

HTTP testing in R

Scott Chamberlain, Mlle Salmon

2021-04-02

Contents

1	Preamble	9
I	Introduction	11
2	HTTP in R 101	13
2.1	What is HTTP?	13
2.2	HTTP requests in R: what package?	14
3	Graceful HTTP R packages	19
3.1	Choose the HTTP resource wisely	19
3.2	User-facing grace (how your package actually works)	20
3.3	Graceful vignettes and examples	20
3.4	Graceful code	21
3.5	Graceful tests	21
3.6	Conclusion	21
4	Packages for HTTP testing	23
4.1	Why do we need special packages for HTTP testing?	23
4.2	webmockr	23
4.3	What is vcr?	24
4.4	What is httptest?	24
4.5	What is webfakes?	25
4.6	testthat	25
4.7	Conclusion	25

II	Whole Game(s)	27
5	Introduction	29
5.1	Our example package	29
6	Use vcr (& webmockr)	33
6.1	Setup	33
6.2	Actual testing	35
6.3	Also testing for real interactions	38
6.4	Summary	39
6.5	PS: Where to put use_cassette()	39
7	Use httptest	41
7.1	Setup	41
7.2	Actual testing	42
7.3	Also testing for real interactions	44
7.4	Summary	44
8	vcr and httptest	47
8.1	Setting up the infrastructure	47
8.2	Calling mock files	47
8.3	Naming mock files	48
8.4	Matching requests	48
8.5	Handling secrets	48
8.6	Recording, playing back	49
8.7	Testing for API errors	50
8.8	Conclusion	50
9	Use webfakes	51
9.1	Setup	51
9.2	Actual testing	52
9.3	Also testing for real interactions	55
9.4	Summary	55

<i>CONTENTS</i>	5
10 vcr (& webmockr), httptest, webfakes	57
10.1 What HTTP client can you use (curl, httr, crul)	57
10.2 Sustainability of the packages	57
10.3 Test writing experience	58
10.4 Test debugging experience	59
10.5 Conclusion	59
III Advanced Topics	61
11 Making real requests	63
11.1 What can change?	63
11.2 How to make real requests	63
11.3 Why not make only or too many real requests?	64
11.4 A complement to real requests: API news!	64
12 CRAN- (and Bioconductor) preparedness for your tests	67
12.1 Running tests on CRAN?	67
12.2 Skipping a few tests on CRAN?	68
12.3 Skipping all tests on CRAN?	68
12.4 Stress-test your package	68
13 Security	69
13.1 Managing secrets securely	69
13.2 Sensitive recorded responses?	73
13.3 Further resources	73
14 Faking HTTP errors	75
14.1 How to test for API errors (e.g. 503)	75
14.2 How to test for sequence of responses (e.g. 503 then 200)	75

15 Contributor friendliness	77
15.1 Taking notes about encryption	77
15.2 Providing a sandbox	77
15.3 Switching between accounts depending on the development mode	78
15.4 Documenting HTTP testing	78
IV Conclusion	79
16 Conclusion	81
V webmockr details	83
17 Mocking HTTP Requests	85
17.1 Package documentation	85
17.2 Features	85
17.3 How webmockr works in detail	85
17.4 Basic usage	87
18 stubs	89
18.1 Writing to disk	91
19 testing	93
20 utilities	95
20.1 Managing stubs	95
20.2 Managing stubs	95
20.3 Managing requests	95
VI vcr details	97
21 Caching HTTP requests	99
21.1 Package documentation	99
21.2 Terminology	99

<i>CONTENTS</i>	7
21.3 Basic usage	100
21.4 vcr enabled testing	102
22 Advanced vcr usage	105
22.1 Mocking writing to disk	105
23 Configure vcr	109
23.1 Set configuration variables	111
23.2 Re-set to defaults	112
23.3 Details on some of the config options	112
24 Record modes	121
24.1 once	121
24.2 none	121
24.3 new_episodes	122
24.4 all	122
25 Request matching	123
25.1 Matching	124
26 Debugging your tests that use vcr	125
26.1 An HTTP request has been made that vcr does not know how to handle	125
27 Security with vcr	129
27.1 Keeping secrets safe	129
27.2 API keys and tests run in varied contexts	130
27.3 Other security	131
28 Turning vcr on and off	133
28.1 turned_off	134
28.2 turn_off/turn_on	134
28.3 turned_on	135
28.4 Environment variables	136

29 Managing cassettes	137
29.1 Why edit cassettes?	137
29.2 Risks related to cassette editing	137
29.3 Example 1: test using an edited cassette with a 503	138
29.4 Example 2: test using an edited cassette with a 503 then a 200	139
29.5 gitignore cassettes	142
29.6 Rbuildignore cassettes	143
29.7 sharing cassettes	143
29.8 deleting cassettes	143
29.9 cassette file types	143
30 Gotchas	145
30.1 Correct line identification	146
31 Session info	147
31.1 Session info	147
31.2 Full session info	149

Chapter 1

Preamble

Are you working on a R package accessing resources on the web, be it a cat facts API, a scientific data source or your system for Customer relationship management? As for all other packages, appropriate unit testing can make your code more robust. The unit testing of a package interacting with web resources, however, brings special challenges: dependence of tests on a good internet connection, testing in the absence of authentication secrets, etc. Having tests fail due to resources being down or slow, during development or on CRAN, means a time loss for everyone involved (slower development, messages from CRAN). Although some packages accessing remote resources are well tested, there is a lack of resources around best practices.

This book is meant to be a free, central reference for developers of R packages accessing web resources, to help them have a faster and more robust development. Our aim is to develop an useful guidance to go with the great recent tools that `{vcr}`, `{webmockr}`, `{httptest}` and `{webfakes}` are.

We expect you to know [package development basics](#), and [git](#).

Note related to previous versions: this book was intended as a detailed guide to using a particular suite of packages for HTTP mocking and testing in R code and/or packages, namely those maintained by Scott Chamberlain (`{curl}`), `{webmockr}`, `{vcr}`) but its scope has been extended to generalize the explanation of concepts to similar packages.

You can also read the [PDF version](#) or [epub version](#) of this book.

Thanks to contributors to the book: [Alex Whan](#), [Aurèle](#), [Christophe Dervieux](#), [Daniel Posseriede](#), [Hugo Gruson](#), [Lluís Revilla Sancho](#), [Xavier A.](#)

Project funded by [rOpenSci](#) (Scott Chamberlain's work) & the [R Consortium](#) (Maëlle Salmon's work).

Part I

Introduction

Chapter 2

HTTP in R 101

2.1 What is HTTP?

HTTP means HyperText Transport Protocol, but you were probably not just looking for a translation of the abbreviation. HTTP is a way for you to exchange information with a remote server. In your package, if information is going back and forth between the R session and internet, you are using some sort of HTTP tooling. Your package is making *requests* and receives *responses*.

2.1.1 HTTP requests

The HTTP request is what your package makes. It has a method (are you fetching information via GET? are you sending information via POST?), different parts of an URL (domain, endpoint, query string), headers (containing e.g. your secret identifiers). It can contain a body, for instance you might be sending data as JSON. In that case one of the header will describe the content.

How do you know what request to make from your package? Hopefully you are interacting with a well documented web resource that will explain to you what methods are associated with what endpoints.

2.1.2 HTTP responses

The HTTP response is what the remote server provides, and what your package parses. A response has a status code indicating whether the request succeeded, response headers, and (optionally) a response body.

Hopefully the documentation of the web API or web resource you are working with shows good examples of responses. In any case you'll find yourself experimenting with different requests to see what the response "looks like".

2.1.3 More resources about HTTP

How do you get started with interacting with HTTP in R?

2.1.3.1 General HTTP resources

- [Mozilla Developer Network docs about HTTP](#) (recommended in the zine mentioned thereafter)
- (*not free*) [Julia Evans' Zine "HTTP: Learn your browser's language!"](#)
- The docs of the web API you are aiming to work with, and a search engine to understand the words that are new.

2.1.4 HTTP with R

- The docs of the R package you end up choosing!
- Digging into the source code of another package that does similar things.

2.2 HTTP requests in R: what package?

In R, to interact with web resources, it is recommended to use `{curl}`; or its higher-level interfaces `{httr}` (*pronounced [hitter](#) or [h-t-t-r](#)*) or `{crul}`.

Do not use `RCurl`, because it is not actively maintained!

When writing a package interacting with web resources, you will probably use either `{httr}` or `{crul}`.

- `httr` is the most popular and oldest of the two, and supports OAuth. `httr` docs feature a vignette called [Best practices for API packages](#)
- `crul` does not support OAuth but it uses an object-oriented interface, which you might like. `crul` has a set of [clients, or ways to perform requests](#), that might be handy. `crul` also has a vignette about [API package best practices](#).

Below we will try to programmatically access the [status of GitHub](#), the open-source platform provided by the company of the same name. We will access the same information with `httr` and `crul`. If you decide for the low-level `curl`, feel free to contribute an example.

```
github_url <- "https://kctbh9vrtwdw.statuspage.io/api/v2/status.json"
```

The URL above leaves no doubt as to what format the data is provided in, JSON!

Let's first use httr.

```
response <- httr::GET(github_url)
```

```
# Check the response status  
httr::http_status(response)
```

```
## $category  
## [1] "Success"  
##  
## $reason  
## [1] "OK"  
##  
## $message  
## [1] "Success: (200) OK"
```

```
# Or in a package you'd just write  
httr::stop_for_status(response)
```

```
# Parse the content  
httr::content(response)
```

```
## $page  
## $page$id  
## [1] "kctbh9vrtwdw"  
##  
## $page$name  
## [1] "GitHub"  
##  
## $page$url  
## [1] "https://www.githubstatus.com"  
##  
## $page$time_zone  
## [1] "Etc/UTC"  
##  
## $page$updated_at  
## [1] "2021-04-01T23:00:14.985Z"  
##  
##  
## $status
```

```
## $status$indicator
## [1] "none"
##
## $status$description
## [1] "All Systems Operational"
```

```
# In case you wonder, the format was obtained from a header
httr::headers(response)$`content-type`
```

```
## [1] "application/json; charset=utf-8"
```

Now, the same with `crul`.

```
# Create a client and get a response
client <- crul::HttpClient$new(github_url)
response <- client$get()
```

```
# Check the response status
response$status_http()
```

```
## <Status code: 200>
## Message: OK
## Explanation: Request fulfilled, document follows
```

```
# Or in a package you'd just write
response$raise_for_status()
```

```
# Parse the content
response$parse()
```

```
## No encoding supplied: defaulting to UTF-8.
```

```
## [1] "{\"page\":{\"id\":\"kctbh9vrtwd\"},\"name\":\"GitHub\", \"url\":\"https://www.g
```

```
jsonlite::fromJSON(response$parse())
```

```
## No encoding supplied: defaulting to UTF-8.
```

```
## $page
## $page$id
## [1] "kctbh9vrtwd"
```



```
##
## $page$name
## [1] "GitHub"
##
## $page$url
## [1] "https://www.githubstatus.com"
##
## $page$time_zone
## [1] "Etc/UTC"
##
## $page$updated_at
## [1] "2021-04-01T23:00:14.985Z"
##
##
## $status
## $status$indicator
## [1] "none"
##
## $status$description
## [1] "All Systems Operational"
```

Hopefully these very short snippets give you an idea of what syntax to expect when choosing one of those packages.

Note that the choice of a package will constrain the HTTP testing tools you can use. However, the general ideas will remain the same. You could switch your package backend from say `crul` to `httr` *without changing your tests*, if your tests do not test too many specificities of internals.

Chapter 3

Graceful HTTP R packages

Based on the previous chapter, your package interacting with a web resource has a dependency on `{curl}`, `{httr}` or `{crul}`. You have hopefully read the docs of the dependency you chose, including, in the case of `httr` and `crul`, the vignette about best practice for HTTP packages. Now, in this chapter we want to give more tips aimed at making your HTTP R package graceful, part of which you'll learn more about in this very book!

Why write a *graceful* HTTP R package? First of all, *graceful* is a nice adjective. ☹️ Then, *graceful* is the adjective used in [CRAN repository policy](#) “*Packages which use Internet resources should fail gracefully with an informative message if the resource is not available or has changed (and not give a check warning nor error).*” Therefore, let's review how to make your R package graceful from this day forward, in success and in failure.

3.1 Choose the HTTP resource wisely

First of all, your life and the life of your package's users will be easier if the web service you're wrapping is well maintained and well documented. When you have a choice, try not to rely on a fragile web service. Moreover, if you can, try to communicate with the API providers (telling them about your package; reporting feature requests and bug reports in their preferred way).

3.2 User-facing grace (how your package actually works)

0. If you can, do not request the API every time the user asks for something but cache data instead. No API call, no API call failure! ☒ To remember answers within a session check out [memoise](#). To remember answers across sessions, see [webmiddens](#), and approaches presented in the R-hub blog post “[Persistent config and data for R packages](#)”. Caching behavior should be well documented for users, and there should probably be an expiration time for caches that’s based on how often data is updated on the remote service.
1. Try to send correct requests by knowing what the API expects and validating user inputs; at the correct rate.
 - For instance, don’t even try interacting with a web API requiring authentication if the user does not provide authentication information.
 - For limiting rate i.e. not sending too many requests, automatically wait or, if the API docs allow you to define an ideal or maximal rate, set the request rate in advance using the [ratelimitr](#) package.
2. If there’s a status API i.e. a separate API indicating whether the web resource is up or down, use it. If it tells you the API is down, `stop()` with an informative error message.
3. If the API indicates an error, depending on the actual error,
 - If the *server* seems to be having issues, [re-try with an exponential back-off](#).
 - Otherwise, [transform the error into an useful error](#).
 - If you used retry and nothing was sent after the maximal number of re-tries, have an informative error message.

That was it for aspects the user will care about. Now, what might be more problematic for your package’s fate on CRAN are the automatic checks that happen there at submission and then [regularly](#).

3.3 Graceful vignettes and examples

4. [Pre-compute vignettes](#) in some way. Don’t use them as tests, they are a showcase. Of course have a system to prevent them from going stale, maybe even simple reminders (potentially in the [unexported `release_questions\(\)` function](#)). Don’t let vignettes run on a system where a failure has bad consequences.

5. Don't run [examples](#) on CRAN. Now, for a first submission, CRAN maintainers might complain if there is no example. In that case, you might want to add some minimal example, e.g.

```
if (crul::ok("some-url")) {  
  foo_bar() # some eg that uses some-url  
}
```

These two precautions ensure that CRAN checks won't end with some WARNINGS e.g. because an example failed when the API was down.

3.4 Graceful code

For simplifying your own life and those of contributors, make sure to re-use code in your package by e.g. defining helper functions for making requests, handling responses etc. It will make it easier for you to support interactions with more parts of the web API. Writing DRY (don't repeat yourself) code means less lines of code to test, less API calls to make or fake!

Also, were you to export a function à la `gh : gh()`, you'll help users call any endpoint of the web API even if you haven't written any high-level helper for it yet.

3.5 Graceful tests

We're getting closer to the actual topic of this book!

6. Read the rest of this book! Your tests should ideally run without needing an actual internet connection nor the API being up. Your tests that do need to interact with the API should be skipped on CRAN. `testthat::skip_on_cran()` will ensure that.
7. Do not only test success behavior! Test for the behavior of your package in case of API errors, which shall also be covered later in the book.

3.6 Conclusion

In summary, to have a graceful HTTP package, make the most of current best practice for the user interface; escape examples and vignettes on CRAN; make tests independent of actual HTTP requests. Do not forget CRAN's "graceful failure" policy is mostly about ensuring a clean R CMD check result on CRAN platforms (0 ERROR, 0 WARNING, 0 NOTE) even when the web service you're wrapping has some hiccups.

Chapter 4

Packages for HTTP testing

A brief presentation of packages you'll "meet" again later in this book!

4.1 Why do we need special packages for HTTP testing?

Packages for HTTP testing are useful because there are challenges to HTTP testing. Packages for HTTP testing help you solve these challenges, rather than letting you solve them with some homegrown solutions (you can still choose to do that, of course).

What are the challenges of HTTP testing?

- Having tests depend on an internet connection is not ideal.
- Having tests depend on having secrets for authentication at hand is not ideal.
- Having tests for situations that are hard to trigger (e.g. the failure of a remote server) is tricky.

4.2 webmockr

`{webmockr}`, maintained by Scott Chamberlain, is an R package to help you "mock" HTTP requests. What does mock mean? Mock refers to the fact that we're faking the response. Here is how it works:

- You "stub" a request. That is, you set rules for what HTTP request you'd like to respond to with a fake. E.g. a rule might be a method, or an URL.

- You also can set rules for what fake you'd like to respond with, if anything (if nothing, then we give you `NULL`).
- Then you make HTTP requests, and those that match your stub i.e. set of rules will return what you requested be returned.
- While webmockr is in use, real HTTP interactions are not allowed. Therefore you need to stub all possible HTTP requests happening via your code. You'll get error messages for HTTP requests not covered by any stub.
- There is no recording interactions to disk at all, just mocked responses given as the user specifies in the R session.

webmockr works with both the `crul` package and the `httr` package.

webmockr is quite low-level and not the first tool you'll use directly in your day-to-day HTTP testing. You may never use it directly but if you use `vcr` it's one of its foundations.

webmockr was inspired by [Ruby webmock gem](#).

4.3 What is vcr?

The short version is `{vcr}`, maintained by Scott Chamberlain, helps you stub HTTP requests so you don't have to repeat HTTP requests, mostly in your unit tests. It uses the power of webmockr, with a higher level interface.

