem 'blorgh', path: "vendor/engines/blorgh"
```

Don't forget to run `bundle install` as usual. By specifying it as a gem within the `Gemfile`, Bundler will load it as such, parsing this `blorgh.gemspec` file and requiring a file within the `lib` directory called `lib/blorgh.rb`. This file requires the `blorgh/engine.rb` file (located at `lib/blorgh/engine.rb`) and defines a base module called `Blorgh`.

```
require "blorgh/engine"
```

```
module Blorgh
end
```

Some engines choose to use this file to put global configuration options for their engine. It's a relatively good idea, so if you want to offer configuration options, the file where your engine's module is defined is perfect for that. Place the methods inside the module and you'll be good to go.

Within `lib/blorgh/engine.rb` is the base class for the engine:

```
module Blorgh
  class Engine < ::Rails::Engine
    isolate_namespace Blorgh
  end
end
```

By inheriting from the `Rails::Engine` class, this gem notifies Rails that there's an engine at the specified path, and will correctly mount the engine inside the application, performing tasks such as adding the `app` directory of the engine to the load path for models, mailers, controllers and views.

The `isolate_namespace` method here deserves special notice. This call is responsible for isolating the controllers, models, routes and other things into their own namespace, away from similar components inside the application. Without this, there is a possibility that the engine's components could "leak" into the application, causing unwanted disruption, or that important engine components could be overridden by similarly named things within the application. One of the examples of such conflicts is helpers. Without calling `isolate_namespace`, the engine's helpers would be included in an application's controllers.

It is **highly** recommended that the `isolate_namespace` line be left within the `Engine` class definition. Without it, classes generated in an engine **may** conflict with an application.

What this isolation of the namespace means is that a model generated by a call to `bin/rails g model`, such as `bin/rails g model article`, won't be called `Article`, but instead be namespaced and called `Blorgh::Article`. In addition, the table for the model is namespaced, becoming `blorgh_articles`, rather than simply `articles`. Similar to the model namespacing, a controller called `ArticlesController` becomes `Blorgh::ArticlesController` and the views for that controller will not be at `app/views/articles`, but `app/views/blorgh/articles` instead. Mailers are namespaced as well.

Finally, routes will also be isolated within the engine. This is one of the most important parts about namespacing, and is discussed later in the Routes section of this guide.

2.1.2 app Directory Inside the `app` directory are the standard `assets`, `controllers`, `helpers`, `mailers`, `models` and `views` directories that you should be familiar with from an application. The `helpers`, `mailers` and `models` directories are empty, so they aren't described in this section. We'll look more into models in a future section, when we're writing the engine.

Within the `app/assets` directory, there are the `images`, `javascripts` and `stylesheets` directories which, again, you should be familiar with due to their similarity to an application. One difference here, however, is that each directory contains a sub-directory with the engine name. Because this engine is going to be namespaced, its assets should be too.

Within the `app/controllers` directory there is a `blorgh` directory that contains a file called `application_controller.rb`. This file will provide any common functionality for the controllers of the engine. The `blorgh` directory is where the other controllers for the engine will go. By placing them within this namespaced directory, you prevent them from possibly clashing with identically-named controllers within other engines or even within the application.

The `ApplicationController` class inside an engine is named just like a Rails application in order to make it easier for you to convert your applications into engines.

Lastly, the `app/views` directory contains a `layouts` folder, which contains a file at `blorgh/application.html.erb`. This file allows you to specify a layout for the engine. If this engine is to be used as a stand-alone engine, then you would add any customization to its layout in this file, rather than the application's `app/views/layouts/application.html.erb` file.

If you don't want to force a layout on to users of the engine, then you can delete this file and reference a different layout in the controllers of your engine.

2.1.3 bin Directory This directory contains one file, `bin/rails`, which enables you to use the `rails` sub-commands and generators just like you would within an application. This means that you will be able to generate new controllers and models for this engine very easily by running commands like this:

```
$ bin/rails g model
```

Keep in mind, of course, that anything generated with these commands inside of an engine that has `isolate_namespace` in the `Engine` class will be namespaced.

2.1.4 test Directory The `test` directory is where tests for the engine will go. To test the engine, there is a cut-down version of a Rails application embedded within it at `test/dummy`. This application will mount the engine in the `test/dummy/config/routes.rb` file:

```
Rails.application.routes.draw do
  mount Blorgh::Engine => "/blorgh"
end
```

This line mounts the engine at the path `/blorgh`, which will make it accessible through the application only at that path.

Inside the test directory there is the `test/integration` directory, where integration tests for the engine should be placed. Other directories can be created in the `test` directory as well. For example, you may wish to create a `test/models` directory for your model tests.

3 Providing engine functionality

The engine that this guide covers provides submitting articles and commenting functionality and follows a similar thread to the Getting Started Guide, with some new twists.

3.1 Generating an Article Resource

The first thing to generate for a blog engine is the `Article` model and related controller. To quickly generate this, you can use the Rails scaffold generator.