When using `vcr` in tests, the first time you run a test, the API response is stored in a YAML file. All subsequent runs of the test use that local file instead of really calling the API. Therefore tests work independently of an internet connection.

`vcr` was inspired by [Ruby vcr gem](#).

`vcr` works for packages using `httr` or `crul`.

[Direct link to vcr \(& webmockr\) demo](#)

4.4 What is httptest?

`{httptest}` maintained by Neal Richardson, like `{vcr}`, uses mocked API responses. It *“enables one to test all of the logic on the R sides of the API in your package without requiring access to the remote service.”*

Contrary to `{vcr}`, `{httptest}` also lets you define mock files by hand (copying from API docs, or dumbing down real responses), whereas with `vcr` all mock files come from recording real interactions (although you can choose to [edit vcr mock files](#) after recording).

`httptest` works for packages using `httr`.

[Direct link to httptest demo](#)

The differences and similarities between `httptest` and `vcr` will become clearer in the chapters where we provide the whole games for both of them.

With both `vcr` and `httptest` the tests will use some sort of fake API responses.

In `vcr` they are called **fixtures** or **cassettes**. In `httptest` they are called **mock files**.

4.5 What is webfakes?

`webfakes`, maintained by Gábor Csárdi, provides an alternative (complementary?) tool for HTTP testing. It will let you fake a whole web service, potentially outputting responses from mock files you'll have created. It does not help with recording fake responses. Because it runs a fake web service, you can even interact with said web service in your browser or with `curl` in the command line.

`webfakes` works with packages using any HTTP package i.e. it works with `curl`, `crul`, `httr`.

[Direct link to webfakes demo](#)

4.6 testthat

`testthat` maintained by Hadley Wickham is not a package specifically for HTTP testing, it is a package for general purpose unit testing of R packages. In this book we will assume that is what you use, because of its popularity.

Now, if you use an alternative like `tinytest`,

- `httptest` won't work as it's specifically designed as a complement to `testthat`;
- `vcr` might work;
- `webfakes` can work.

4.7 Conclusion

Now that you have an idea of the tools we can use for HTTP testing, we'll now create a minimal package and then amend it in three versions tested with

- `{vcr}` and `{webmockr}`;
- `{httptest}`;

- {webfakes}.

Our minimal package will use `httr`. However, it will help you understand concepts even if you end up using `crul` or `curl`.¹

¹If you end up using `crul`, you can use `vcr` and `webmockr`; or `webfakes`; but not `httpptest`. If you end up using `curl` you can only use `webfakes`.

Part II

Whole Game(s)

Chapter 5

Introduction

Similar to the [Whole Game chapter in the R packages book by Hadley Wickham and Jenny Bryan](#), we shall go through how to add HTTP tests to a minimal package. However, we will do it *three* times to present alternative approaches: with [vcr](#), [httptest](#), [webfakes](#). After that exercise, we shall compare approaches: we will compare both packages that involve mocking i.e. [vcr vs. httptest](#); and all three HTTP packages in a [last chapter](#). The next section will then present single topics such as “how to deal with authentication” in further details.

5.1 Our example package

Our minimal package, [exempligratia](#), accesses the GitHub status API and one endpoint of GitHub V3 REST API. It is named after the Latin phrase *exempli gratia* that means “for instance”, with an H for GH. If you really need to interact with GitHub V3 API, we recommend the [gh](#) package. We also recommend looking at the source of the [gh](#) package, and at the docs of GitHub V3 API, in particular about [authentication](#).

Our example package calls web *APIs* but the tools and concepts are applicable to packages wrapping any web resource, even poorly documented ones.¹

GitHub V3 API works without authentication too, but at a lower rate. For the sake of having an example of a package *requiring* authentication we shall assume the API is *not* usable without authentication. Authentication is, here, the setting of a token in a HTTP header (so quite simple, compared to e.g. OAuth).

GitHub Status API, on the contrary, does not necessitate authentication at all.

So we shall create two functions, one that works without authentication, one that works with authentication.

¹An interesting post to read about an R package wrapping an undocumented web API is “[One-Hour Package](#)” by Neal Richardson.

How did we create the package? You are obviously free to use your own favorite workflow tools, but below we share our workflow using the [usethis package](#).

- We followed [usethis setup article](#).

Then we ran

- `usethis::create_package("path/to/folder/exemplighratia")` to create and open the package project;
- `usethis::use_mit_license()` to add an MIT license;
- `usethis::use_package("httr")` to add a dependency on `httr`;
- `usethis::use_package("purrr")` to add a dependency on `purrr`;
- `use_r("api-status.R")` to add the first function whose code is written below;

```
status_url <- function() {
  "https://kctbh9vrtwdw.statuspage.io/api/v2/components.json"
}

#' GitHub APIs status
#'
#' @description Get the status of requests to GitHub APIs
#'
#' @return A character vector, one of "operational", "degraded_performance",
#' "partial_outage", or "major_outage."
#'
#' @details See details in https://www.githubstatus.com/api#components.
#' @export
#'
#' @examples
#' \dontrun{
#' gh_api_status()
#' }
gh_api_status <- function() {
  response <- httr::GET(status_url())

  # Check status
  httr::stop_for_status(response)

  # Parse the content
  content <- httr::content(response)

  # Extract the part about the API status
  components <- content$components
  api_status <- components[purrr::map_chr(components, "name") == "API Requests"][[1]]
}
```

```
# Return status
api_status$status
}
```

- `use_test("api-status")` (and using [testthat latest edition](#) so setting `Config/testthat/edition: 3` in DESCRIPTION) to add a simple test whose code is below.

```
test_that("gh_api_status() works", {
  testthat::expect_type(gh_api_status(), "character")
})
```

- `use_r("organizations.R")` to add a second function. Note that an ideal version of this function would have some sort of callback in the retry, to call the `gh_api_status()` function (which seems easier to implement with [crul's retry method](#)).

```
gh_v3_url <- function() {
  "https://api.github.com/"
}

#' GitHub organizations
#'
#' @description Get logins of GitHub organizations.
#'
#' @param since The integer ID of the last organization that you've seen.
#'
#' @return A character vector of at most 30 elements.
#' @export
#'
#' @details Refer to https://developer.github.com/v3/orgs/#list-organizations
#'
#' @examples
#' \dontrun{
#' gh_organizations(since = 42)
#' }
gh_organizations <- function(since = 1) {
  url <- httr::modify_url(
    gh_v3_url(),
    path = "organizations",
    query = list(since = since)
  )
```

```
token <- Sys.getenv("GITHUB_PAT")

if (!nchar(token)) {
  stop("No token provided! Set up the GITHUB_PAT environment variable please.")
}

response <- httr::RETRY(
  "GET",
  url,
  httr::add_headers("Authorization" = paste("token", token))
)

httr::stop_for_status(response)

content <- httr::content(response)

purrr::map_chr(content, "login")
}
```

- `use_test("organizations")` to add a simple test.

```
test_that("gh_organizations works", {
  testthat::expect_type(gh_organizations(), "character")
})
```

All good, now our package has 100% test coverage and passes R CMD Check (granted, our tests could be more thorough, but remember this is a minimal example). But what if we try working without a connection? In the following chapters, we'll add more robust testing infrastructure to this minimal package, and we will do that *three* times to compare packages/approaches: once with `vcr`, once with `httptest`, and once with `webfakes`.

Chapter 6

Use vcr (& webmockr)

In this chapter we aim at adding HTTP testing infrastructure to `exemplighratia` using `vcr` (& `webmockr`).

[Corresponding pull request to exemplighratia](#). Feel free to fork the repository to experiment yourself!

6.1 Setup

Before working on all this, we need to install `{vcr}`.

First, we need to run `vcr::use_vcr()` (in the `exemplighratia` directory) which has a few effects:

- Adding `vcr` as a dependency to `DESCRIPTION`, under `Suggests` just like `testthat`.
- Creating an example test file for us to look at. This is useful the first few times you setup `vcr` in another package, after a while you might even delete it without reading it.
- Adding a `.gitattributes` file with the line `tests/fixtures/**/* -diff` which will hide the changes in your cassettes from the git diff. It makes your git diff easier to deal with.¹
- Creating a setup file under `tests/testthat/setup-exemplighratia`,

```
library("vcr")
invisible(vcr::vcr_configure(
  dir = vcr::vcr_test_path("fixtures")
))
vcr::check_cassette_names()
```

¹However, if you change something related to handling secrets in your code or tests, please check again your new cassettes do not include secrets.

When testthat runs tests, files whose name starts with “setup” are always run first. The setup file created by vcr

- loads vcr,
- indicates where mocked responses are saved ("./fixtures" which translates, from the root of the package, to tests/fixtures),
- and checks that you are not using the same name twice for cassettes (mock files).

We have to tweak vcr setup a bit for our needs.

- We do not want our API token to appear in the mock responses, and we know it's used in the Authorization header of the requests, so we use the `filter_request_headers` argument of `vcr::vcr_configure()`. For other secret filtering one can use `filter_response_headers` and `filter_sensitive_data` (for a regular expression purging, on the whole saved interactions).
- We need to ensure that we set up a fake API key when there is no API token around. Why? Because if you remember well, the code of our function `gh_organizations()` checks for the presence of a token. With mock responses around, we don't need a token but we still need to fool our own package in contexts where there is no token (e.g. in continuous integration checks for a fork of a GitHub repository).

Below is the updated setup file saved under `tests/testthat/setup-exemplighratia`.

```
library("vcr")

vcr_dir <- vcr::vcr_test_path("fixtures")

if (!nzchar(Sys.getenv("GITHUB_PAT"))) {
  if (dir.exists(vcr_dir)) {
    # Fake API token to fool our package
    Sys.setenv("GITHUB_PAT" = "foobar")
  } else {
    # If there's no mock files nor API token, impossible to run tests
    stop("No API key nor cassettes, tests cannot be run.",
        call. = FALSE)
  }
}

invisible(vcr::vcr_configure(
  dir = vcr_dir,
  # Filter the request header where the token is sent, make sure you know
```

```
# how authentication works in your case and read the Security chapter :-)
filter_request_headers = list(Authorization = "My bearer token is safe")
))
```

So this was just setup, now on to adapting our tests!

6.2 Actual testing

The most important function will be `vcr::use_cassette("cassette-informative-and-unique-name", {code-block})` which tells vcr to create a mock file to store all API responses for API calls occurring in the code block.

Let's tweak the test for `gh_api_status`, it now becomes

```
test_that("gh_api_status() works", {
  vcr::use_cassette("gh_api_status", {
    status <- gh_api_status()
  })
  testthat::expect_type(status, "character")
})
```

We only had to wrap the code involving interactions with the API, `status <- gh_api_status()`, in `vcr::use_cassette()`.

If we run this test (in RStudio clicking on "Run test"),

- the first time, vcr creates a cassette (mock file) under `tests/testthat/fixtures/gh_api_status.yml` where it stores the API response. It contains all the information related to requests and responses, headers included.

```
http_interactions:
- request:
  method: get
  uri: https://kctbh9vrtwdw.statuspage.io/api/v2/components.json
  body:
    encoding: ''
    string: ''
  headers:
    Accept: application/json, text/xml, application/xml, */*
  response:
    status:
      status_code: 200
      category: Success
```

```

reason: OK
message: 'Success: (200) OK'
headers:
  vary: Accept,Accept-Encoding,Fastly-SSL
  cache-control: max-age=0, private, must-revalidate
  x-cache: MISS
  content-type: application/json; charset=utf-8
  content-encoding: gzip
  strict-transport-security: max-age=259200
  date: Thu, 15 Oct 2020 11:59:23 GMT
  x-request-id: d9888435-3f04-4401-be5c-b9d1bfdfa015
  x-download-options: noopen
  x-xss-protection: 1; mode=block
  x-runtime: '0.037254'
  x-permitted-cross-domain-policies: none
  access-control-allow-origin: '*'
  accept-ranges: bytes
  x-content-type-options: nosniff
  etag: W/"gz[a479c9894f51b7db286dc31cd922e7bf]"
  x-statuspage-skip-logging: 'true'
  x-statuspage-version: fd137a4bb14c20ce721393e5b6540ea6eebff3a3
  referrer-policy: strict-origin-when-cross-origin
  age: '0'
body:
  encoding: UTF-8
  file: no
  string: '{"page":{"id":"kctbh9vrtdwd","name":"GitHub","url":"https://www.githubs
    Operations","status":"operational","created_at":"2017-01-31T20:05:05.370Z","up
    of git clones, pulls, pushes, and associated operations","showcase":false,"sta
    Requests","status":"operational","created_at":"2017-01-31T20:01:46.621Z","upda
    for GitHub APIs","showcase":false,"start_date":null,"group_id":null,"page_id":
    time HTTP callbacks of user-generated and system events","showcase":false,"sta
    www.githubstatus.com for more information","status":"operational","created_at"
    for Issues on GitHub.com","showcase":false,"start_date":null,"group_id":null,"
    Requests","status":"operational","created_at":"2020-09-02T15:39:06.329Z","upda
    for Pull Requests on GitHub.com","showcase":false,"start_date":null,"group_id"
    Actions","status":"operational","created_at":"2019-11-13T18:02:19.432Z","updat
    Compute and Orchestration for GitHub Actions","showcase":false,"start_date":nu
    Packages","status":"operational","created_at":"2019-11-13T18:02:40.064Z","upda
    requests and webhook delivery for GitHub Packages","showcase":false,"start_dat
    Pages","status":"operational","created_at":"2017-01-31T20:04:33.923Z","updated
    application and API servers for Pages builds","showcase":false,"start_date":nu
  recorded_at: 2020-10-15 11:59:23 GMT
  recorded_with: vcr/0.5.4, webmockr/0.7.0

```

- all the times after that, unless we delete the mock file, vcr simply uses the mock files instead of actually calling the API.

Let's tweak our other test, of `gh_organizations()`. Here things get more exciting or complicated, as we also set out to adding a test of the error behavior. This inspired us to change error behavior a bit with a slightly more specific error message i.e. `httr::stop_for_status(response)` became `httr::stop_for_status(response, task = "get data from the API, oops")`.

The test file `tests/testthat/test-organizations.R` is now:

```
test_that("gh_organizations works", {
  vcr::use_cassette("gh_organizations", {
    orgs <- gh_organizations()
  })
  testthat::expect_type(orgs, "character")
})

test_that("gh_organizations errors when the API doesn't behave", {
  webmockr::enable()
  stub <- webmockr::stub_request("get", "https://api.github.com/organizations?since=1")
  webmockr::to_return(stub, status = 502)
  expect_error(gh_organizations(), "oops")
  webmockr::disable()
})
```

The first test is similar to what we did for `gh_api_status()`. In the second test there is more to unpack.

- We enable the use of `{webmockr}` at the beginning with `webmockr::enable()`. Why `webmockr`? Because it can help mock a failure scenario.
- We explicitly write that a request to `https://api.github.com/organizations?since=1` should return a status of 502.

```
stub <- webmockr::stub_request("get", "https://api.github.com/organizations?since=1")
webmockr::to_return(stub, status = 502)
```

- We then test for the error message with `expect_error(gh_organizations(), "oops")`.
- We disable `webmockr` with `webmockr::disable()`.

Instead of using `webmockr` for creating a fake API error, we could have

- recorded a normal cassette;

- edited it to replace the status code.

Read pros and cons of this approach in the vcr vignette [Why and how edit your vcr cassettes?](#), especially if you don't find the webmockr approach enjoyable.

Without the HTTP testing infrastructure, testing for behavior of the package in case of API errors would be more difficult.

Regarding our secret API token, the first time we run the test file, vcr creates a cassette where we notice these lines

```
http_interactions:
- request:
  method: get
  uri: https://api.github.com/organizations?since=1
  body:
    encoding: ''
    string: ''
  headers:
    Accept: application/json, text/xml, application/xml, */*
    Content-Type: ''
    Authorization: My bearer token is safe
```

Our API token has been replaced with the string we indicated in `vcr::vcr_configure()`, `My bearer token is safe`.

6.3 Also testing for real interactions

What if the API responses change? Hopefully we'd notice that thanks to following API news. However, sometimes web APIs change without any notice. Therefore it is important to run tests against the real web service once in a while.

The vcr package provides various methods to turn vcr use on and off to allow real requests i.e. ignoring mock files. See `?vcr::lightswitch`.

In the case of *exempligratia*, we added a [GitHub Actions workflow](#) that will run on schedule once a week, for which one of the build has vcr turned off via the `VCR_TURN_OFF` environment variable. We chose to have one build with vcr turned on and otherwise the same configuration to make it easier to assess what broke in case of failure (if both builds fail, the web API is probably not the culprit). Compared to continuous integration builds where vcr is turned on, this one build needs to have access to a `GITHUB_PAT` secret environment variable. Furthermore, it is slower.

One could imagine other strategies:

- Always having one continuous integration build with `vcr` turned off but skipping it in contexts where there isn't any token (pull requests from forks for instance?);
- Only running tests with `vcr` turned off locally once in a while.

6.4 Summary

- We set up `vcr` usage in our package `exempligratia` by running `use_vcr()` and tweaking the setup file to protect our secret API key and to fool our own package that needs an API token.
- Inside `test_that()` blocks, we wrapped parts of the code into `vcr::use_cassette()` and ran the tests a first time to generate mock files that hold all information about the API interactions.
- In one of the tests, we used `webmockr` to create an environment where only fake requests are allowed. We defined that the request that `gh_organizations()` makes should get a 502 status. We were therefore able to test for the error message `gh_organizations()` returns in such cases.

Now, how do we make sure this works?

- Turn off wifi, run the tests again. It works! Turn on wifi again.
- Open `.Renviron` (`usethis::edit_r_environ()`), edit "GITHUB_PAT" into "byeGITHUB_PAT", re-start R, run the tests again. It works! Fix your "GITHUB_PAT" token in `.Renviron`.

So we now have tests that no longer rely on an internet connection nor on having API credentials.

We also added a continuous integration workflow for having a build using real interactions once every week, as it is important to regularly make sure the package still works against the latest API responses.

For the full list of changes applied to `exempligratia` in this chapter, see [the pull request diff on GitHub](#).

How do we get there with other packages? Let's try `httptest` in the next chapter!

6.5 PS: Where to put `use_cassette()`

Where do we put the `vcr::use_cassette()` call? Well, as written in the manual page of that function, *There's a few ways to get correct line numbers for failed tests and one way to not get correct line numbers*: What's correct?

- Wrapping the whole `testthat::test_that()` call;

```
vcr::use_cassette("thing", {
  testthat::test_that("thing", {
    lala <- get_foo()
    expect_true(lala)
  })
})
```

- Wrapping a few lines inside `testthat::test_that()` **excluding the expectations** `expect_blabla()`

```
testthat::test_that("thing", {
  vcr::use_cassette("thing", {
    lala <- get_foo()
  })
  expect_true(lala)
})
```

What's incorrect?

```
testthat::test_that("thing", {
  vcr::use_cassette("thing", {
    lala <- get_foo()
    expect_true(lala)
  })
})
```

We used the solution of only wrapping the lines containing API calls in `vcr::use_cassette()`, but it is up to you to choose what you prefer.

Chapter 7

Use httptest

In this chapter we aim at adding HTTP testing infrastructure to `exempligratia` using `httptest`. For this, we start from the initial state of `exempligratia` again. Back to square one!

Note that the `httptest::with_mock_dir()` function is only available in `httptest` version $\geq 4.0.0$ (released on CRAN on 2021-02-01).

[Corresponding pull request to `exempligratia`](#) Feel free to fork the repository to experiment yourself!

7.1 Setup

Before working on all this, we need to install `{httptest}`.

First, we need to run `httptest::use_httptest()` which has a few effects:

- Adding `httptest` as a dependency to `DESCRIPTION`, under `Suggests` just like `testthat`.
- Creating a setup file under `tests/testthat/setup`,

```
library(httptest)
```

When `testthat` runs tests, [files whose name starts with “setup”](#) are always run first. The setup file added by `httptest` loads `httptest`.

We shall tweak it a bit to fool our package into believing there is an API token around in contexts where there is not. Since tests will use recorded responses when we are not recording, we do not need an actual API token when not recording, but we need `gh_organizations()` to not stop because `Sys.getenv("GITHUB_PAT")` returns nothing.

```
library(httptest)

# for contexts where the package needs to be fooled
# (CRAN, forks)
# this is ok because the package will use recorded responses
# so no need for a real secret
if (!nzchar(Sys.getenv("GITHUB_PAT"))) {
  Sys.setenv(GITHUB_PAT = "foobar")
}
```

So this was just setup, now on to adapting our tests!

7.2 Actual testing

The key function will be `httptest::with_mock_dir("dir", {code-block})` which tells `httptest` to create mock files under `tests/testthat/dir` to store all API responses for API calls occurring in the code block. We are allowed to tweak the mock files by hand, and we will do that in some cases.

Let's tweak the test file for `gh_status_api`, it becomes

```
with_mock_dir("gh_api_status", {
  test_that("gh_api_status() works", {
    testthat::expect_type(gh_api_status(), "character")
    testthat::expect_equal(gh_api_status(), "operational")
  })
})
```

We only had to wrap the whole test in `httptest::with_mock_dir()`.

If we run this test (in RStudio clicking on “Run test”),

- the first time, `httptest` creates a mock file under `tests/testthat/gh_api_status/kctbh9vrtdwd.sta` where it stores the API response. We however dumbed it down by hand, to

```
{"components": [{"name": "API Requests", "status": "operational"}]}
```

- all the times after that, `httptest` simply uses the mock file instead of actually calling the API.

Let's tweak our other test, of `gh_organizations()`.

Here things get more exciting or complicated, as we also set out to adding a test of the error behavior. This inspired us to change error behavior a bit with a slightly

more specific error message i.e. `httr::stop_for_status(response)` became `httr::stop_for_status(response, task = "get data from the API, oops")`.

The test file `tests/testthat/test-organizations.R` is now:

```
with_mock_dir("gh_organizations", {
  test_that("gh_organizations works", {
    testthat::expect_type(gh_organizations(), "character")
  })
})

with_mock_dir("gh_organizations_error", {
  test_that("gh_organizations errors if the API doesn't behave", {
    testthat::expect_error(gh_organizations())
  })
},
simplify = FALSE)
```

The first test is similar to what we did for `gh_api_status()` except we didn't touch the mock file this time, out of laziness. In the second test there is more to unpack: how do we get a mock file corresponding to an error?

- We first run the test as is. It fails because there is no error, which we expected. Note the `simplify = FALSE` that means the mock file also contains headers for the response.
- We replaced 200L with 502L and removed the body, to end up with a very simple mock file under `tests/testthat/gh_organizations_error/api.github.com/organizations-5377e8.R`.

```
structure(list(url = "https://api.github.com/organizations?since=1",
  status_code = 502L, headers = NULL), class = "response")
```

- We re-run the tests. We got the expected error message.

Without the HTTP testing infrastructure, testing for behavior of the package in case of API errors would be more difficult.

Regarding our secret API token, since `httptest` doesn't save the requests, and since the responses don't contain the token, it is safe without our making any effort.

In this demo we used `httptest::with_mock_dir()` but there are other ways to use `httptest`, e.g. using `httptest::with_mock_api()` that does not require naming a directory (you'd still need to use a separate directory for mocking the error response).

Find out more in the [main `httptest` vignette](#).

7.3 Also testing for real interactions

What if the API responses change? Hopefully we'd notice that thanks to following API news. However, sometimes web APIs change without any notice. Therefore it is important to run tests against the real web service once in a while.

As with `vcr` we setup a [GitHub Actions workflow](#) that runs once a week with tests against the real web service. The difference is what and where these tests are. As some tests with custom made mock files can be more specific (e.g. testing for actual values, whereas the latest responses from the API will have different values), instead of turning off mock files usage, we use our old original tests that we put in a folder called `real-tests`. Most of the time `real-tests` is `.Rbuildignore`d but in the scheduled run, before checking the package we replace the content of `tests` with `real-tests`. An alternative would be to use `testthat::test_dir()` on that directory but in case of failures we would not get artifacts as we do with `R CMD check` (at least not without further effort).

Again, one could imagine other strategies, but in all cases it is important to keep checking the package against the real web service fairly regularly.

7.4 Summary

- We set up `httpptest` usage in our package `exemplighratia` by running `use_httpptest()` and tweaking the setup file to fool our own package that needs an API token.
- We wrapped `test_that()` into `httpptest::with_mock_dir()` and ran the tests a first time to generate mock files that hold all information about the API responses. In some cases we modified these mock files to make them smaller or to make them correspond to an API error.

Now, how do we make sure this works?

- Turn off wifi, run the tests again. It works! Turn on wifi again.
- Open `.Renv` (`usethis::edit_r_env()`), edit `"GITHUB_PAT"` into `"byeGITHUB_PAT"`, re-start R, run the tests again. It works! Fix your `"GITHUB_PAT"` token in `.Renv`.

So we now have tests that no longer rely on an internet connection nor on having API credentials.

We also added a continuous integration workflow for having a build using real interactions once every week, as it is important to regularly make sure the package still works against the latest API responses.

For the full list of changes applied to `exempligratia` in this chapter, see [the pull request diff on GitHub](#).

How do we get there with yet another package? We'll try `webfakes` but first let's compare `vcr` and `httptest` as they both use mocking.

Chapter 8

vcr and httptest

We have just followed very similar processes to add HTTP testing infrastructure involving mock files to exempligratia

- Adding a package as a Suggests dependency;
- Creating a helper file that in particular loads this package before each test;
- Tweaking tests, in some cases wrapping our tests into functions that allows to record API responses in mock files and to play them back from said mock files; in other cases (only with httptest), creating mock files ourselves.

Now, there were a few differences. We won't end up advocating for one package in particular since both have their merits, but we do hope to help you differentiate the two packages.

8.1 Setting up the infrastructure

To set up the HTTP testing infrastructure, in one case you need to run `vcr::use_vcr()` and in another case you need to run `httptest::use_httptest()`. Not too hard to remember.

8.2 Calling mock files

As mentioned before, vcr and httptest both use mock files but they call them differently.

In vcr they are called both **fixtures** and **cassettes**. In httptest they are called **mock files**. Note that fixtures is not as specific as cassettes and mock files: cassettes and

mock files are fixtures, but anything (a csv file of input for instance) you use to consistently test your package is a fixture.

8.3 Naming mock files

With `vcr` the `use_cassette()` call needs to include a name that will be used to create the filename of the mock file. The help of `?use_cassette` explains some criteria for naming them, such as the fact that cassette names need to be unique. Now if you wrap your whole `test_that()` block in them you might just as well use a name similar to the test name, and you already make those meaningful, right?

With `httptest` the mock filepaths are translated from requests according to several rules that incorporate the request method, URL, query parameters, and body. If you use `with_mock_dir()` you need a name for the directory under which the mock files are saved, and you can make it meaningful.

Also note that with `vcr` one file can (but does not have to) contain several HTTP interactions (requests and responses) whereas with `httptest` one file contains one response only (and the filename helps matching it to a request).

8.4 Matching requests

With `httptest` as the mock file name includes everything that's potentially varying about a request, each mock file corresponds to one request only.

With `vcr`, there are different possible [configurations for matching a request to a saved interaction](#) but by default you can mostly expect that one saved interaction corresponds to one request only.

8.5 Handling secrets

With `vcr`, since everything from the HTTP interactions is recorded, you always need to add some sort of configuration to be sure to wipe your API tokens from the mock files.

With `httptest`, only responses are saved, and most often, only their bodies. Most often, responses don't contain secrets e.g. they don't contain your API token. If the response contains secrets, refer to `httptest`'s article about "[Redacting sensitive information](#)".

8.6 Recording, playing back

When using mock files for testing, first you need to record responses in mock files; and then you want to use the mock files instead of real HTTP interactions (that's the whole point).

With `vcr`, the recording vs playing back modes happen automatically depending on the existence of the cassette. If you write `vcr : use_cassette("blabla",)` and there's no cassette called `blabla`, `vcr` will create it. Note that if you change the HTTP interactions in the code block, you'll have to re-record the cassette which is as simple as deleting it then running the test. *Note that you can also change the way `vcr` behaves by looking into `?vcr : vcr_configure`'s "Cassette Options".*

With `httptest`, there is a lot of flexibility around how to record mock files. It is because `httptest` doesn't assume that every API mock came from a real request to a real server; maybe you copy some of the mocks directly from the API docs.

Note that nothing prevents you from editing `vcr` cassettes by hand, but you'll have to be careful not re-recording them by mistake.

`httptest` flexibility comes from [original design principles of `httptest`](#)

"[httptest] doesn't assume that every API mock came from a real request to a real server, and it is designed so that you are able to see and modify test fixtures. Among the considerations:

- 1. In many cases, API responses contain way more content than is necessary to test your R code around them: 100 records when 2 will suffice, request metadata that you don't care about and can't meaningfully assert things about, and so on. In the interest of minimally reproducible examples, and of making tests readable, it often makes sense to take an actual API response and delete a lot of its content, or even to fabricate one entirely.*
- 2. And then it's good to keep that API mock fixed so you know exactly what is in it. If I re-recorded a Twitter API response of, say, the most recent 10 tweets with `#rstats`, the specific content will change every time I record it, so my tests can't say much about what is in the response without having to rewrite them every time too.*
- 3. Some conditions (rate limiting, server errors, e.g.) are difficult to test with real responses, but if you can hand-create a API mock with, say, a 503 response status code and test how your code handles it, you can have confidence of how your package will respond when that rare event happens with the real API.*
- 4. Re-recording all responses can make for a huge code diff, which can blow up your repository size and make code review harder."*

Now, creating mock files by hand (or inventing some custom scripts to create them) involves more elbow grease, so it's a compromise.

8.7 Testing for API errors

In your test suite you probably want to check how things go if the server returns 502 or so, and you cannot trigger such a response to record it.

With `httptest`, to test for API errors, you need to create one or several fake mock file(s). The easiest way to do that might be to use `httptest::with_mock_dir()` that will create mock files with the expected filenames and locations, that you can then tweak. Or reading the error message of `httptest::with_mock_ap()` helps you know where to create a mock file.

With `vcr`, you either

- use `webmockr` as we showed in our demo. On the one hand it's more compact than creating a fake mock file, on the other hand it's a way to test that's different from the `vcr` cassette.

```
test_that("gh_organizations errors when the API doesn't behave", {
  webmockr::enable()
  stub <- webmockr::stub_request("get", "https://api.github.com/organizations?since=1")
  webmockr::to_return(stub, status = 502)
  expect_error(gh_organizations(), "oops")
  webmockr::disable()
})
```

- or you edit a cassette by hand which would be similar to testing for API errors with `httptest`. If you did that, you'd need to skip the test when `vcr` is off, as when `vcr` is off real requests are made. For that you can use `vcr::skip_if_vcr_off()`.

8.8 Conclusion

Both `vcr` and `httptest` are similar packages in that they use mock files for allowing easier HTTP testing. They are a bit different in their design philosophy and features, which might help you choose one of them.

And now, to make things even more complex, or fun, we shall explore a third HTTP testing package that does not *mock* requests but instead spins up a local fake web service.

Chapter 9

Use webfakes

In this chapter we aim at adding HTTP testing infrastructure to `exempligratia` using `webfakes`.

9.1 Setup

Before working on all this, we need to install `webfakes`, with `install.packages("webfakes")`.

Then, we need to add `webfakes` as a `Suggests` dependency of our package, potentially via running `usethis::use_package("webfakes", type = "Suggests")`.

Last but not least, we create a setup file at `tests/testthat/setup.R`. When `testthat` runs tests, [files whose name starts with “setup” are always run first](#). We need to ensure that we set up a fake API key when there is no API token around. Why? Because if you remember well, the code of our function `gh_organizations()` checks for the presence of a token. When using our own fake web service, we obviously don't really need a token but we still need to fool our own package in contexts where there is no token (e.g. in continuous integration checks for a fork of a GitHub repository).

```
if(!nzchar(Sys.getenv("REAL_REQUESTS"))) {  
  Sys.setenv("GITHUB_PAT" = "foobar")  
}
```

The setup file could also load `webfakes`, but in our demo we will namespace `webfakes` functions instead.

9.2 Actual testing

With webfakes we will be spinning local fake web services that we will want our package to interact with instead of the real APIs. Therefore, we first need to amend the code of functions returning URLs to services to be able to change them via an environment variable. They become:

```
status_url <- function() {
  env_url <- Sys.getenv("EXEMPLIGHRATIA_GITHUB_STATUS_URL")

  if (nzchar(env_url)) {
    return(env_url)
  }

  "https://kctbh9vrttdwd.statuspage.io/api/v2/components.json"
}
```

and

```
gh_v3_url <- function() {
  api_url <- Sys.getenv("EXEMPLIGHRATIA_GITHUB_API_URL")

  if (nzchar(api_url)) {
    return(api_url)
  }

  "https://api.github.com/"
}
```

Having these two switches is crucial.

Then, let's tweak our test of `gh_api_status()`.

```
test_that("gh_api_status() works", {
  if (!nzchar(Sys.getenv("REAL_REQUESTS"))) {
    app <- webfakes::new_app()
    app$get("/", function(req, res) {
      res$send_json(
        list( components =
              list(
                list(
                  name = "API Requests",
                  status = "operational"
                )
              )
            )
      )
    })
  }
})
```

```

    )
  )
),
  auto_unbox = TRUE
)
})
web <- webfakes::local_app_process(app, start = TRUE)
web$local_env(list(EXEMPLIGHRATIA_GITHUB_STATUS_URL = "{url}"))
}

testthat::expect_type(gh_api_status(), "character")
})

```

So what's happening here?

- When we're not asking for requests to the real service to be made (`Sys.getenv("REAL_REQUESTS")`), we prepare a new app via `webfakes::new_app()`. It's a very simple one, that returns, for GET requests, a list corresponding to what we're used to getting out of the status API, except that a) it's much smaller and b) the "operational" status is hard-coded.
- When then create a local app process via `webfakes::local_app_process(, start = TRUE)`. It will start right away thanks to `start=TRUE` but we could have chosen to start it later via calling e.g. `web$url()` (see `?webfakes::local_app_process`); and most importantly it will be stopped automatically after the test. No mess made!
- We set the `EXEMPLIGHRATIA_GITHUB_STATUS_URL` variable to the URL of the local app process. This is what connects our code to our fake web service.

It might seem like a lot of overhead code but

- It means no real requests are made which is our ultimate goal.
- We will get used to it.
- We can write helper code in a `testthat` helper file to not repeat ourselves in further test files; there could even be an app shared between all test files depending on your package.

Now, let's add a test for error behavior. This inspired us to change error behavior a bit with a slightly more specific error message i.e. `httr::stop_for_status(response)` became `httr::stop_for_status(response, task = "get API status, ouch!")`.

```

test_that("gh_api_status() errors when the API does not behave", {
  app <- webfakes::new_app()
  app$get("/", function(req, res) {

```

```

    res$send_status(502L)
  })
  web <- webfakes::local_app_process(app, start = TRUE)
  web$local_env(list(EXEMPLIGHRATIA_GITHUB_STATUS_URL = "{url}"))
  testthat::expect_error(gh_api_status(), "ouch")
})

```

It's a similar process to the earlier test:

- setting up a new app;
- having it return something we chose, in this case a 502 status;
- launching a local app process;
- connecting our code to it via setting the EXEMPLIGHRATIA_GITHUB_STATUS_URL environment variable to the URL of the fake service;
- test.