```
$ bin/rails generate scaffold article title:string text:text
```

This command will output this information:

```
invoke active_record
create db/migrate/[timestamp]_create_blorgh_articles.rb
create app/models/blorgh/article.rb
invoke test_unit
create test/models/blorgh/article_test.rb
create test/fixtures/blorgh/articles.yml
invoke resource_route
route resources :articles
invoke scaffold_controller
create app/controllers/blorgh/articles_controller.rb
invoke erb
create app/views/blorgh/articles
create app/views/blorgh/articles/index.html.erb
```

```

create      app/views/lorgh/articles/edit.html.erb
create      app/views/lorgh/articles/show.html.erb
create      app/views/lorgh/articles/new.html.erb
create      app/views/lorgh/articles/_form.html.erb
invoke      test_unit
create      test/controllers/lorgh/articles_controller_test.rb
invoke      helper
create      app/helpers/lorgh/articles_helper.rb
invoke      assets
invoke      js
create      app/assets/javascripts/lorgh/articles.js
invoke      css
create      app/assets/stylesheets/lorgh/articles.css
invoke      css
create      app/assets/stylesheets/scaffold.css

```

The first thing that the scaffold generator does is invoke the `active_record` generator, which generates a migration and a model for the resource. Note here, however, that the migration is called `create.lorgh_articles` rather than the usual `create_articles`. This is due to the `isolate_namespace` method called in the `Blorgh::Engine` class's definition. The model here is also namespaced, being placed at `app/models/lorgh/article.rb` rather than `app/models/article.rb` due to the `isolate_namespace` call within the `Engine` class.

Next, the `test_unit` generator is invoked for this model, generating a model test at `test/models/lorgh/article_test.rb` (rather than `test/models/article_test.rb`) and a fixture at `test/fixtures/lorgh/articles.yml` (rather than `test/fixtures/articles.yml`).

After that, a line for the resource is inserted into the `config/routes.rb` file for the engine. This line is simply `resources :articles`, turning the `config/routes.rb` file for the engine into this:

```

Blorgh::Engine.routes.draw do
  resources :articles
end

```

Note here that the routes are drawn upon the `Blorgh::Engine` object rather than the `YourApp::Application` class. This is so that the engine routes are confined to the engine itself and can be mounted at a specific point as shown in the test directory section. It also causes the engine's routes to be isolated from those routes that are within the application. The Routes section of this guide describes it in detail.

Next, the `scaffold_controller` generator is invoked, generating a controller called `Blorgh::ArticlesController` (at `app/controllers/lorgh/articles_controller.rb`) and its related views at `app/views/lorgh/articles`. This generator also generates a test for the controller (`test/controllers/lorgh/articles_controller_test.rb`) and a helper (`app/helpers/lorgh/articles_controller.rb`).

Everything this generator has created is neatly namespaced. The controller's class is defined within the `Blorgh` module:

```

module Blorgh
  class ArticlesController < ApplicationController
    ...
  end
end

```

The `ApplicationController` class being inherited from here is the `Blorgh::ApplicationController` , not an application's `ApplicationController` .

The helper inside `app/helpers/blorgh/articles_helper.rb` is also namespaced:

```
module Blorgh
  module ArticlesHelper
    ...
  end
end
```

This helps prevent conflicts with any other engine or application that may have an article resource as well.

Finally, the assets for this resource are generated in two files: `app/assets/javascripts/blorgh/articles.js` and `app/assets/stylesheets/blorgh/articles.css` . You'll see how to use these a little later.

By default, the scaffold styling is not applied to the engine because the engine's layout file, `app/views/layouts/blorgh/application.html.erb` , doesn't load it. To make the scaffold styling apply, insert this line into the `<head>` tag of this layout:

```
<%= stylesheet_link_tag "scaffold" %>
```

You can see what the engine has so far by running `rake db:migrate` at the root of our engine to run the migration generated by the scaffold generator, and then running `rails server` in `test/dummy` . When you open `http://localhost:3000/blorgh/articles` you will see the default scaffold that has been generated. Click around! You've just generated your first engine's first functions.

If you'd rather play around in the console, `rails console` will also work just like a Rails application. Remember: the `Article` model is namespaced, so to reference it you must call it as `Blorgh::Article` .