Last but not least let's convert our test of `gh_organizations()`,

```

test_that("gh_organizations works", {

  if (!nzchar(Sys.getenv("REAL_REQUESTS"))) {
    app <- webfakes::new_app()
    app$get("/organizations", function(req, res) {
      res$send_json(
        jsonlite::read_json(
          testthat::test_path(
            file.path("responses", "organizations.json")
          )
        ),
        auto_unbox = TRUE
      )
    })
    web <- webfakes::local_app_process(app, start = TRUE)
    web$local_env(list(EXEMPLIGHRATIA_GITHUB_API_URL = "{url}"))
  }

  testthat::expect_type(gh_organizations(), "character")
})

```

As before we

- create a new app;
- have it returned something we chose for a GET request of the `/organizations` endpoint. In this case, we have it return the content of a JSON file we created

at `tests/testthat/responses/organizations.json` by copy-pasting a real response from the API;

- launch a local app process;
- set its URL as the `EXEMPLIGHRATIA_GITHUB_API_URL` environment variable;
- test.

9.3 Also testing for real interactions

What if the API responses change? Hopefully we'd notice that thanks to following API news. However, sometimes web APIs change without any notice. Therefore it is important to run tests against the real web service once in a while.

In our tests we have used the condition

```
if (!nzchar(Sys.getenv("REAL_REQUESTS"))) {
```

before launching the app and using its URL as URL for the service. So if our tests are generic enough, we can add a CI build where the environment variable `REAL_REQUESTS` is set to true. If they are not generic enough, we can use the approach exemplified in the chapter about `httptest`.

- set up a folder `real-tests` with tests interacting with the real web service;
- add it to `Rbuildignore`;
- in a CI build, delete `tests/testthat` and replace it with `real-tests`, before running R CMD check.

9.4 Summary

- We set up webfakes usage in our package `exemplighratia` by adding a dependency on `webfakes` and by adding a setup file to fool our own package that needs an API token.
- We created and launched fake apps in our test files.

Now, how do we make sure this works?

- Turn off wifi, run the tests again. It works! Turn on wifi again.
- Open `.Renviron` (`usethis::edit_r_environ()`), edit `"GITHUB_PAT"` into `"byeGITHUB_PAT"`, re-start R, run the tests again. It works! Fix your `"GITHUB_PAT"` token in `.Renviron`.

So we now have tests that no longer rely on an internet connection nor on having API credentials.

For the full list of changes applied to `exempligratia` in this chapter, see [the pull request diff on GitHub](#).

In the next chapter, we shall compare the three approaches to HTTP testing we've demo-ed.

Chapter 10

vcr (& webmockr), httptest, webfakes

We're now at a nice stage where we have made a demo of usage for each of the HTTP testing packages, in our `exempligratia` package. Of course, the choice of strategy in the demo is a bit subjective, but we hope it showed the best of each tool.

A first message that's important to us: if you're learning about HTTP testing and using it in a branch of your own package sounds daunting, create a minimal package for playing!

10.1 What HTTP client can you use (curl, httr, crul)

- `httptest` only works with `httr` (the most popular HTTP R client);
- `vcr` (& `webmockr`) works with both `httr` and `crul` (the two “high-level” HTTP R clients);
- `webfakes` works with any R HTTP client, even base R if you wish.

10.2 Sustainability of the packages

All packages (`vcr`, `webmockr`, `httptest`, `webfakes`) are actively maintained. During the writing of this book, issues and pull requests were tackled rather quickly, and always in a very nice way.

10.3 Test writing experience

In all cases having HTTP tests, i.e. tests that work independently from any internet connection, depends on

- setup, which is mainly adding a dependency on the HTTP testing packages in DESCRIPTION, and a setup or helper file;
- providing responses from the API.

The difference between packages, the *test writing experience* depends on how you can provide responses from the API, both real ones and fake ones.

With `vcr` and `httptest` for tests testing *normal behavior*, after set up (for which there is a helper function), testing is just a function away (`vcr::use_cassette()`, `httptest::with_mock_dir()`, `httptest::with_mock_api()`). Recording happens automatically during the first run of tests. You might also provide fake recorded response or dumb down the existing ones. For creating *API errors*, and API sequence of responses (e.g. 502 then 200), you end up either using `webmockr`, or amending mock files, see [vcr](#) and `httptest` related docs.¹

With `webfakes` you need to create an app. There could be one per test, per test file or for the whole test suite. It might seem like more overhead code but being able to share an app between different tests reduces this effort. You can test for an [API sequence of responses \(e.g. 502 then 200\)](#) by following an how-to. The one thing that's not supported in `webfakes` yet is a smooth workflow for recording responses, so at the time of writing you might need to write your own workflow for recording responses.

In general setup&test writing might be easier for packages with mocking (`vcr` and `httptest`) but you might be able to replicate more complex behavior with `webfakes` (such as an [OAuth dance](#)).

10.3.1 The special case of secrets

With `webfakes` as no authentication is needed at any point, you have less chance of exposing a secret.

With `httptest` only the body of responses is saved, so unless it contains secrets, no further effort is needed. If you *need* to redact mock files, see [the corresponding vignette](#).

With `vcr` as all HTTP interactions, including request URLs and headers, are saved to disk, you will most often have to use the `filter_sensitive_data`, `filter_request_header` and/or `filter_response_header` arguments of `vcr::vcr_configure()`.

¹Sequence of requests are not [supported smoothly yet](#) by `httptest`.

10.3.2 How about making real requests

In all three cases, switching back to real requests might be an environment variable away (turning `vcr` off, setting the URL of the real web service as `URL` to be connected to instead of a `webfakes` fake web service). However, your tests using fixed/fake responses / a fake web service might not work with real requests as you can't trigger an API error, and as you might test for specific values in your tests using mock files whereas the API returns something different every day. Therefore, and it's a challenge common to all three packages, you might need to choose to have *distinct* tests as integration tests/[contract tests](#). See also our chapter about [making real requests](#).

10.4 Test debugging experience

Sadly sometimes one needs to run code from the tests in an interactive session, either to debug tests after making a code change, or to learn how to write HTTP tests.

With `webfakes`, debugging works this way: load the helper or test file where

- the app is created,
- the environment variable connecting your package code to the fake web service is changed.

Then run your code. To debug *webfakes apps*, follow the [guidance](#).

With `vcr`, refer to the [debugging](#) vignette: you'll have to load the helper file or source the setup file after making sure the paths use in it work both from `tests/testthat/` and the package root (see `?vcr::vcr_test_path`), and then use `vcr::inject_cassette()`; don't forget to run `vcr::eject_cassette()` afterwards. With `webmockr` debugging is quite natural, run the code that's in the test, in particular `webmockr::enable()` and `webmockr::disable()`.

With `httptest`, the process is similar as with `vcr` except the key functions are

- `use_mock_api()`
- `.mockPaths`.

10.5 Conclusion

In this chapter we compared the three R packages that make HTTP testing easier. If you are still unsure which one to pick, first try packages out without commitment, in branches or so, but then choose one and [commit to your lock-in](#).

”Every piece of code written in a given language or framework is a step away from any other language, and five more minutes you’ll have to spend migrating it to something else. That’s fine. You just have to decide what you’re willing to be locked into.

(...)

Code these days becomes obsolete so quickly, regardless of what’s chosen. By the time your needs change, by the time the latest framework is obsolete, all of the code will be rotten anyway

(...)

The most dangerous feature about these articles examining cloud lock-in is that they introduce a kind of paralysis into teams that result in applications never being completely fleshed out or finished.”

Vicki Boykis, “Commit to your lock-in”.

Part III

Advanced Topics

Chapter 11

Making real requests

As touched upon in the Whole Games section, it's good to have *some* tests against the real API. Indeed, the web resource can change.

11.1 What can change?

What can happen?

- An API introducing rate-limiting;
- A web resource disappearing;
- etc.

11.2 How to make real requests

Maybe you can just run the same tests without using the mock files.

- with `vcr`, this behavior is [one environment variable away](#) (namely, `VCR_TURN_OFF`).
- with `httptest` you can create the [same kind of behavior](#).
- with `webfakes` you can also [create that behavior](#).

Now this means assuming *all* your tests work with real requests.

- If a few tests won't work with real requests (say they have a fixture mimicking an [API error](#), or specific answer as if today were a given date) then you can skip

these tests when mocking/faking the web service is off. With `vcr` this means using `vcr::skip_if_vcr_off()`; with `httptest` and `webfakes` you'd create your custom skipper.

- If most tests won't work with real requests, then creating a different folder for tests making real requests makes sense. It might be less unit-y as you could view these tests as integration/contract tests. Maybe they could use [testthat's snapshot testing](#) (so you could view what's different in the API).

11.2.1 When to make real requests?

Locally, you might want to make real requests once in a while, in particular before a CRAN release.

On continuous integration you have to learn how to trigger workflows and configure build matrices to e.g.

- Have one build in your build matrix using real requests at each commit (this might be too much, see next section);
- Have one scheduled workflow once a day or once a week using real requests.

11.3 Why not make only or too many real requests?

The reasons why you can't only make real requests in your tests are the [reasons why you are reading these book](#):

- they are slower;
- you can't test for API errors;
- etc.

Now no matter what your setup is you don't want to make *too many* real requests as it can be bad for the web resource and bad for you (e.g. using all your allowed requests!). Regarding allowed requests, if possible you could however increase them by requesting for some sort of special development account if such a thing exists for the API you are working with.

11.4 A complement to real requests: API news!

Running real requests is important to notice if something changes in the API (expected requests, responses). Now, you can and should also follow the news of the web resource you are using in case there is something in place.

- Subscribe to the API newsletter if there's one;
- Read API changelogs if they are public;
- In particular, if the API is developed on GitHub/GitLab/etc. you could watch the repo or subscribe to releases, so that you might ^automatically get notified.

Chapter 12

CRAN- (and Bioconductor) preparedness for your tests

There is no one right answer to how to manage your tests for CRAN, except that you do want a [clean check result on CRAN at all times](#). This probably applies to Bioconductor too. The following is a discussion of the various considerations - which should give you enough information to make an educated decision.

12.1 Running tests on CRAN?

You can run `vcr/httpptest/webfakes` enabled tests on CRAN. CRAN is okay with files associated with tests, and so in general on CRAN you can run your tests that use cassettes, mock files or recorded responses on CRAN. Another aspect is the presence of dependencies: make sure the HTTP testing package you use is listed as `Suggests` dependency in `DESCRIPTION`! With `webfakes` this might mean also listing [optional dependencies](#) in `DESCRIPTION`. With `webfakes`, your tests if run on CRAN should not assume the availability of a given port.

When running HTTP tests on CRAN, be aware of a few things:

- If your tests require any secret environment variables or R options (apart from the “foobar” ones used to fool your package when using a saved response), they won’t be available on CRAN. In these cases you likely want to skip these tests with `testthat::skip_on_cran()`.
- If your tests have cassettes, mock files or recorded responses with sensitive information in them, you probably do not want to have those cassettes on the internet, in which case you won’t be running `vcr` enabled tests on CRAN either.

In the case of sensitive information, you might want to `Rbuildignore` the cassettes, mock files or recorded responses (and to `gitignore` them or make your package development repository private).

- There is a maximal size for package sources so you will want your cassettes, mock files or recorded responses to not be too big. There are three ways to limit their size
 - Make requests that do not generate a huge response (e.g. tweak the time range);
 - Edit the recorded responses (why not even copy-paste responses from the API docs as those are often short) — see [vcr docs about editing cassettes for pros and cons](#);
 - Share [cassettes](#) / mock files / recorded responses between tests.

Do *not* [compress cassettes, mock files or recorded responses](#): CRAN submissions are already compressed; compressed files will make git diffs hard to use.

12.2 Skipping a few tests on CRAN?

If you are worried at all about problems with HTTP tests on CRAN you can use `testthat::skip_on_cran()` to skip specific tests. Make sure your tests run somewhere else (on continuous integration) regularly!

We'd recommend not running tests making real requests on CRAN, even when interacting with an API without authentication.

12.3 Skipping all tests on CRAN?

If you have a good continuous integration setup (several operating systems, scheduled runs, etc.) why not skip all tests on CRAN?

12.4 Stress-test your package

To stress-test your package before a CRAN submission, use `rhub::check_for_cran()` without passing any environment variable to the function, and use [WinBuilder](#).

Chapter 13

Security

When developing a package that uses secrets (API keys, OAuth tokens) and produces them (OAuth tokens, sensitive data),

- You want the secrets to be usable by you, collaborators and CI services, without being readable by anyone else;
- You want tests and checks (e.g. vignette building) that use the secrets to be turned off in environments where secrets won't be available (CRAN, forks of your development repository).

Your general attitude should be to think about:

- what are my secrets (an API key, an OAuth2.0 access token and the refresh token, etc.) and where/how exactly are there used (in the query part of an URL? as a header? which header, Authentication or something else?) – packages like `httr` might abstract some of the complexity for you but you need to really know where secrets are used and could be leaked,
- what could go wrong (e.g. your token ending up being published),
- how to prevent that (save your unedited token outside of your package, make sure it is not printed in logs or present in package check artefacts),
- how to fix mistakes (how do you deactivate a token and how do you check no one used it in the meantime).

13.1 Managing secrets securely

13.1.1 Follow best practice when developing your package

This book is about testing but security starts with how you develop your package. To better protect your users' secret,

- It might be best not to let users pass API keys as parameters. It's best to have them save them in `.Renvirom` or e.g. using the [keyring package](#). This way, API keys are not in scripts. The [opencage package](#) might provide some inspiration.
- If the API you are working with lets you pass keys either in the request headers or query string, prefer to use request headers.

13.1.2 Share secrets with continuous integration services

You need to share secrets with continuous integration services... for real requests only! For tests using `vcr`, `httptest` or `webfakes`, you at most need a fake secret, e.g. "foobar" as API key – except for recording cassettes and mock files, but that is something you do locally.

In GitHub repositories, when storing a new secret, do not save it with quotes. I.e. if your secret is "blabla", the field should contain `blabla`, not `"blabla"` nor `'blabla'`.

```
knitr::include_graphics("secret.png")
```

13.1.2.1 API keys

For API keys, you can use something like GitHub repo secrets if you use GitHub Actions. Then for the secret to be accessible as environment variable from your workflow in GitHub Actions [as explained in gargle docs](#) you need to add a line like

```
env:
  PACKAGE_PASSWORD: ${ secrets.PACKAGE_PASSWORD }
```

13.1.2.2 More complex objects

If your secret is an OAuth token, you might be able to re-create it from pieces, where the pieces are strings you can store as repo secrets much like what you'd do for an API key. E.g. if your secret is an OAuth token, the [actual secrets](#) are the access token and refresh token.

```
env:
  ACCESS_TOKEN: ${ secrets.ACCESS_TOKEN }
  REFRESH_TOKEN: ${ secrets.REFRESH_TOKEN }
```

Therefore you could re-create it using e.g. the `credentials` argument of `httr::oauth2.0_token()`. The re-creation using environment variables `Sys.getenv("ACCESS_TOKEN")` and `Sys.getenv("REFRESH_TOKEN")` would happen in a [testthat helper file](#).

13.1.3 Secret files

For files, you will need to use encryption and to store a text-version of the encryption key/passwords as GitHub repo secret if you use GitHub Actions. Read the documentation of the continuous integration service you are using to find out how secrets are protected and how you can use them in your builds.

See [gargle vignette about securely managing tokens](#).

The approach is:

- Create your OAuth token locally, either outside of your package folder, or inside of it if you really want to, but **gitignored and Rbuildignored**.
- Encrypt it using e.g. the [user-friendly cyphr package](#). Save the code for this and for the step before in a file e.g. `inst/secrets.R` for when you need to re-create a token as even refresh tokens expire.
- For encrypting you need some sort of password. You will want to save it securely as `text` in your [user-level .Renviron](#) and in your GitHub repo secrets (or equivalent secret place for other CI services). E.g. create a key via `sodium_key <- sodium::keygen()` and get its text equivalent via `sodium::bin2hex(sodium_key)`. E.g. the latter command might give me `e46b7faf296e3f0624e6240a6efafe3dfb17b92ae0089c7e51952934b60749f2` and I will save this in `.Renviron`

```
MEETUPR_PWD="e46b7faf296e3f0624e6240a6efafe3dfb17b92ae0089c7e51952934b60749f2"
```

Example of a script creating and encrypting an OAuth token (for the Meetup API).

```
# thanks Jenny Bryan https://github.com/r-lib/gargle/blob/4fcf142fde43d107c6a20f905052f24859133c3
token_path <- testthat::test_path(".meetup_token.rds")
use_build_ignore(token_path)
use_git_ignore(token_path)

meetupr::meetup_auth(
  token = NULL,
  cache = TRUE,
  set_renv = FALSE,
  token_path = token_path
)

# sodium_key <- sodium::keygen()
# set_renv("MEETUPR_PWD" = sodium::bin2hex(sodium_key))
# set_renv being an internal function taken from rtweet
# that saves something to .Renviron
```

```
# get key from environment variable
key <- cyphr::key_sodium(sodium::hex2bin(Sys.getenv("MEETUPR_PWD")))

cyphr::encrypt_file(
  token_path,
  key = key,
  dest = testthat::test_path("secret.rds")
)
```

- In tests you have a setup file with code like below.

```
key <- cyphr::key_sodium(sodium::hex2bin(Sys.getenv("MEETUPR_PWD")))

temptoken <- tempfile(fileext = ".rds")

cyphr::decrypt_file(
  testthat::test_path("secret.rds"),
  key = key,
  dest = temptoken
)
```

Now what happens in contexts where `MEETUPR_PWD` is not available? Well there should be no tests using it! See [our chapter about making real requests](#).

13.1.4 Do not store secrets in the cassettes, mock files, recorded responses

- With `vcr` make sure to [configure vcr correctly](#).
- With `httptest` only the response body (and headers, but not by default) are recorded. If those contains secrets, refer to the documentation about [redacting sensitive information](#).
- With `webfakes` you will be creating recorded responses yourself, make sure this process does not leak secrets. If you test something related to authentication, use fake secrets.

If the API you are interacting with uses OAuth for instance, make sure you are not leaking access tokens nor *refresh tokens*.

13.1.5 Escape tests that require secrets

This all depends on your setup for testing [real requests](#). You have to be sure no test requiring secrets will be run on [CRAN](#) for instance.

13.2 Sensitive recorded responses?

In that case you might want to gitignore the cassettes / mock files / recorded responses, and skip the tests using them on continuous integration (e.g. `testthat::skip_on_ci()` or something more involved). You'd also [Rbuildignore](#) the cassettes / mock files / recorded responses, as you do not want to release them to CRAN.

13.3 Further resources

Some tools might help you detect leaks or prevent them.

- [shhgit](#)'s goal is "Find secrets in your code. Secrets detection for your GitHub, GitLab and Bitbucket repositories".
- [Yelp's detect-secret](#) is "An enterprise friendly way of detecting and preventing secrets in code."
- [git-secret](#) is a "bash tool to store your private data inside a git repo".

Chapter 14

Faking HTTP errors

With HTTP testing you can test the behavior of your package in case of an API error without having to actually trigger an API error. This is important for testing your package's gracefulness (informative error message for the user) and robustness (if you e.g. use retrying in case of API errors).

14.1 How to test for API errors (e.g. 503)

Different possibilities:

- Use [webmockr](#) as in [our demo](#).
- [Edit a vcr cassette](#); be careful to skip this test when vcr is off with `vcr::skip_if_vcr_is_off()`.
- With `httptest`, edit a mock file as in [our demo](#), or create it from scratch.
- With `webfakes`, choose what to return, have a specific app for the test, see [our demo](#).

14.2 How to test for sequence of responses (e.g. 503 then 200)

Different possibilities:

- Use [webmockr](#).
- [Edit a vcr cassette](#); be careful to skip this test when vcr is off with `vcr::skip_if_vcr_is_off()`
- With `httptest`, this is [not easy yet](#)

- With webfakes, follow [the docs](#). Also have a specific app for the test as this is not the behavior you want in all your tests.

Chapter 15

Contributor friendliness

How do you make your package wrapping an HTTP resource contributor-friendly?

rOpenSci has some general advice on [contributor-friendliness](#).

Now, there are some more aspects when dealing with HTTP testing.

15.1 Taking notes about encryption

In your contributing guide, make sure you note how you e.g. created an encrypted token for the tests. Link to a script that one could run to re-create it. Good for future contributors including yourself!

15.2 Providing a sandbox

It might be very neat to provide a **sandbox**, even if just for yourself.

- If interacting with say Twitter API you might want to create a Twitter account dedicated to this.
- If interacting with some sort of web platform you might want to create an account special for storing test data.
- Some web APIs provide a test API key, a test account that one can request access to.

Make sure to take notes on how to create / request access to a sandbox, in your contributing guide.

15.3 Switching between accounts depending on the development mode

Your package might have some behaviour to load a default token for instance, placed in an app dir. Now, for testing, you might want it to load another token, and you probably also want the token choice to be as automatic as possible.

The `rtweet` package has such logic.

- It [detects testing/dev mode](#).

```
is_testing <- function() {
  identical(Sys.getenv("TESTTHAT"), "true")
}
is_dev_mode <- function() {
  exists(".__DEVTOOLS__", .getNamespace("rtweet"))
}
```

- If some environment variables are present it is able to [create a testing token](#).

```
rtweet_test <- function() {
  access_token <- Sys.getenv("RTWEET_ACCESS_TOKEN")
  access_secret <- Sys.getenv("RTWEET_ACCESS_SECRET")

  if (identical(access_token, "") || identical(access_secret, "")) {
    return()
  }

  rtweet_bot(
    "7rX1CfEY0jrtZenmBhjljPz03",
    "rM3HOLDqmjWzr9UN4cvscchlkFprPNNg99zJJU5R8iYtpCOP0q",
    access_token,
    access_secret
  )
}
```

- The [testing token](#) or a default token is loaded depending on the development mode.

15.4 Documenting HTTP testing

Contributors to the package might not be familiar with the HTTP testing package(s) you use (this is true of any non-trivial test setup). Make sure your contributing guide mentions pre-requisites and link to resources (maybe even this very book?).

Part IV

Conclusion

Chapter 16

Conclusion

Once you get here you will have read about basic HTTP (testing) concepts in R, discovered four great packages in demos (`vcr` & `webmockr`, `httptest`, `webfakes`), and dived into more advanced topics like security.

What's next? Applying those tools in your package(s), of course!

- Pick one or several HTTP testing package(s) for your package. Examples of combinations:
 - `vcr` for testing normal behavior, `webmockr` for testing behavior in case of web resource errors.
 - `vcr` or `httptest` for most tests, `webfakes` for more advanced things like OAuth2.0 flows or slow internet connection.
- Read all the docs of the HTTP testing package(s) you choose – a very worthy use of your time. For `vcr` and `webmockr` you can even stay here in this book and take advantage of the “`vcr` details” and “`webmockr` details” sections.
- For more examples, you could also look at the reverse dependencies of the HTTP testing package(s) you use to see how they are used by other developers.
- Follow developments of the HTTP testing package(s) you choose. As all four packages we presented are developed on GitHub, you could e.g. release-watch their repositories. They are also all distributed on CRAN, so you might use your usual channel for learning about CRAN updates.
- Participate in the development of the HTTP testing package(s) you choose? Your bug reports, feature requests, contributions might be helpful. Make sure to read the contributing guide and to look at current activity in the repositories.
- Report any feedback about this book, your experience HTTP testing, tips, etc.

- in the [GitHub repository](#) of the book,
- or in [rOpenSci forum](#).

Happy HTTP testing!

Part V

webmockr details

Chapter 17

Mocking HTTP Requests

The very very short version is: [webmockr](#) helps you stub HTTP requests so you don't have to repeat yourself.

17.1 Package documentation

Check out <https://docs.ropensci.org/webmockr/> for documentation on `webmockr` functions.

17.2 Features

- Stubbing HTTP requests at low http client lib level
- Setting and verifying expectations on HTTP requests
- Matching requests based on method, URI, headers and body
- Support for `testthat` via [vcr](#)
- Can be used for testing or outside of a testing context

17.3 How `webmockr` works in detail

You tell `webmockr` what HTTP request you want to match against and if it sees a request matching your criteria it doesn't actually do the HTTP request. Instead, it gives back the same object you would have gotten back with a real request, but only with the bits it knows about. For example, we can't give back the actual data you'd get from a real HTTP request as the request wasn't performed.

In addition, if you set an expectation of what `webmockr` should return, we return that. For example, if you expect a request to return a 418 error (I'm a Teapot), then that's what you'll get.

What you can match against

- HTTP method (required)

Plus any single or combination of the following:

- URI
 - Right now, we can match directly against URI's, and with regex URI patterns. Eventually, we will support RFC 6570 URI templates.
 - We normalize URI paths so that URL encoded things match URL un-encoded things (e.g. `hello world` to `hello%20world`)
- Query parameters
 - We normalize query parameter values so that URL encoded things match URL un-encoded things (e.g. `message = hello world` to `message = hello%20world`)
- Request headers
 - We normalize headers and treat all forms of same headers as equal. For example, the following two sets of headers are equal:


```
* list(H1 = "value1", content_length = 123, X_CuStOm_hEAdEr = "foo")
* list(h1 = "value1", "Content-Length" = 123, "x-cuSTOM-HeAder" = "foo")
```
- Request body

Real HTTP requests

There's a few scenarios to think about when using `webmockr`:

After doing

```
library(webmockr)
```

`webmockr` is loaded but not turned on. At this point `webmockr` doesn't change anything.

Once you turn on `webmockr` like

```
webmockr::enable()
```

webmockr will now by default not allow real HTTP requests from the http libraries that adapters are loaded for (crul and httr).

You can optionally allow real requests via `webmockr_allow_net_connect()`, and disallow real requests via `webmockr_disable_net_connect()`. You can check whether you are allowing real requests with `webmockr_net_connect_allowed()`.

Certain kinds of real HTTP requests allowed: We don't support this yet, but you can allow localhost HTTP requests with the `allow_localhost` parameter in the `webmockr_configure()` function.

Storing actual HTTP responses

webmockr doesn't do that. Check out [vcr](#)

17.4 Basic usage

```
library("webmockr")  
# enable webmockr  
webmockr::enable()
```

Stubbed request based on uri only and with the default response

```
stub_request("get", "https://httpbin.org/get")
```

```
#> <webmockr stub>  
#> method: get  
#> uri: https://httpbin.org/get  
#> with:  
#> query:  
#> body:  
#> request_headers:  
#> to_return:
```

```
library("crul")  
x <- HttpClient$new(url = "https://httpbin.org")  
x$get('get')
```

```
#> <crul response>  
#> url: https://httpbin.org/get
```

```
#> request_headers:  
#>   User-Agent: libcurl/7.64.1 r-curl/4.3 crul/1.1.0  
#>   Accept-Encoding: gzip, deflate  
#>   Accept: application/json, text/xml, application/xml, */*  
#> response_headers:  
#> status: 200
```


Chapter 18

stubs

```
library("webmockr")
```

set return objects

```
stub_request("get", "https://httpbin.org/get") %>%  
  with(  
    query = list(hello = "world")) %>%  
    to_return(status = 418)
```

```
#> <webmockr stub>  
#> method: get  
#> uri: https://httpbin.org/get  
#> with:  
#>   query: hello=world  
#>   body:  
#> request_headers:  
#> to_return:  
#> - status: 418  
#>   body:  
#> response_headers:  
#> should_timeout: FALSE  
#> should_raise: FALSE
```

```
x$get('get', query = list(hello = "world"))
```

```
#> <crul response>  
#> url: https://httpbin.org/get
```

```
#> request_headers:
#>   User-Agent: libcurl/7.64.1 r-curl/4.3 crul/1.1.0
#>   Accept-Encoding: gzip, deflate
#>   Accept: application/json, text/xml, application/xml, */*
#> response_headers:
#> status: 418
```

Stubbing requests based on method, uri and query params

```
stub_request("get", "https://httpbin.org/get") %>%
  with(query = list(hello = "world"),
       headers = list('User-Agent' = 'libcurl/7.51.0 r-curl/2.6 crul/0.3.6',
                      'Accept-Encoding' = "gzip, deflate"))
```

```
#> <webmockr stub>
#> method: get
#> uri: https://httpbin.org/get
#> with:
#>   query: hello=world
#>   body:
#>   request_headers: User-Agent=libcurl/7.51.0 r-cur..., Accept-Encoding=gzip, defl
#> to_return:
```

```
stub_registry()
```

```
#> <webmockr stub registry>
#> Registered Stubs
#>   GET: https://httpbin.org/get
#>   GET: https://httpbin.org/get?hello=world | to_return: with status 418
#>   GET: https://httpbin.org/get?hello=world with headers {"User-Agent":"libcurl/7
```

```
x <- HttpClient$new(url = "https://httpbin.org")
x$get('get', query = list(hello = "world"))
```

```
#> <crul response>
#> url: https://httpbin.org/get
#> request_headers:
#>   User-Agent: libcurl/7.64.1 r-curl/4.3 crul/1.1.0
#>   Accept-Encoding: gzip, deflate
#>   Accept: application/json, text/xml, application/xml, */*
#> response_headers:
#> status: 418
```

Stubbing requests and set expectation of a timeout

```

stub_request("post", "https://httpbin.org/post") %>% to_timeout()
x <- HttpClient$new(url = "https://httpbin.org")
x$post('post')
#> Error: Request Timeout (HTTP 408).
#> - The client did not produce a request within the time that the server was prepared
#>   to wait. The client MAY repeat the request without modifications at any later time.

```

Stubbing requests and set HTTP error expectation

```

library(fauxpas)
stub_request("get", "https://httpbin.org/get?a=b") %>% to_raise(HTTPBadRequest)
x <- HttpClient$new(url = "https://httpbin.org")
x$get('get', query = list(a = "b"))
#> Error: Bad Request (HTTP 400).
#> - The request could not be understood by the server due to malformed syntax.
#>   The client SHOULD NOT repeat the request without modifications.

```

18.1 Writing to disk

There are two ways to deal with mocking writing to disk. First, you can create a file with the data you'd like in that file, then tell crul or httr where that file is. Second, you can simply give webmockr a file path (that doesn't exist yet) and some data, and webmockr can take care of putting the data in the file.

Here's the first method, where you put data in a file as your mock, then pass the file as a connection (with `file(<file path>)`) to `to_return()`.

```

## make a temp file
f <- tempfile(fileext = ".json")
## write something to the file
cat("{\"hello\":\"world\"}\n", file = f)
## make the stub
invisible(stub_request("get", "https://httpbin.org/get") %>%
  to_return(body = file(f)))
## make a request
out <- HttpClient$new("https://httpbin.org/get")$get(disk = f)
## view stubbed file content
readLines(file(f))

```

```
#> [1] "{\"hello\":\"world\"}"
```

With the second method, use `webmockr::mock_file()` to have `webmockr` handle file and contents.

```
g <- tempfile(fileext = ".json")
## make the stub
invisible(stub_request("get", "https://httpbin.org/get?a=b") %>%
  to_return(body = mock_file(path = g, payload = "{\\"hello\\":\\"mars\\"}\n")))
## make a request
out <- crul::HttpClient$new("https://httpbin.org/get?a=b")$get(disk = g)
## view stubbed file content
readLines(out$content)
```

```
#> [1] "{\\"hello\\":\\"mars\\"}" ""
```

`webmockr` also supports `httr::write_disk()`, here letting `webmockr` handle the mock file creation:

```
library(httr)
httr_mock()
## make a temp file
f <- tempfile(fileext = ".json")
## make the stub
invisible(stub_request("get", "https://httpbin.org/get?cheese=swiss") %>%
  to_return(
    body = mock_file(path = f, payload = "{\\"foo\\": \\"bar\\"}"),
    headers = list('content-type' = "application/json")
  ))
## make a request
out <- GET("https://httpbin.org/get?cheese=swiss", write_disk(f, TRUE))
## view stubbed file content
readLines(out$content)
```

```
#> [1] "{\\"foo\\": \\"bar\\"}"
```

Chapter 19

testing

```
library("webmockr")
library("crul")
library("testthat")

stub_registry_clear()

# make a stub
stub_request("get", "https://httpbin.org/get") %>%
  to_return(body = "success!", status = 200)
```

```
#> <webmockr stub>
#>   method: get
#>   uri: https://httpbin.org/get
#>   with:
#>     query:
#>     body:
#>     request_headers:
#>   to_return:
#> - status: 200
#>   body: success!
#>   response_headers:
#>   should_timeout: FALSE
#>   should_raise: FALSE
```

```
# check that it's in the stub registry
stub_registry()
```

```
#> <webmockr stub registry>
```

```
#> Registered Stubs
#> GET: https://httpbin.org/get | to_return: with body "success!" with status 200
```

```
# make the request
z <- crul::HttpClient$new(url = "https://httpbin.org")$get("get")

# run tests (nothing returned means it passed)
expect_is(z, "HttpResponse")
expect_equal(z$status_code, 200)
expect_equal(z$parse("UTF-8"), "success!")
```

Chapter 20

utilities

```
library("webmockr")
```

20.1 Managing stubs

- `enable()`
- `enabled()`
- `disable()`
- `httr_mock()`

20.2 Managing stubs

- `stub_registry()`
- `stub_registry_clear()`
- `remove_request_stub()`

20.3 Managing requests

- `request_registry()`

Part VI
vcr details

Chapter 21

Caching HTTP requests

Record HTTP calls and replay them

21.1 Package documentation

Check out <https://docs.ropensci.org/vcr/> for documentation on `vcr` functions.