```
>> Blorgh::Article.find(1)
=> #<Blorgh::Article id: 1 ...>
```

One final thing is that the `articles` resource for this engine should be the root of the engine. Whenever someone goes to the root path where the engine is mounted, they should be shown a list of articles. This can be made to happen if this line is inserted into the `config/routes.rb` file inside the engine:

```
root to: "articles#index"
```

Now people will only need to go to the root of the engine to see all the articles, rather than visiting `/articles` . This means that instead of `http://localhost:3000/blorgh/articles` , you only need to go to `http://localhost:3000/blorgh` now.

3.2 Generating a Comments Resource

Now that the engine can create new articles, it only makes sense to add commenting functionality as well. To do this, you'll need to generate a comment model, a comment controller and then modify the articles scaffold to display comments and allow people to create new ones.

From the application root, run the model generator. Tell it to generate a `Comment` model, with the related table having two columns: a `article_id` integer and `text` text column.

```
$ bin/rails generate model Comment article_id:integer text:text
```

This will output the following:

```
invoke  active_record
create  db/migrate/[timestamp]_create_blorgh_comments.rb
create  app/models/lorgh/comment.rb
invoke  test_unit
create  test/models/lorgh/comment_test.rb
create  test/fixtures/lorgh/comments.yml
```

This generator call will generate just the necessary model files it needs, namespacing the files under a `lorgh` directory and creating a model class called `Blorgh::Comment`. Now run the migration to create our `lorgh_comments` table:

```
$ rake db:migrate
```

To show the comments on an article, edit `app/views/lorgh/articles/show.html.erb` and add this line before the “Edit” link:

```
<h3>Comments</h3>
<%= render @article.comments %>
```

This line will require there to be a `has_many` association for comments defined on the `Blorgh::Article` model, which there isn’t right now. To define one, open `app/models/lorgh/article.rb` and add this line into the model:

```
has_many :comments
```

Turning the model into this:

```
module Blorgh
  class Article < ActiveRecord::Base
    has_many :comments
  end
end
```

Because the `has_many` is defined inside a class that is inside the `Blorgh` module, Rails will know that you want to use the `Blorgh::Comment` model for these objects, so there’s no need to specify that using the `:class_name` option here.

Next, there needs to be a form so that comments can be created on an article. To add this, put this line underneath the call to `render @article.comments` in `app/views/lorgh/articles/show.html.erb`:

```
<%= render "lorgh/comments/form" %>
```

Next, the partial that this line will render needs to exist. Create a new directory at `app/views/lorgh/comments` and in it a new file called `_form.html.erb` which has this content to create the required partial:

```

<h3>New comment</h3>
<%= form_for [@article, @article.comments.build] do |f| %>
  <p>
    <%= f.label :text %><br>
    <%= f.text_area :text %>
  </p>
  <%= f.submit %>
<% end %>

```

When this form is submitted, it is going to attempt to perform a POST request to a route of `/articles/:article_id/comments` within the engine. This route doesn't exist at the moment, but can be created by changing the `resources :articles` line inside `config/routes.rb` into these lines:

```

resources :articles do
  resources :comments
end

```

This creates a nested route for the comments, which is what the form requires.

The route now exists, but the controller that this route goes to does not. To create it, run this command from the application root:

```
$ bin/rails g controller comments
```

This will generate the following things:

```

create  app/controllers/blrgh/comments_controller.rb
invoke  erb
  exist  app/views/blrgh/comments
invoke  test_unit
create  test/controllers/blrgh/comments_controller_test.rb
invoke  helper
create  app/helpers/blrgh/comments_helper.rb
invoke  assets
invoke  js
create  app/assets/javascripts/blrgh/comments.js
invoke  css
create  app/assets/stylesheets/blrgh/comments.css

```

The form will be making a POST request to `/articles/:article_id/comments`, which will correspond with the `create` action in `Blorgh::CommentsController`. This action needs to be created, which can be done by putting the following lines inside the class definition in `app/controllers/blrgh/comments_controller.rb`:

```

def create
  @article = Article.find(params[:article_id])
  @comment = @article.comments.create(comment_params)
  flash[:notice] = "Comment has been created!"
end

```

```

    redirect_to articles_path
  end

  private
  def comment_params
    params.require(:comment).permit(:text)
  end

```

This is the final step required to get the new comment form working. Displaying the comments, however, is not quite right yet. If you were to create a comment right now, you would see this error:

```

Missing partial blorgh/comments/comment with {:handlers=>[:erb, :builder],
:formats=>[:html], :locale=>[:en, :en]}. Searched in: *
"/Users/ryan/Sites/side_projects/blorgh/test/dummy/app/views" *
"/Users/ryan/Sites/side_projects/blorgh/app/views"

```

The engine is unable to find the partial required for rendering the comments. Rails looks first in the application's (`test/dummy`) `app/views` directory and then in the engine's `app/views` directory. When it can't find it, it will throw this error. The engine knows to look for `blorgh/comments/comment` because the model object it is receiving is from the `Blorgh::Comment` class.

This partial will be responsible for rendering just the comment text, for now. Create a new file at `app/views/blorgh/comments/_comment.html.erb` and put this line inside it:

```
<%= comment_counter + 1 %>. <%= comment.text %>
```

The `comment_counter` local variable is given to us by the `<%= render @article.comments %>` call, which will define it automatically and increment the counter as it iterates through each comment. It's used in this example to display a small number next to each comment when it's created.

That completes the comment function of the blogging engine. Now it's time to use it within an application.

4 Hooking Into an Application

Using an engine within an application is very easy. This section covers how to mount the engine into an application and the initial setup required, as well as linking the engine to a `User` class provided by the application to provide ownership for articles and comments within the engine.

4.1 Mounting the Engine

First, the engine needs to be specified inside the application's `Gemfile`. If there isn't an application handy to test this out in, generate one using the `rails new` command outside of the engine directory like this:

```
$ rails new unicorn
```

Usually, specifying the engine inside the `Gemfile` would be done by specifying it as a normal, everyday gem.

```
gem 'devise'
```

However, because you are developing the `blorgh` engine on your local machine, you will need to specify the `:path` option in your `Gemfile`:

```
gem 'blorgh', path: "/path/to/blorgh"
```

Then run `bundle` to install the gem.

As described earlier, by placing the gem in the `Gemfile` it will be loaded when Rails is loaded. It will first require `lib/blorgh.rb` from the engine, then `lib/blorgh/engine.rb`, which is the file that defines the major pieces of functionality for the engine.

To make the engine's functionality accessible from within an application, it needs to be mounted in that application's `config/routes.rb` file:

```
mount Blorgh::Engine, at: "/blog"
```

This line will mount the engine at `/blog` in the application. Making it accessible at `http://localhost:3000/blog` when the application runs with `rails server`.

Other engines, such as `Devise`, handle this a little differently by making you specify custom helpers (such as `devise_for`) in the routes. These helpers do exactly the same thing, mounting pieces of the engines' functionality at a pre-defined path which may be customizable.

4.2 Engine setup

The engine contains migrations for the `blorgh_articles` and `blorgh_comments` table which need to be created in the application's database so that the engine's models can query them correctly. To copy these migrations into the application use this command:

```
$ rake blorgh:install:migrations
```

If you have multiple engines that need migrations copied over, use `railties:install:migrations` instead:

```
$ rake railties:install:migrations
```

This command, when run for the first time, will copy over all the migrations from the engine. When run the next time, it will only copy over migrations that haven't been copied over already. The first run for this command will output something such as this:

```
Copied migration [timestamp_1]_create_blorgh_articles.rb from blorgh
Copied migration [timestamp_2]_create_blorgh_comments.rb from blorgh
```

The first timestamp (`[timestamp_1]`) will be the current time, and the second timestamp (`[timestamp_2]`) will be the current time plus a second. The reason for this is so that the migrations for the engine are run after any existing migrations in the application.

To run these migrations within the context of the application, simply run `rake db:migrate`. When accessing the engine through `http://localhost:3000/blog`, the articles will be empty. This is because the table created inside the application is different from the one created within the engine. Go ahead, play around with the newly mounted engine. You'll find that it's the same as when it was only an engine.

If you would like to run migrations only from one engine, you can do it by specifying `SCOPE`:

```
rake db:migrate SCOPE=blorgh
```

This may be useful if you want to revert engine's migrations before removing it. To revert all migrations from blorgh engine you can run code such as:

```
rake db:migrate SCOPE=blorgh VERSION=0
```

4.3 Using a Class Provided by the Application

4.3.1 Using a Model Provided by the Application When an engine is created, it may want to use specific classes from an application to provide links between the pieces of the engine and the pieces of the application. In the case of the `blorgh` engine, making articles and comments have authors would make a lot of sense.

A typical application might have a `User` class that would be used to represent authors for an article or a comment. But there could be a case where the application calls this class something different, such as `Person`. For this reason, the engine should not hardcode associations specifically for a `User` class.

To keep it simple in this case, the application will have a class called `User` that represents the users of the application (we'll get into making this configurable further on). It can be generated using this command inside the application:

```
rails g model user name:string
```

The `rake db:migrate` command needs to be run here to ensure that our application has the `users` table for future use.

Also, to keep it simple, the articles form will have a new text field called `author_name`, where users can elect to put their name. The engine will then take this name and either create a new `User` object from it, or find one that already has that name. The engine will then associate the article with the found or created `User` object.

First, the `author_name` text field needs to be added to the `app/views/blorgh/articles/_form.html.erb` partial inside the engine. This can be added above the `title` field with this code:

```
<div class="field">
  <%= f.label :author_name %><br>
  <%= f.text_field :author_name %>
</div>
```

Next, we need to update our `Blorgh::ArticleController#article_params` method to permit the new form parameter:

```
def article_params
  params.require(:article).permit(:title, :text, :author_name)
end
```

The `Blorgh::Article` model should then have some code to convert the `author_name` field into an actual `User` object and associate it as that article's `author` before the article is saved. It will also need to have an `attr_accessor` set up for this field, so that the setter and getter methods are defined for it.

To do all this, you'll need to add the `attr_accessor` for `author_name`, the association for the author and the `before_save` call into `app/models/blorgh/article.rb`. The `author` association will be hard-coded to the `User` class for the time being.

```

attr_accessor :author_name
belongs_to :author, class_name: "User"

before_save :set_author

private
  def set_author
    self.author = User.find_or_create_by(name: author_name)
  end
end

```

By representing the `author` association's object with the `User` class, a link is established between the engine and the application. There needs to be a way of associating the records in the `blorgh_articles` table with the records in the `users` table. Because the association is called `author`, there should be an `author_id` column added to the `blorgh_articles` table.

To generate this new column, run this command within the engine:

```
$ bin/rails g migration add_author_id_to_blorgh_articles author_id:integer
```

Due to the migration's name and the column specification after it, Rails will automatically know that you want to add a column to a specific table and write that into the migration for you. You don't need to tell it any more than this.

This migration will need to be run on the application. To do that, it must first be copied using this command:

```
$ rake blorgh:install:migrations
```

Notice that only *one* migration was copied over here. This is because the first two migrations were copied over the first time this command was run.

```
NOTE Migration [timestamp]_create_blorgh_articles.rb from blorgh has been
skipped. Migration with the same name already exists. NOTE Migration
[timestamp]_create_blorgh_comments.rb from blorgh has been skipped. Migration
with the same name already exists. Copied migration
[timestamp]_add_author_id_to_blorgh_articles.rb from blorgh
```

Run the migration using:

```
$ rake db:migrate
```

Now with all the pieces in place, an action will take place that will associate an author - represented by a record in the `users` table - with an article, represented by the `blorgh_articles` table from the engine.

Finally, the author's name should be displayed on the article's page. Add this code above the "Title" output inside `app/views/blorgh/articles/show.html.erb`:

```

<p>
  <b>Author:</b>
  <%= @article.author %>
</p>

```

By outputting `@article.author` using the `<%=` tag, the `to_s` method will be called on the object. By default, this will look quite ugly:

```
#<User:0x00000100ccb3b0>
```

This is undesirable. It would be much better to have the user's name there. To do this, add a `to_s` method to the `User` class within the application:

```
def to_s
  name
end
```

Now instead of the ugly Ruby object output, the author's name will be displayed.

4.3.2 Using a Controller Provided by the Application Because Rails controllers generally share code for things like authentication and accessing session variables, they inherit from `ApplicationController` by default. Rails engines, however are scoped to run independently from the main application, so each engine gets a scoped `ApplicationController`. This namespace prevents code collisions, but often engine controllers need to access methods in the main application's `ApplicationController`. An easy way to provide this access is to change the engine's scoped `ApplicationController` to inherit from the main application's `ApplicationController`. For our Blorgh engine this would be done by changing `app/controllers/blorgh/application_controller.rb` to look like:

```
class Blorgh::ApplicationController < ApplicationController
end
```

By default, the engine's controllers inherit from `Blorgh::ApplicationController`. So, after making this change they will have access to the main application's `ApplicationController`, as though they were part of the main application.

This change does require that the engine is run from a Rails application that has an `ApplicationController`.

4.4 Configuring an Engine

This section covers how to make the `User` class configurable, followed by general configuration tips for the engine.

4.4.1 Setting Configuration Settings in the Application The next step is to make the class that represents a `User` in the application customizable for the engine. This is because that class may not always be `User`, as previously explained. To make this setting customizable, the engine will have a configuration setting called `author_class` that will be used to specify which class represents users inside the application.

To define this configuration setting, you should use a `matr_accessor` inside the `Blorgh` module for the engine. Add this line to `lib/blorgh.rb` inside the engine:

```
matr_accessor :author_class
```

This method works like its brothers, `attr_accessor` and `attr_accessor`, but provides a setter and getter method on the module with the specified name. To use it, it must be referenced using `Blorgh.author_class`.

The next step is to switch the `Blorgh::Article` model over to this new setting. Change the `belongs_to` association inside this model (`app/models/blorgh/article.rb`) to this:

```
belongs_to :author, class_name: Blorgh.author_class
```

The `set_author` method in the `Blorgh::Article` model should also use this class:

```
self.author = Blorgh.author_class.constantize.find_or_create_by(name: author_name)
```

To save having to call `constantize` on the `author_class` result all the time, you could instead just override the `author_class` getter method inside the `Blorgh` module in the `lib/blorgh.rb` file to always call `constantize` on the saved value before returning the result:

```
def self.author_class
  @@author_class.constantize
end
```

This would then turn the above code for `set_author` into this:

```
self.author = Blorgh.author_class.find_or_create_by(name: author_name)
```

Resulting in something a little shorter, and more implicit in its behavior. The `author_class` method should always return a `Class` object.

Since we changed the `author_class` method to return a `Class` instead of a `String`, we must also modify our `belongs_to` definition in the `Blorgh::Article` model:

```
belongs_to :author, class_name: Blorgh.author_class.to_s
```

To set this configuration setting within the application, an initializer should be used. By using an initializer, the configuration will be set up before the application starts and calls the engine's models, which may depend on this configuration setting existing.

Create a new initializer at `config/initializers/blorgh.rb` inside the application where the `blorgh` engine is installed and put this content in it:

```
Blorgh.author_class = "User"
```

It's very important here to use the `String` version of the class, rather than the class itself. If you were to use the class, Rails would attempt to load that class and then reference the related table. This could lead to problems if the table wasn't already existing. Therefore, a `String` should be used and then converted to a class using `constantize` in the engine later on.

Go ahead and try to create a new article. You will see that it works exactly in the same way as before, except this time the engine is using the configuration setting in `config/initializers/blorgh.rb` to learn what the class is.

There are now no strict dependencies on what the class is, only what the API for the class must be. The engine simply requires this class to define a `find_or_create_by` method which returns an object of that class, to be associated with an article when it's created. This object, of course, should have some sort of identifier by which it can be referenced.

4.4.2 General Engine Configuration Within an engine, there may come a time where you wish to use things such as initializers, internationalization or other configuration options. The great news is that these things are entirely possible, because a Rails engine shares much the same functionality as a Rails application. In fact, a Rails application's functionality is actually a superset of what is provided by engines!

If you wish to use an initializer - code that should run before the engine is loaded - the place for it is the `config/initializers` folder. This directory's functionality is explained in the Initializers section of the Configuring guide, and works precisely the same way as the `config/initializers` directory inside an application. The same thing goes if you want to use a standard initializer.

For locales, simply place the locale files in the `config/locales` directory, just like you would in an application.

5 Testing an engine

When an engine is generated, there is a smaller dummy application created inside it at `test/dummy`. This application is used as a mounting point for the engine, to make testing the engine extremely simple. You may extend this application by generating controllers, models or views from within the directory, and then use those to test your engine.

The `test` directory should be treated like a typical Rails testing environment, allowing for unit, functional and integration tests.

5.1 Functional Tests

A matter worth taking into consideration when writing functional tests is that the tests are going to be running on an application - the `test/dummy` application - rather than your engine. This is due to the setup of the testing environment; an engine needs an application as a host for testing its main functionality, especially controllers. This means that if you were to make a typical GET to a controller in a controller's functional test like this:

```
module Blorgh
  class FooControllerTest < ActionController::TestCase
    def test_index
      get :index
      ...
    end
  end
end
```

It may not function correctly. This is because the application doesn't know how to route these requests to the engine unless you explicitly tell it **how**. To do this, you must set the `@routes` instance variable to the engine's route set in your setup code:

```
module Blorgh
  class FooControllerTest < ActionController::TestCase
    setup do
      @routes = Engine.routes
    end
  end
end
```