21.2 Terminology

- *vcr*: the name comes from the idea that we want to record something and play it back later, like a vcr
- *cassette*: A *thing* to record HTTP interactions to. Right now the only option is the file system (writing to files), but in the future could be other things, e.g. a key-value store like Redis
- *fixture*: A fixture is something used to consistently test a piece of software. In this case, a cassette (just defined above) is a fixture - used in unit tests. If you use our setup function `vcr_setup()` the default directory created to hold cassettes is called `fixtures/` as a signal as to what the folder contains.
- *Persisters*: how to save requests - currently only option is the file system
- *serialize*: translating data into a format that can be stored; here, translate HTTP request and response data into a representation on disk to read back later
- *Serializers*: how to serialize the HTTP response - currently only option is YAML; other options in the future could include e.g. JSON
- *insert cassette*: create a cassette (all HTTP interactions will be recorded to this cassette)
- *eject cassette*: eject the cassette (no longer recording to that cassette)

- *replay*: refers to using a cached result of an http request that was recorded earlier

21.3 Basic usage

21.3.1 In tests

In your tests, for whichever tests you want to use `vcr`, wrap them in a `vcr::use_cassette()` call like:

```
library(testthat)
vcr::use_cassette("rl_citation", {
  test_that("my test", {
    aa <- rl_citation()

    expect_is(aa, "character")
    expect_match(aa, "IUCN")
    expect_match(aa, "www.iucnredlist.org")
  })
})
```

OR put the `vcr::use_cassette()` block on the inside, but put `testthat` expectations outside of the `vcr::use_cassette()` block:

```
library(testthat)
test_that("my test", {
  vcr::use_cassette("rl_citation", {
    aa <- rl_citation()
  })

  expect_is(aa, "character")
  expect_match(aa, "IUCN")
  expect_match(aa, "www.iucnredlist.org")
})
```

Don't wrap the `use_cassette()` block inside your `test_that()` block with `testthat` expectations inside the `use_cassette()` block, as you'll only get the line number that the `use_cassette()` block starts on on failures.

The first time you run the tests, a "cassette" i.e. a file with recorded HTTP interactions, is created at `tests/fixtures/rl_citation.yml`. The times after that, the cassette will be used. If you change your code and more HTTP interactions are needed in the code wrapped by `vcr::use_cassette("rl_citation"`, delete `tests/fixtures/rl_citation.yml` and run the tests again for re-recording the cassette.

21.3.2 Outside of tests

If you want to get a feel for how vcr works, although you don't need too.

```
library(vcr)
library(crul)

cli <- crul::HttpClient$new(url = "https://eu.httpbin.org")
system.time(
  use_cassette(name = "helloworld", {
    cli$get("get")
  })
)
```

The request gets recorded, and all subsequent requests of the same form used the cached HTTP response, and so are much faster

```
system.time(
  use_cassette(name = "helloworld", {
    cli$get("get")
  })
)
```

Importantly, your unit test deals with the same inputs and the same outputs - but behind the scenes you use a cached HTTP response - thus, your tests run faster.

The cached response looks something like (condensed for brevity):

```
http_interactions:
- request:
  method: get
  uri: https://eu.httpbin.org/get
  body:
    encoding: ''
    string: ''
  headers:
    User-Agent: libcurl/7.54.0 r-curl/3.2 crul/0.5.2
response:
  status:
    status_code: '200'
    message: OK
    explanation: Request fulfilled, document follows
  headers:
    status: HTTP/1.1 200 OK
    connection: keep-alive
```

```

body:
  encoding: UTF-8
  string: "{\n  \"args\": {}, \n  \"headers\": {\n    \"Accept\": \"application/javascript/xml, application/xml, */*\", \n    \"Accept-Encoding\": \"gzip, deflate\", \n    \"Connection\": \"close\", \n    \"Host\": \"httpbin.org\", \n    \"User-Agent\": \"libcurl/7.54.0 r-curl/3.2 crul/0.5.2\", \n    \"origin\": \"111.222.333.44\", \n    \"url\": \"https://eu.httpbin.org/get\"}\n}\n"
  recorded_at: 2018-04-03 22:55:02 GMT
  recorded_with: vcr/0.1.0, webmockr/0.2.4, crul/0.5.2

```

All components of both the request and response are preserved, so that the HTTP client (in this case `crul`) can reconstruct its own response just as it would if it wasn't using `vcr`.

21.3.3 Less basic usage

For tweaking things to your needs, make sure to read the docs about [configuration](#) (e.g., where are the fixtures saved? can they be re-recorded automatically regularly?) and [request matching](#) (how does `vcr` match a request to a recorded interaction?)

All components of both the request and response are preserved, so that the HTTP client (in this case `crul`) can reconstruct its own response just as it would if it wasn't using `vcr`.

21.4 vcr enabled testing

21.4.1 check vs. test

TLDR: Run `devtools::test()` before running `devtools::check()` for recording your cassettes.

When running tests or checks of your whole package, note that you'll get different results with `devtools::check()` (check button of RStudio build pane) vs. `devtools::test()` (test button of RStudio build pane). This arises because `devtools::check()` runs in a temporary directory and files created (vcr cassettes) are only in that temporary directory and thus don't persist after `devtools::check()` exits.

However, `devtools::test()` does not run in a temporary directory, so files created (vcr cassettes) are in whatever directory you're running it in.

Alternatively, you can run `devtools::test_file()` (or the "Run test" button in RStudio) to create your vcr cassettes one test file at a time.

21.4.2 CI sites: GitHub Actions, Appveyor, etc.

Refer to [the security chapter](#).

Chapter 22

Advanced vcr usage

Now that we've covered basic vcr usage, it's time for some more advanced usage topics.

```
library("vcr")
```

22.1 Mocking writing to disk

If you have http requests for which you write the response to disk, then use `vcr_configure()` to set the `write_disk_path` option. See more about the [write_disk_path configuration option](#).

Here, we create a temporary directory, then set the fixtures

```
tmpdir <- tempdir()
vcr_configure(
  dir = file.path(tmpdir, "fixtures"),
  write_disk_path = file.path(tmpdir, "files")
)
```

```
#> <vcr configuration>
#>   Cassette Dir: /var/folders/24/8k48jl6d249_n_qfxwsl6xvm0000gn/T//Rtmp6UOL19/fixtures
#>   Record: once
#>   Serialize with: yaml
#>   URI Parser: crul::url_parse
#>   Match Requests on: method, uri
#>   Preserve Bytes?: FALSE
#>   Logging?: FALSE
```

```
#> ignored hosts:
#> ignore localhost?: FALSE
#> Write disk path: /var/folders/24/8k48jl6d249_n_qfxws16xvm0000gn/T//Rtmp6UOL19/fil
```

Then pass a file path (that doesn't exist yet) to `curl`'s `disk` parameter. `vcr` will take care of handling writing the response to that file in addition to the cassette.

```
library(curl)
## make a temp file
f <- tempfile(fileext = ".json")
## make a request
cas <- use_cassette("test_write_to_disk", {
  out <- HttpClient$new("https://httpbin.org/get")$get(disk = f)
})
file.exists(out$content)
```

```
#> [1] TRUE
```

```
out$parse()
```

```
#> [1] "{\n  \"args\": {},\n  \"headers\": {\n    \"Accept\": \"application/json, text
```

This also works with `httr`. The only difference is that you write to disk with a function `httr::write_disk(path)` rather than a parameter.

Note that when you write to disk when using `vcr`, the cassette is slightly changed. Instead of holding the http response body itself, the cassette has the file path with the response body.

```
http_interactions:
- request:
  method: get
  uri: https://httpbin.org/get
  response:
    headers:
      status: HTTP/1.1 200 OK
      access-control-allow-credentials: 'true'
    body:
      encoding: UTF-8
      file: yes
      string: /private/var/folders/fc/n7g_vrvn0sx_st0p81xb3ts40000gn/T/Rtmp5W4o1r/files
```

And the file has the response body that otherwise would have been in the `string` yaml field above:

```
{
  "args": {},
  "headers": {
    "Accept": "application/json, text/xml, application/xml, */*",
    "Accept-Encoding": "gzip, deflate",
    "Host": "httpbin.org",
    "User-Agent": "libcurl/7.54.0 r-curl/4.3 crul/0.9.0"
  },
  "origin": "24.21.229.59, 24.21.229.59",
  "url": "https://httpbin.org/get"
}
```


Chapter 23

Configure vcr

```
library("vcr")
```

You can also get the default configuration variables via `vcr_config_defaults()`

```
vcr_config_defaults()
```

```
#> $verbose_errors
#> [1] FALSE
#>
#> $write_disk_path
#> NULL
#>
#> $filter_response_headers
#> NULL
#>
#> $filter_request_headers
#> NULL
#>
#> $filter_sensitive_data
#> NULL
#>
#> $log_opts
#> $log_opts$file
#> [1] "vcr.log"
#>
#> $log_opts$log_prefix
#> [1] "Cassette"
```

```
#>
#> $log_opts$date
#> [1] TRUE
#>
#>
#> $log
#> [1] FALSE
#>
#> $linked_context
#> NULL
#>
#> $cassettes
#> list()
#>
#> $allow_http_connections_when_no_cassette
#> [1] FALSE
#>
#> $clean_outdated_http_interactions
#> [1] FALSE
#>
#> $re_record_interval
#> NULL
#>
#> $turned_off
#> [1] FALSE
#>
#> $preserve_exact_body_bytes
#> [1] FALSE
#>
#> $uri_parser
#> [1] "crul::url_parse"
#>
#> $ignore_request
#> NULL
#>
#> $ignore_localhost
#> [1] FALSE
#>
#> $ignore_hosts
#> NULL
#>
#> $persist_with
#> [1] "FileSystem"
#>
#> $json_pretty
#> [1] FALSE
```

```
#>
#> $serialize_with
#> [1] "yaml"
#>
#> $allow_unused_http_interactions
#> [1] TRUE
#>
#> $match_requests_on
#> [1] "method" "uri"
#>
#> $record
#> [1] "once"
#>
#> $dir
#> [1] "."
```

These defaults are set when you load `vcr` - you can override any of them as described below.

23.1 Set configuration variables

Use `vcr_configure()` to set configuration variables.

For example, set a single variable:

```
vcr_configure(
  dir = "foobar/vcr_cassettes"
)
```

```
#> <vcr configuration>
#>  Cassette Dir: foobar/vcr_cassettes
#>  Record: once
#>  Serialize with: yaml
#>  URI Parser: crul::url_parse
#>  Match Requests on: method, uri
#>  Preserve Bytes?: FALSE
#>  Logging?: FALSE
#>  ignored hosts:
#>  ignore localhost?: FALSE
#>  Write disk path:
```

Or many at once:

```
vcr_configure(  
  dir = "foobar/vcr_cassettes",  
  record = "all"  
)
```

```
#> <vcr configuration>  
#>  Cassette Dir: foobar/vcr_cassettes  
#>  Record: all  
#>  Serialize with: yaml  
#>  URI Parser: crul::url_parse  
#>  Match Requests on: method, uri  
#>  Preserve Bytes?: FALSE  
#>  Logging?: FALSE  
#>  ignored hosts:  
#>  ignore localhost?: FALSE  
#>  Write disk path:
```

23.2 Re-set to defaults

```
vcr_configure_reset()
```

23.3 Details on some of the config options

23.3.1 dir

Directory where cassettes are stored

```
vcr_configure(dir = "new/path")
```

```
#> <vcr configuration>  
#>  Cassette Dir: new/path  
#>  Record: once  
#>  Serialize with: yaml  
#>  URI Parser: crul::url_parse  
#>  Match Requests on: method, uri  
#>  Preserve Bytes?: FALSE  
#>  Logging?: FALSE  
#>  ignored hosts:  
#>  ignore localhost?: FALSE  
#>  Write disk path:
```


23.3.2 record

The record mode

One of: 'all', 'none', 'new_episodes', 'once'. See `?recording` for info on the options

```
vcr_configure(record = "new_episodes")
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
#>  Serialize with: yaml
#>  URI Parser: crul::url_parse
#>  Match Requests on: method, uri
#>  Preserve Bytes?: FALSE
#>  Logging?: FALSE
#>  ignored hosts:
#>  ignore localhost?: FALSE
#>  Write disk path:
```

23.3.3 match_requests_on

Customize how vcr matches requests

```
vcr_configure(match_requests_on = c('query', 'headers'))
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
#>  Serialize with: yaml
#>  URI Parser: crul::url_parse
#>  Match Requests on: query, headers
#>  Preserve Bytes?: FALSE
#>  Logging?: FALSE
#>  ignored hosts:
#>  ignore localhost?: FALSE
#>  Write disk path:
```

23.3.4 allow_unused_http_interactions

Allow HTTP connections when no cassette

Default is TRUE, and thus does not error when http interactions are unused. You can set to FALSE in which case vcr errors when a cassette is ejected and not all http interactions have been used.

```
vcr_configure(allow_unused_http_interactions = FALSE)
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
#>  Serialize with: yaml
#>  URI Parser: crul::url_parse
#>  Match Requests on: query, headers
#>  Preserve Bytes?: FALSE
#>  Logging?: FALSE
#>  ignored hosts:
#>  ignore localhost?: FALSE
#>  Write disk path:
```

23.3.5 serialize_with

Which serializer to use. Right now only option is “yaml”

```
vcr_configure(serialize_with = "yaml")
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
#>  Serialize with: yaml
#>  URI Parser: crul::url_parse
#>  Match Requests on: query, headers
#>  Preserve Bytes?: FALSE
#>  Logging?: FALSE
#>  ignored hosts:
#>  ignore localhost?: FALSE
#>  Write disk path:
```

23.3.6 persist_with

Which persister to use. Right now only option is “FileSystem”

```
vcr_configure(persist_with = "FileSystem")
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
```

```
#> Serialize with: yaml
#> URI Parser: crul::url_parse
#> Match Requests on: query, headers
#> Preserve Bytes?: FALSE
#> Logging?: FALSE
#> ignored hosts:
#> ignore localhost?: FALSE
#> Write disk path:
```

23.3.7 ignoring some requests

ignore_hosts

Specify particular hosts to ignore. By ignore, we mean that real HTTP requests to the ignored host will be allowed to occur, while all others will not.

```
vcr_configure(ignore_hosts = "google.com")
```

```
#> <vcr configuration>
#> Cassette Dir: new/path
#> Record: new_episodes
#> Serialize with: yaml
#> URI Parser: crul::url_parse
#> Match Requests on: query, headers
#> Preserve Bytes?: FALSE
#> Logging?: FALSE
#> ignored hosts: google.com
#> ignore localhost?: FALSE
#> Write disk path:
```

ignore_localhost

Ignore all localhost requests

```
vcr_configure(ignore_localhost = TRUE)
```

```
#> <vcr configuration>
#> Cassette Dir: new/path
#> Record: new_episodes
#> Serialize with: yaml
#> URI Parser: crul::url_parse
#> Match Requests on: query, headers
#> Preserve Bytes?: FALSE
#> Logging?: FALSE
```

```
#> ignored hosts: google.com
#> ignore localhost?: TRUE
#> Write disk path:
```

ignore_request

THIS DOESN'T WORK YET

How to ignore requests

For ignoring requests, you can for example, have real http requests go through (ignored by vcr) while other requests are handled by vcr. For example, let's say you want requests to google.com to be ignored:

```
vcr_configure(ignore_hosts = "google.com")
use_cassette("foo_bar", {
  crul::HttpClient$new("https://httpbin.org/get")$get()
  crul::HttpClient$new("https://google.com")$get()
})
```

The request to httpbin.org will be handled by vcr, a cassette created for the request/response to that url, while the google.com request will be ignored and not cached at all.

Note: ignoring requests only works for the crul package for now; it should work for httr in a later vcr version.

23.3.8 uri_parse

Which uri parser to use

By default we use crul::url_parse, but you can use a different one. Remember to pass in the function quoted, and namespaced.

```
vcr_configure(uri_parser = "urltools::url_parse")
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
#>  Serialize with: yaml
#>  URI Parser: urltools::url_parse
#>  Match Requests on: query, headers
#>  Preserve Bytes?: FALSE
#>  Logging?: FALSE
#>  ignored hosts: google.com
#>  ignore localhost?: TRUE
#>  Write disk path:
```

23.3.9 `preserve_exact_body_bytes`

Some HTTP servers are not well-behaved and respond with invalid data. Set `preserve_exact_body_bytes` to `TRUE` to base64 encode the result body in order to preserve the bytes exactly as-is. `vcr` does not do this by default, since base64-encoding the string removes the human readability of the cassette.

```
vcr_configure(preserve_exact_body_bytes = TRUE)
```

```
#> <vcr configuration>
#>  Cassette Dir: new/path
#>  Record: new_episodes
#>  Serialize with: yaml
#>  URI Parser: urltools::url_parse
#>  Match Requests on: query, headers
#>  Preserve Bytes?: TRUE
#>  Logging?: FALSE
#>  ignored hosts: google.com
#>  ignore localhost?: TRUE
#>  Write disk path:
```

23.3.10 `filter_sensitive_data`

A named list of values to replace. Sometimes your package or script is working with sensitive tokens/keys, which you do not want to accidentally share with the world.

Before recording (writing to a cassette) we do the replacement and then when reading from the cassette we do the reverse replacement to get back to the real data.

```
vcr_configure(
  filter_sensitive_data = list("<some_api_key>" = Sys.getenv('MY_API_KEY'))
)
```

Before recording to disk, the env var `MY_API_KEY` is retrieved from your machine, and we find instances of it, and replace with `<some_api_key>`. When replaying to create the HTTP response object we put the real value of the env var back in place.

To target specific request or response headers see `filter_request_headers` and `filter_response_headers`.

23.3.11 `filter_request_headers`

Expects a character vector or a named list. If a character vector, or any unnamed element in a list, the request header is removed before being written to the cassette.

If a named list is passed, the name is the header and the value is the value with which to replace the real value.

A request header you set to remove or replace is only removed/replaced from the cassette, and any requests using a cassette, but will still be in your `crul` or `httr` response objects on a real request that creates the cassette.

Examples:

```
vcr_configure(  
  filter_request_headers = "Authorization"  
)  
vcr_configure(  
  filter_request_headers = c("Authorization", "User-Agent")  
)  
vcr_configure(  
  filter_request_headers = list(Authorization = "<<<not-my-bearer-token>>>")  
)
```

23.3.12 filter_response_headers

Expects a character vector or a named list. If a character vector, or any unnamed element in a list, the response header is removed before being written to the cassette.

If a named list is passed, the name is the header and the value is the value with which to replace the real value.

A response header you set to remove or replace is only removed/replaced from the cassette, and any requests using a cassette, but will still be in your `crul` or `httr` response objects on a real request that creates the cassette.

Examples:

```
vcr_configure(  
  filter_response_headers = "server"  
)  
vcr_configure(  
  filter_response_headers = c("server", "date")  
)  
vcr_configure(  
  filter_response_headers = list(server = "fake-server")  
)
```

23.3.13 filter_query_parameters

Expects a character vector or a named list. If a character vector, or any unnamed element in a list, the query parameter is removed (both parameter name and value)

before being written to the cassette.

If a named list is passed, the name is the query parameter name and the value is the value with which to replace the real value.

A response header you set to remove or replace is only removed/replaced from the cassette, and any requests using a cassette, but will still be in your `crul` or `httr` response objects on a real request that creates the cassette.

Beware of your `match_requests_on` option when using this filter. If you filter out a query parameter it's probably a bad idea to match on query given that there is no way for `vcr` to restore the exact http request from your cassette after one or more query parameters is removed or changed. One way you could filter a query parameter and still match on query or at least on the complete uri is to use replacement behavior (a named list), but instead of `list(a="b")` use two values `list(a=c("b", "c"))`, where "c" is the string to be stored in the cassette. You could of course replace those values with values from environment variables so that you obscure the real values if your code is public.

Examples:

```
# completely drop parameter "user"
vcr_configure(
  filter_query_parameters = "user"
)
# completely drop parameters "user" and "api_key"
vcr_configure(
  filter_query_parameters = c("user", "api_key")
)
# replace the value of parameter "api_key" with "fake-api-key"
# NOTE: in this case there's no way to put back any value on
# subsequent requests, so we have to match by dropping this
# parameter value before comparing URIs
vcr_configure(
  filter_query_parameters = list(api_key = "fake-api-key")
)
# replace the value found at Sys.getenv("MY_API_KEY") of parameter
# "api_key" with the value "foo". When using a cassette on subsequent
# requests, we can replace "foo" with the value at Sys.getenv("MY_API_KEY")
# before doing the URI comparison
vcr_configure(
  filter_query_parameters = list(api_key = c(Sys.getenv("MY_API_KEY"), "foo"))
)
```


Chapter 24

Record modes

Record modes dictate under what circumstances http requests/responses are recorded to cassettes (disk). Set the recording mode with the parameter `record` in the `use_cassette()` and `insert_cassette()` functions.

24.1 `once`

The `once` record mode will:

- Replay previously recorded interactions.
- Record new interactions if there is no cassette file.
- Cause an error to be raised for new requests if there is a cassette file.

It is similar to the `new_episodes` record mode, but will prevent new, unexpected requests from being made (i.e. because the request URI changed or whatever).

`once` is the default record mode, used when you do not set one.

24.2 `none`

The `none` record mode will:

- Replay previously recorded interactions.
- Cause an error to be raised for any new requests.

This is useful when your code makes potentially dangerous HTTP requests. The `none` record mode guarantees that no new HTTP requests will be made.

24.3 new_episodes

The `new_episodes` record mode will:

- Record new interactions.
- Replay previously recorded interactions.

It is similar to the `once` record mode, but will **always** record new interactions, even if you have an existing recorded one that is similar (but not identical, based on the `match_request_on` option).

24.4 all

The `all` record mode will:

- Record new interactions.
- Never replay previously recorded interactions.

This can be temporarily used to force `vcr` to re-record a cassette (i.e. to ensure the responses are not out of date) or can be used when you simply want to log all HTTP requests.

Chapter 25

Request matching

To match previously recorded requests, `vcr` has to try to match new HTTP requests to a previously recorded one. By default, we match on HTTP method (e.g., `GET`) and URI (e.g., `http://foo.com`), following Ruby's `VCR` gem.

You can customize how we match requests with one or more of the following options, some of which are on by default, some of which can be used together, and some alone.

- `method`: Use the **method** request matcher to match requests on the HTTP method (i.e. `GET`, `POST`, `PUT`, `DELETE`, etc). You will generally want to use this matcher. The **method** matcher is used (along with the **uri** matcher) by default if you do not specify how requests should match.
- `uri`: Use the **uri** request matcher to match requests on the request URI. The **uri** matcher is used (along with the **method** matcher) by default if you do not specify how requests should match.
- `host`: Use the **host** request matcher to match requests on the request host. You can use this (alone, or in combination with **path**) as an alternative to **uri** so that non-deterministic portions of the URI are not considered as part of the request matching.
- `path`: Use the **path** request matcher to match requests on the path portion of the request URI. You can use this (alone, or in combination with **host**) as an alternative to **uri** so that non-deterministic portions of the URI
- `query`: Use the **query** request matcher to match requests on the query string portion of the request URI. You can use this (alone, or in combination with others) as an alternative to **uri** so that non-deterministic portions of the URI are not considered as part of the request matching.
- `body`: Use the **body** request matcher to match requests on the request body.
- `headers`: Use the **headers** request matcher to match requests on the request headers.

You can set your own options by tweaking the `match_requests_on` parameter in `use_cassette()`:

```
library(vcr)
```

```
use_cassette(name = "foo_bar", {  
  cli$post("post", body = list(a = 5))  
},  
  match_requests_on = c('method', 'headers', 'body')  
)
```

25.1 Matching

25.1.1 headers

```
library(crul)  
library(vcr)  
cli <- crul::HttpClient$new("https://httpbin.org/get",  
  headers = list(foo = "bar"))  
use_cassette(name = "nothing_new", {  
  one <- cli$get()  
},  
  match_requests_on = 'headers'  
)  
cli$headers$foo <- "stuff"  
use_cassette(name = "nothing_new", {  
  two <- cli$get()  
},  
  match_requests_on = 'headers'  
)  
one$request_headers  
two$request_headers
```

Chapter 26

Debugging your tests that use vcr

Sometimes your tests using a vcr cassette will fail and you will want to debug them.

26.1 An HTTP request has been made that vcr does not know how to handle

If you get an error starting with “An HTTP request has been made that vcr does not know how to handle:” when running your tests, it means that the code in your test makes an HTTP request for which there is no matching information in the cassette you are using. You might have added a request, or changed one slightly.

The easy fix is: delete the cassette and re-run the test to re-record the cassette. Run the test a second time to ensure all is well. If not, escalate to the next paragraph.

Maybe you didn’t actually want to change the request you are making. Make sure the requests do not contain something random, or something related to e.g. what time it is now, in the URI (`http://foo.com?time=13`). To make sure things are not varying, you might want to use mocking (of e.g. a function returning the current time), setting a random seed, using `withhr` (for e.g. setting an option to a certain value in your test).

26.1.1 Actual debugging

Ideally you will want to run the code of the tests as if it were run inside tests, in particular, using the same vcr cassette.

26.1.2 Prepare your debugging environment

You will first need to load either the vcr helper `tests/testthat/helper-vcr.R` (e.g. via `devtools::load_all()`) or source the vcr setup file `tests/testthat/setup-vcr.R` i.e. the file with these lines (and maybe others)

```
library("vcr")
invisible(vcr::vcr_configure(
  dir = vcr::vcr_test_path("fixtures"),
  filter_sensitive_data = list("<<github_api_token>>" = Sys.getenv('GITHUB_PAT'))
))
vcr::check_cassette_names()
```

If instead of `vcr::vcr_test_path("fixtures")` you see `../fixtures`, replace `../fixtures` with `vcr::vcr_test_path("fixtures")`, as `vcr::vcr_test_path()` is a function that is meant to help exactly what you will want: have the path to `tests/fixtures/` work from `tests` and from the root (which is where you will be running the code to debug it).

So that is one step (loading the vcr helper or sourcing the vcr setup file), or maybe two (if you also had to replace `../fixtures` with `vcr::vcr_test_path("fixtures")`).

26.1.3 Debugging itself

Now look at the test whose code you are trying to debug e.g.

```
foo <- function() crul::ok('https://httpbin.org/get')

test_that("foo works", {
  vcr::use_cassette("testing", {
    x <- foo()
  })
  expect_true(x)
})
```

If you want to run the code as if you were in the test,

```
foo <- function() crul::ok('https://httpbin.org/get')
vcr::insert_cassette("testing") # it will be created if needed
x <- foo()
x
# further interactive debugging and fixes
vcr::eject_cassette("testing")
```

26.1.4 Logging

You can use vcr's built in logging to help in your debugging process. To configure logging, use the `vcr::vcr_configure()` function, and set `log=TRUE` and set options for logging on the `log_opts` parameter as a named list. See `?vcr_configure` for details.

Here, we are setting our log file to be a temporary file that will be cleaned up at the end of the R session. Here, the file extension is `.log`, but the file extension does not matter.

```
vcr::vcr_configure(
  dir = vcr::vcr_test_path("fixtures"),
  log = TRUE,
  log_opts = list(file = file.path(tempdir(), "vcr.log"))
)
```

With `log=TRUE` you can continue with debugging. Open the log file you set in a text editor or other location; or examine in your shell/terminal.

As an example, after running the block above

```
foo <- function() crul::ok('https://httpbin.org/get')

test_that("foo works", {
  vcr::use_cassette("testing", {
    x <- foo()
  })
  expect_true(x)
})
```

If we open the log file we'll see the logs for each step vcr takes in handling an HTTP request. The logs have information on what cassette was used, what exact time it was recorded, what matchers were in use, the cassette options, and how a request is handled.

```
[Cassette: 'testing'] - 2020-11-24 16:05:17 - Init. HTTPInteractionList w/ request matchers [meth
[Cassette: 'testing'] - 2020-11-24 16:05:17 - Initialized with options: {name: testing, record: c
[Cassette: 'testing'] - 2020-11-24 16:05:17 - Handling request: head https://httpbin.org/get (dis
[Cassette: 'testing'] - 2020-11-24 16:05:17 - Identified request type: (recordable) for head http
[Cassette: 'testing'] - 2020-11-24 16:05:17 - Recorded HTTP interaction: head https://httpbin.
```

Logging isn't meant to be turned on all the time - rather only for debugging/informational purposes.

26.1.5 Return to normal development

Make sure you ejected the cassette you were using!

Unless your vcr helper/setup file tweaked more things than you would like, you do not even need to re-start R, but you could, just to be on the safe side.

Chapter 27

Security with vcr

Refer to [the security chapter](#) for more general guidance.

27.1 Keeping secrets safe

To keep your secrets safe, you need to use parameters of `vcr : vcr_configure()` that tell vcr either where secrets are (and what to put in their place), or what secrets are (and what to put in their place). It is best if you know how secrets are used in requests: e.g. is the API key passed as a header or part of the query string? Maybe you will need different strategies for the different secrets (e.g. an OAuth2.0 access token will be set as Authorization header but an OAuth2.0 refresh token might be in a query string).

In all cases, it is crucial to look at your cassettes before putting them on the public web, just to be sure you got the configuration right!

27.1.1 If the secret is in a request header

You can use `filter_request_headers!`

There are different ways to use it.

```
# Remove one header from the cassettes
vcr_configure(
  filter_request_headers = "Authorization"
)

# Remove two headers from the cassettes
```

```
vcr_configure(
  filter_request_headers = c("Authorization", "User-Agent")
)

# Replace one header with a given string
vcr_configure(
  filter_request_headers = list(Authorization = "<<<not-my-bearer-token>>>")
)
```

27.1.2 If the secret is in a response header

You can use `filter_response_headers` that works like `filter_request_headers`.

27.1.3 If the secret is somewhere else

In this case you need to tell vcr what the secret string is via `filter_sensitive_data`. Do not write the secret string directly in the configuration, that'd defeat the purpose of protecting it! Have the secret in an environment variable for instance and tell vcr to read it from there.

The configuration parameter `filter_sensitive_data` accepts a named list.

Each element in the list should be of the following format:

```
thing_to_replace_it_with = thing_to_replace
```

We replace all instances of `thing_to_replace` with `thing_to_replace_it_with`.

Before recording (writing to a cassette) we do the replacement and then when reading from the cassette we do the reverse replacement to get back to the real data.

```
vcr_configure(
  filter_sensitive_data = list("<<<my_api_key>>>" = Sys.getenv('API_KEY'))
)
```

You want to make the string that replaces your sensitive string something that won't be easily found elsewhere in the response body/headers/etc.

27.2 API keys and tests run in varied contexts

- For real requests a real API key is needed.
- For requests using cassettes a fake API key is needed to fool your package. That is why in our [demo of vcr](#) we set a fake API key in a test setup file.

27.3 Other security

Let us know about any other security concerns! Surely there's things we haven't considered yet.

Chapter 28

Turning vcr on and off

Sometimes you may need to turn off `vcr`, either for individual function calls, individual test blocks, whole test files, or for the entire package. The following attempts to break down all the options.

`vcr` has the following four exported functions:

- `turned_off()` - Turns `vcr` off for the duration of a code block
- `turn_off()` - Turns `vcr` off completely, so that it no longer handles every HTTP request
- `turn_on()` - turns `vcr` on; the opposite of `turn_off()`
- `turned_on()` - Asks if `vcr` is turned on, returns a boolean

Instead of using the above four functions, you could use environment variables to achieve the same thing. This way you could enable/disable `vcr` in non-interactive environments such as continuous integration, Docker containers, or running R non-interactively from the command line. The full set of environment variables `vcr` uses, all of which accept only `TRUE` or `FALSE`:

- `VCR_TURN_OFF`: turn off `vcr` altogether; set to `TRUE` to skip any `vcr` usage; default: `FALSE`
- `VCR_TURNED_OFF`: set the `turned_off` internal package setting; this does not turn off `vcr` completely as does `VCR_TURN_OFF` does, but rather is looked at together with `VCR_IGNORE_CASSETTES`
- `VCR_IGNORE_CASSETTES`: set the `ignore_cassettes` internal package setting; this is looked at together with `VCR_TURNED_OFF`

28.1 turned_off

`turned_off()` lets you temporarily make a real HTTP request without completely turning `vcr` off, unloading it, etc.

What happens internally is we turn off `vcr`, run your code block, then on exit turn `vcr` back on - such that `vcr` is only turned off for the duration of your code block. Even if your code block errors, `vcr` will be turned back on due to use of `on.exit(turned_on())`

```
library(vcr)
library(crul)
turned_off({
  con <- HttpClient$new(url = "https://httpbin.org/get")
  con$get()
})
```

```
#> <crul response>
#> url: https://httpbin.org/get
#> request_headers:
#>   User-Agent: libcurl/7.54.0 r-curl/4.3 crul/0.9.0
#>   Accept-Encoding: gzip, deflate
#>   Accept: application/json, text/xml, application/xml, */*
#> response_headers:
#>   status: HTTP/1.1 200 OK
#>   date: Fri, 14 Feb 2020 19:44:46 GMT
#>   content-type: application/json
#>   content-length: 365
#>   connection: keep-alive
#>   server: gunicorn/19.9.0
#>   access-control-allow-origin: *
#>   access-control-allow-credentials: true
#> status: 200
```

28.2 turn_off/turn_on

`turn_off()` is different from `turned_off()` in that `turn_off()` is not aimed at a single call block, but rather it turns `vcr` off for the entire package. `turn_off()` does check first before turning `vcr` off that there is not currently a cassette in use. `turn_off()` is meant to make R ignore `vcr::insert_cassette()` and `vcr::use_cassette()` blocks in your test suite - letting the code in the block run as if they were not wrapped in `vcr` code - so that all you have to do to run your tests with cached requests/responses AND with real HTTP requests is toggle a single R function or environment variable.

```

library(vcr)
vcr_configure(dir = tempdir())
# real HTTP request works - vcr is not engaged here
crul::HttpClient$new(url = "https://eu.httpbin.org/get")$get()
# wrap HTTP request in use_cassette() - vcr is engaged here
use_cassette("foo_bar", {
  crul::HttpClient$new(url = "https://eu.httpbin.org/get")$get()
})
# turn off & ignore cassettes - use_cassette is ignored, real HTTP request made
turn_off(ignore_cassettes = TRUE)
use_cassette("foo_bar", {
  crul::HttpClient$new(url = "https://eu.httpbin.org/get")$get()
})
# if you turn off and don't ignore cassettes, error thrown
turn_off(ignore_cassettes = FALSE)
use_cassette("foo_bar", {
  res2=crul::HttpClient$new(url = "https://eu.httpbin.org/get")$get()
})
# vcr back on - now use_cassette behaves as before
turn_on()
use_cassette("foo_bar3", {
  res2=crul::HttpClient$new(url = "https://eu.httpbin.org/get")$get()
})

```

28.3 turned_on

turned_on() does what it says on the tin - it tells you if vcr is turned on or not.

```

library(vcr)
turn_on()
turned_on()

```

```
#> [1] TRUE
```

```

turn_off()
turned_on()

```

```
#> [1] FALSE
```

28.4 Environment variables

The `VCR_TURN_OFF` environment variable can be used within R or on the command line to turn off `vcr`. For example, you can run tests for a package that uses `vcr`, but ignore any `use_cassette/insert_cassette` usage, by running this on the command line in the root of your package:

```
VCR_TURN_OFF=true Rscript -e "devtools::test()"
```

Or, similarly within R:

```
Sys.setenv(VCR_TURN_OFF = TRUE)
devtools::test()
```

The `VCR_TURNED_OFF` and `VCR_IGNORE_CASSETTES` environment variables can be used in combination to achieve the same thing as `VCR_TURN_OFF`:

```
VCR_TURNED_OFF=true VCR_IGNORE_CASSETTES=true Rscript -e "devtools::test()"
```


Chapter 29

Managing cassettes

29.1 Why edit cassettes?

By design vcr is very good at recording HTTP interactions that actually took place. Now sometimes when testing/demo-ing your package you will want to use *fake* HTTP interactions. For instance:

- What happens if the web API returns a 503 code? Is there an informative error?
- What happens if it returns a 503 and then a 200 code? Does the retry work?
- What if the API returns too much data for even simple queries and you want to make your cassettes smaller?