```

    def test_index
      get :index
      ...
    end
  end
end
end

```

This tells the application that you still want to perform a `GET` request to the `index` action of this controller, but you want to use the engine's route to get there, rather than the application's one.

This also ensures that the engine's URL helpers will work as expected in your tests.

6 Improving engine functionality

This section explains how to add and/or override engine MVC functionality in the main Rails application.

6.1 Overriding Models and Controllers

Engine model and controller classes can be extended by open classing them in the main Rails application (since model and controller classes are just Ruby classes that inherit Rails specific functionality). Open classing an Engine class redefines it for use in the main application. This is usually implemented by using the decorator pattern.

For simple class modifications, use `Class#class_eval`. For complex class modifications, consider using `ActiveSupport::Concern`.

6.1.1 A note on Decorators and Loading Code Because these decorators are not referenced by your Rails application itself, Rails' autoloading system will not kick in and load your decorators. This means that you need to require them yourself.

Here is some sample code to do this:

```

# lib/bling/engine.rb
module Bling
  class Engine < ::Rails::Engine
    isolate_namespace Bling

    config.to_prepare do
      Dir.glob(Rails.root + "app/decorators/**/*_decorator*.rb").each do |c|
        require_dependency(c)
      end
    end
  end
end
end

```

This doesn't apply to just Decorators, but anything that you add in an engine that isn't referenced by your main application.

6.1.2 Implementing Decorator Pattern Using Class#class_eval Adding Article#time_since_created:

```
# MyApp/app/decorators/models/blorgh/article_decorator.rb
```

```
Blorgh::Article.class_eval do
  def time_since_created
    Time.current - created_at
  end
end
```

```
# Blorgh/app/models/article.rb
```

```
class Article < ActiveRecord::Base
  has_many :comments
end
```

Overriding Article#summary:

```
# MyApp/app/decorators/models/blorgh/article_decorator.rb
```

```
Blorgh::Article.class_eval do
  def summary
    "#{title} - #{truncate(text)}"
  end
end
```

```
# Blorgh/app/models/article.rb
```

```
class Article < ActiveRecord::Base
  has_many :comments
  def summary
    "#{title}"
  end
end
```

6.1.3 Implementing Decorator Pattern Using ActiveSupport::Concern Using Class#class_eval is great for simple adjustments, but for more complex class modifications, you might want to consider using ActiveSupport::Concern. ActiveSupport::Concern manages load order of interlinked dependent modules and classes at run time allowing you to significantly modularize your code.

Adding Article#time_since_created and Overriding Article#summary:

```
# MyApp/app/models/blorgh/article.rb
```

```
class Blorgh::Article < ActiveRecord::Base
  include Blorgh::Concerns::Models::Article
```

```
def time_since_created
  Time.current - created_at
end

def summary
  "#{title} - #{truncate(text)}"
end
end

# Blorgh/app/models/article.rb

class Article < ActiveRecord::Base
  include Blorgh::Concerns::Models::Article
end

# Blorgh/lib/concerns/models/article

module Blorgh::Concerns::Models::Article
  extend ActiveSupport::Concern

  # 'included do' causes the included code to be evaluated in the
  # context where it is included (article.rb), rather than being
  # executed in the module's context (blorgh/concerns/models/article).
  included do
    attr_accessor :author_name
    belongs_to :author, class_name: "User"

    before_save :set_author

    private
    def set_author
      self.author = User.find_or_create_by(name: author_name)
    end
  end

  def summary
    "#{title}"
  end

  module ClassMethods
    def some_class_method
      'some class method string'
    end
  end
end
```

6.2 Overriding Views

When Rails looks for a view to render, it will first look in the `app/views` directory of the application. If it cannot find the view there, it will check in the `app/views` directories of all engines that have this directory.

When the application is asked to render the view for `Blorgh::ArticlesController`'s `index` action, it will first look for the path `app/views/blrorgh/articles/index.html.erb` within the application. If it cannot find it, it will look inside the engine.

You can override this view in the application by simply creating a new file at `app/views/blrorgh/articles/index.html.erb`. Then you can completely change what this view would normally output.

Try this now by creating a new file at `app/views/blrorgh/articles/index.html.erb` and put this content in it:

```
<h1>Articles</h1>
<%= link_to "New Article", new_article_path %>
<% @articles.each do |article| %>
  <h2><%= article.title %></h2>
  <small>By <%= article.author %></small>
  <%= simple_format(article.text) %>
  <hr>
<% end %>
```

6.3 Routes

Routes inside an engine are isolated from the application by default. This is done by the `isolate_namespace` call inside the `Engine` class. This essentially means that the application and its engines can have identically named routes and they will not clash.

Routes inside an engine are drawn on the `Engine` class within `config/routes.rb`, like this:

```
Blorgh::Engine.routes.draw do
  resources :articles
end
```

By having isolated routes such as this, if you wish to link to an area of an engine from within an application, you will need to use the engine's routing proxy method. Calls to normal routing methods such as `articles_path` may end up going to undesired locations if both the application and the engine have such a helper defined.

For instance, the following example would go to the application's `articles_path` if that template was rendered from the application, or the engine's `articles_path` if it was rendered from the engine:

```
<%= link_to "Blog articles", articles_path %>
```

To make this route always use the engine's `articles_path` routing helper method, we must call the method on the routing proxy method that shares the same name as the engine.

```
<%= link_to "Blog articles", blorgh.articles_path %>
```

If you wish to reference the application inside the engine in a similar way, use the `main_app` helper:

```
<%= link_to "Home", main_app.root_path %>
```

If you were to use this inside an engine, it would **always** go to the application’s root. If you were to leave off the `main_app` “routing proxy” method call, it could potentially go to the engine’s or application’s root, depending on where it was called from.

If a template rendered from within an engine attempts to use one of the application’s routing helper methods, it may result in an undefined method call. If you encounter such an issue, ensure that you’re not attempting to call the application’s routing methods without the `main_app` prefix from within the engine.

6.4 Assets

Assets within an engine work in an identical way to a full application. Because the engine class inherits from `Rails::Engine`, the application will know to look up assets in the engine’s ‘app/assets’ and ‘lib/assets’ directories.

Like all of the other components of an engine, the assets should be namespaced. This means that if you have an asset called `style.css`, it should be placed at `app/assets/stylesheets/[engine name]/style.css`, rather than `app/assets/stylesheets/style.css`. If this asset isn’t namespaced, there is a possibility that the host application could have an asset named identically, in which case the application’s asset would take precedence and the engine’s one would be ignored.

Imagine that you did have an asset located at `app/assets/stylesheets/blorgh/style.css`. To include this asset inside an application, just use `stylesheet_link_tag` and reference the asset as if it were inside the engine:

```
<%= stylesheet_link_tag "blorgh/style.css" %>
```

You can also specify these assets as dependencies of other assets using Asset Pipeline require statements in processed files:

```
/*
 *= require blorgh/style
*/
```

Remember that in order to use languages like Sass or CoffeeScript, you should add the relevant library to your engine’s `.gemspec`.

6.5 Separate Assets & Precompiling

There are some situations where your engine’s assets are not required by the host application. For example, say that you’ve created an admin functionality that only exists for your engine. In this case, the host application doesn’t need to require `admin.css` or `admin.js`. Only the gem’s admin layout needs these assets. It doesn’t make sense for the host app to include `"blorgh/admin.css"` in its stylesheets. In this situation, you should explicitly define these assets for precompilation. This tells sprockets to add your engine assets when `rake assets:precompile` is triggered.

You can define assets for precompilation in `engine.rb`:

```
initializer "blorgh.assets.precompile" do |app|
  app.config.assets.precompile += %w(admin.css admin.js)
end
```

For more information, read the Asset Pipeline guide.

6.6 Other Gem Dependencies

Gem dependencies inside an engine should be specified inside the `.gemspec` file at the root of the engine. The reason is that the engine may be installed as a gem. If dependencies were to be specified inside the `Gemfile`, these would not be recognized by a traditional `gem install` and so they would not be installed, causing the engine to malfunction.

To specify a dependency that should be installed with the engine during a traditional `gem install`, specify it inside the `Gem::Specification` block inside the `.gemspec` file in the engine:

```
s.add_dependency "moo"
```

To specify a dependency that should only be installed as a development dependency of the application, specify it like this:

```
s.add_development_dependency "moo"
```

Both kinds of dependencies will be installed when `bundle install` is run inside of the application. The development dependencies for the gem will only be used when the tests for the engine are running.

Note that if you want to immediately require dependencies when the engine is required, you should require them before the engine's initialization. For example:

```
require 'other_engine/engine'
require 'yet_another_engine/engine'

module MyEngine
  class Engine < ::Rails::Engine
    end
end
```

7 Feedback

You're encouraged to help improve the quality of this guide.

Please contribute if you see any typos or factual errors. To get started, you can read our documentation contributions section.

You may also find incomplete content, or stuff that is not up to date. Please do add any missing documentation for master. Make sure to check Edge Guides first to verify if the issues are already fixed or not on the master branch. Check the Ruby on Rails Guides Guidelines for style and conventions.

If for whatever reason you spot something to fix but cannot patch it yourself, please open an issue.

And last but not least, any kind of discussion regarding Ruby on Rails documentation is very welcome in the `rubyonrails-docs` mailing list.