In all these cases, you can edit your cassettes as long as you are aware of the risks!

29.2 Risks related to cassette editing

- If you use a vcr cassette where you replace a 200 code with a 503 code, and vcr is turned off, the test will fail because the API will probably not return an error. Use `vcr::skip_if_vcr_off()`.
- If you edit cassettes by hand you can't re-record them easily, you'd need to re-record them then re-apply your edits.

Therefore you'll need to develop a good workflow.

29.3 Example 1: test using an edited cassette with a 503

First, write your test e.g.

```
vcr::use_cassette("api-error", {
  testthat("Errors are handled well", {
    vcr::skip_if_vcr_off()
    expect_error(call_my_api()), "error message")
  })
})
```

Then run your tests a first time.

1. It will fail
2. It will have created a cassette under `tests/fixtures/api-error.yml` that looks something like

```
http_interactions:
- request:
  method: get
  uri: https://eu.httpbin.org/get
  body:
    encoding: ''
    string: ''
  headers:
    User-Agent: libcurl/7.54.0 r-curl/3.2 crul/0.5.2
  response:
    status:
      status_code: '200'
      message: OK
      explanation: Request fulfilled, document follows
    headers:
      status: HTTP/1.1 200 OK
      connection: keep-alive
    body:
      encoding: UTF-8
      string: "{\n  \"args\": {},\n  \"headers\": {\n    \"Accept\": \"application/javascript/xml, application/xml, */*\",\n    \"Accept-Encoding\": \"gzip, deflate\"\n  },\n  \"connection\": \"close\",\n  \"host\": \"httpbin.org\",\n  \"user-agent\": \"libcurl/7.54.0 r-curl/3.2 crul/0.5.2\"\n},\n  \"origin\": \"111.222.333.444\",\n  \"url\": \"https://eu.httpbin.org/get\"\n}\n"
    recorded_at: 2018-04-03 22:55:02 GMT
    recorded_with: vcr/0.1.0, webmockr/0.2.4, crul/0.5.2
```

You can edit to (new status code)

```
http_interactions:
- request:
  method: get
  uri: https://eu.httpbin.org/get
  body:
    encoding: ''
    string: ''
  headers:
    User-Agent: libcurl/7.54.0 r-curl/3.2 crul/0.5.2
  response:
    status:
      status_code: '503'
```

And run your test again, it should pass! Note the use of `vcr::skip_if_vcr_off()`: if `vcr` is turned off, there is a real API request and most probably this request won't get a 503 as a status code.

29.3.1 The same thing with webmockr

The advantage of the approach involving editing cassettes is that you only learn one thing, which is `vcr`. Now, by using the `webmockr` directly in your tests, you can also test for the behavior of your package in case of errors. Below we assume `api_url()` returns the URL `call_my_api()` calls.

```
testthat("Errors are handled well", {
  webmockr::enable()
  stub <- webmockr::stub_request("get", api_url())
  webmockr::to_return(stub, status = 503)
  expect_error(call_my_api(), "error message")
  webmockr::disable()
})
```

A big pro of this approach is that it works even when `vcr` is turned off. A con is that it's quite different from the `vcr` syntax.

29.4 Example 2: test using an edited cassette with a 503 then a 200

Here we assume your package contains some sort of `retry`.

First, write your test e.g.

```
vcr::use_cassette("api-error", {
  testthat("Errors are handled well", {
    vcr::skip_if_vcr_off()
    expect_message(thing <- call_my_api()), "retry message")
    expect_s4_class(thing, "data.frame")
  })
})
```

Then run your tests a first time.

1. It will fail
2. It will have created a cassette under `tests/fixtures/api-error.yml` that looks something like

```
http_interactions:
- request:
  method: get
  uri: https://eu.httpbin.org/get
  body:
    encoding: ''
    string: ''
  headers:
    User-Agent: libcurl/7.54.0 r-curl/3.2 crul/0.5.2
response:
  status:
    status_code: '200'
    message: OK
    explanation: Request fulfilled, document follows
  headers:
    status: HTTP/1.1 200 OK
    connection: keep-alive
  body:
    encoding: UTF-8
    string: "{\n  \"args\": {}, \n  \"headers\": {\n    \"Accept\": \"application/javascript/xml, application/xml, */*\", \n    \"Accept-Encoding\": \"gzip, deflate\", \n    \"Connection\": \"close\", \n    \"Host\": \"httpbin.org\", \n    \"User-Agent\": \"libcurl/7.54.0 r-curl/3.2 crul/0.5.2\", \n  }, \n  \"origin\": \"111.222.333.444\", \n  \"url\": \"https://eu.httpbin.org/get\"}\n\n"
  recorded_at: 2018-04-03 22:55:02 GMT
  recorded_with: vcr/0.1.0, webmockr/0.2.4, crul/0.5.2
```

You can duplicate the HTTP interaction, and make the first one return a 503 status code. `vcr` will first use the first interaction, then the second one, when making the same request.

```

http_interactions:
- request:
  method: get
  uri: https://eu.httpbin.org/get
  body:
    encoding: ''
    string: ''
  headers:
    User-Agent: libcurl/7.54.0 r-curl/3.2 crul/0.5.2
  response:
    status:
      status_code: '503'
- request:
  method: get
  uri: https://eu.httpbin.org/get
  body:
    encoding: ''
    string: ''
  headers:
    User-Agent: libcurl/7.54.0 r-curl/3.2 crul/0.5.2
  response:
    status:
      status_code: '200'
      message: OK
      explanation: Request fulfilled, document follows
    headers:
      status: HTTP/1.1 200 OK
      connection: keep-alive
    body:
      encoding: UTF-8
      string: "{\n  \"args\": {}, \n  \"headers\": {\n    \"Accept\": \"application/json,\n      text/xml, application/xml, */*\", \n    \"Accept-Encoding\": \"gzip, deflate\", \n    \"Connection\": \"close\", \n    \"Host\": \"httpbin.org\", \n    \"User-Agent\":\n      \"libcurl/7.54.0 r-curl/3.2 crul/0.5.2\"\n  }, \n  \"origin\": \"111.222.333.444\", \n  \"url\": \"https://eu.httpbin.org/get\"\n}"
    recorded_at: 2018-04-03 22:55:02 GMT
    recorded_with: vcr/0.1.0, webmockr/0.2.4, crul/0.5.2

```

And run your test again, it should pass! Note the use of `vcr::skip_if_vcr_off()`: if `vcr` is turned off, there is a real API request and most probably this request won't get a 503 as a status code.

29.4.1 The same thing with webmockr

The advantage of the approach involving editing cassettes is that you only learn one thing, which is `vcr`. Now, by using the `webmockr` directly in your tests, you can also test for the behavior of your package in case of errors. Below we assume `api_url()` returns the URL `call_my_api()` calls.

```
testthat("Errors are handled well", {
  webmockr::enable()
  stub <- webmockr::stub_request("get", api_url())
  stub %>%
  to_return(status = 503) %>%
  to_return(status = 200, body = "{\n  \"args\": {}, \n  \"headers\": {\n    \"Accept\
    text/xml, application/xml, */*\", \n    \"Accept-Encoding\": \"gzip, deflate\
    \n    \"Connection\": \"close\", \n    \"Host\": \"httpbin.org\", \n    \"User-
    \"libcurl/7.54.0 r-curl/3.2 crul/0.5.2\"\n  }, \n  \"origin\": \"111.222.333.4
    \n  \"url\": \"https://eu.httpbin.org/get\"\n}\n", headers = list(b = 6))
  expect_message(thing <- call_my_api(), "retry message")
  expect_s4_class(thing, "data.frame")
  webmockr::disable()
})
```

The pro of this approach is the elegance of the stubbing, with the two different responses. Each `webmockr` function like `to_return()` even has an argument `times` indicating the number of times the given response should be returned.

The con is that on top of being different from `vcr`, in this case where we also needed a good response in the end (the one with a 200 code, and an actual body), writing the mock is much more cumbersome than just recording a `vcr` cassette.

Be aware when you add your cassettes to either `.gitignore` and/or `.Rbuildignore`.

29.5 gitignore cassettes

The `.gitignore` file lets you tell `[git]` what files to ignore - those files are not tracked by git and if you share the git repository to the public web, those files in the `.gitignore` file won't be shared in the public version.

When using `vcr` you may want to include your cassettes in the `.gitignore` file. You may want to when your cassettes contain sensitive data that you don't want to have on the internet & don't want to hide with `filter_sensitive_data`.

You may want to have your cassettes included in your GitHub repo, both to be present when tests run on CI, and when others run your tests.

There's no correct answer on whether to gitignore your cassettes. Think about security implications and whether you want CI and human contributors to use previously created cassettes or to create/use their own.

29.6 Rbuildignore cassettes

The `.Rbuildignore` file is used to tell R to ignore certain files/directories.

There's not a clear use case for why you'd want to add vcr cassettes to your `.Rbuildignore` file, but if you do be aware that will affect your vcr enabled tests.

29.7 sharing cassettes

Sometimes you may want to share or re-use cassettes across tests, for example to [reduce the size for package sources](#) or to test different functionality of your package functions that make the same query under the hood.

To do so, you can use the same cassette name for multiple `vcr::use_cassette()` calls. `vcr::check_cassette_names()` will complain about duplicate cassette names, preventing you from accidentally re-using cassettes, however. To allow duplicates, you can provide a character vector of the cassette names you want to re-use to the `allowed_duplicates` argument of `vcr::check_cassette_names()`. That way you can use the same cassette across multiple tests.

29.8 deleting cassettes

Removing a cassette is as easy as deleting in your file finder, or from the command line, or from within a text editor or RStudio.

If you delete a cassette, on the next test run the cassette will be recorded again.

If you do want to re-record a test to a cassette, instead of deleting the file you can toggle [record modes](#).

29.9 cassette file types

For right now the only persistence option is yaml. So all files have a `.yaml` extension.

When other persist options are added, additional file types may be found. The next persist type is likely to be JSON, so if you use that option, you'd have `.json` files instead of `.yaml` files.

Chapter 30

Gotchas

There's a few things to watch out for when using `vcr`.

- **Security:** Don't put your secure API keys, tokens, etc. on the public web. See the [Security chapter](#) and the [vcr security chapter](#).
- **API key issues:** Running `vcr` enabled tests in different contexts when API keys are used can have some rough edges.
- **Dates:** Be careful when using dates in tests with `vcr`. e.g. if you generate today's date, and pass that in to a function in your package that uses that date for an HTTP request, the date will be different from the one in the matching cassette, causing a `vcr` failure.
- **HTTP errors:** It's a good idea to test failure behavior of a web service in your test suite. Sometimes `vcr` can handle that and sometimes it cannot. Open any issues about this because ideally i think `vcr` could handle all cases of HTTP failures.
- **Very large response bodies:** A few things about large response bodies. First, `vcr` may give you trouble with very large response bodies as we've see yaml parsing problems already. Second, large response bodies means large cassettes on disk - so just be aware of the file size if that's something that matters to you. Third, large response bodies will take longer to load into R, so you may still have a multi second test run even though the test is using a cached HTTP response.
- **Encoding:** We haven't dealt with encoding much yet at all, so we're likely to run into encoding issues. One blunt instrument for this for now is to set `preserve_exact_body_bytes = TRUE` when running `vcr::use_cassette()` or `vcr::insert_cassette()`, which stores the response body as base64.
- **devtools::check vs. devtools::test:** See [\(21.4.1\)](#)
- **ignored files:** See [\(29\)](#)

30.1 Correct line identification

To get the actual lines where failures occur, you can wrap the `test_that` block in a `use_cassette()` block:

```
library(testthat)
vcr::use_cassette("rl_citation", {
  test_that("my test", {
    aa <- rl_citation()

    expect_is(aa, "character")
    expect_match(aa, "IUCN")
    expect_match(aa, "www.iucnredlist.org")
  })
})
```

OR put the `use_cassette()` block on the inside, but make sure to put `testthat` expectations outside of the `use_cassette()` block:

```
library(testthat)
test_that("my test", {
  vcr::use_cassette("rl_citation", {
    aa <- rl_citation()
  })

  expect_is(aa, "character")
  expect_match(aa, "IUCN")
  expect_match(aa, "www.iucnredlist.org")
})
```

Do not wrap the `use_cassette()` block inside your `test_that()` block with `testthat` expectations inside the `use_cassette()` block, as you'll only get the line number that the `use_cassette()` block starts on on failures.

Chapter 31

Session info

31.1 Session info

```
library("magrittr")

dependencies <- attachment::att_from_rmds(".")
dependencies <- dependencies[!dependencies %in% c("attachment", "bookdown", "knitr")]

sessioninfo::package_info(
  pkgs = dependencies
) %>%
  as.data.frame() %>%
  .[, c("package", "ondiskversion")] %>%
  knitr::kable()
```

	package	ondiskversion
askpass	askpass	1.1
assertthat	assertthat	0.2.1
base64enc	base64enc	0.1-3
brio	brio	1.1.1
callr	callr	3.6.0
cli	cli	2.3.1
cpp11	cpp11	0.2.7
crayon	crayon	1.4.1
crul	crul	1.1.0
curl	curl	4.3
desc	desc	1.3.0
diffobj	diffobj	0.3.4
digest	digest	0.6.27
ellipsis	ellipsis	0.3.1
evaluate	evaluate	0.14
fansi	fansi	0.4.2
fauxpas	fauxpas	0.5.0
fs	fs	1.5.0
gh	gh	1.2.1
gitcreds	gitcreds	0.1.1
glue	glue	1.4.2
httpcode	httpcode	0.3.0
httr	httr	1.4.2
ini	ini	0.3.1
jsonlite	jsonlite	1.7.2
lifecycle	lifecycle	1.0.0
magrittr	magrittr	2.0.1
mime	mime	0.10
openssl	openssl	1.4.3
pillar	pillar	1.5.1
pkgbuild	pkgbuild	1.2.0
pkgconfig	pkgconfig	2.0.3
pkgload	pkgload	1.2.0
praise	praise	1.0.0
prettyunits	prettyunits	1.1.1
processx	processx	3.5.0
ps	ps	1.6.0
R6	R6	2.5.0
Rcpp	Rcpp	1.0.6
rematch2	rematch2	2.1.2
rlang	rlang	0.4.10
rprojroot	rprojroot	2.0.2
rstudioapi	rstudioapi	0.13
sessioninfo	sessioninfo	1.1.1
sys	sys	3.4
testthat	testthat	3.0.2
tibble	tibble	3.1.0
triebeard	triebeard	0.3.0
urltools	urltools	1.7.3
utf8	utf8	1.2.1
vcr	vcr	0.6.0
vctrs	vctrs	0.3.7
waldo	waldo	0.2.5

None of `crul`, `webmockr`, `vcr`, `httptest` have compiled code, but an underlying dependency of all of them, `curl` does. See [curl's README](#) for installation instructions in case you run into `curl` related problems. `webfakes` has compiled code.

31.2 Full session info

Session info for this book

```
sessioninfo::session_info()
```

```
#> - Session info -----
#> setting value
#> version R version 4.0.4 (2021-02-15)
#> os      macOS Catalina 10.15.7
#> system x86_64, darwin17.0
#> ui      X11
#> language (EN)
#> collate en_US.UTF-8
#> ctype   en_US.UTF-8
#> tz      UTC
#> date    2021-04-02
#>
#> - Packages -----
#> package * version      date       lib source
#> assertthat 0.2.1      2019-03-21 [1] CRAN (R 4.0.2)
#> attachment 0.2.1      2021-01-21 [1] CRAN (R 4.0.2)
#> base64enc  0.1-3      2015-07-28 [1] CRAN (R 4.0.2)
#> bookdown    0.21.9     2021-04-02 [1] Github (rstudio/bookdown@df6786e)
#> cli         2.3.1      2021-02-23 [1] CRAN (R 4.0.2)
#> crayon      1.4.1      2021-02-08 [1] CRAN (R 4.0.2)
#> crul        * 1.1.0      2021-02-15 [1] CRAN (R 4.0.2)
#> curl        4.3        2019-12-02 [1] CRAN (R 4.0.1)
#> desc        1.3.0      2021-03-05 [1] CRAN (R 4.0.2)
#> digest      0.6.27     2020-10-24 [1] CRAN (R 4.0.2)
#> evaluate    0.14       2019-05-28 [1] CRAN (R 4.0.1)
#> fauxpas    0.5.0      2020-04-13 [1] CRAN (R 4.0.2)
#> fs          1.5.0      2020-07-31 [1] CRAN (R 4.0.2)
#> gh          1.2.1      2021-04-01 [1] CRAN (R 4.0.4)
#> gitcreds    0.1.1      2020-12-04 [1] CRAN (R 4.0.2)
#> glue        1.4.2      2020-08-27 [1] CRAN (R 4.0.2)
#> highr      0.8        2019-03-20 [1] CRAN (R 4.0.2)
#> htmltools   0.5.1.9000 2021-04-02 [1] Github (rstudio/htmltools@9abc3f2)
#> httpcode    0.3.0      2020-04-10 [1] CRAN (R 4.0.2)
```

```
#> httr          * 1.4.2      2020-07-20 [1] CRAN (R 4.0.2)
#> jsonlite      1.7.2      2020-12-09 [1] CRAN (R 4.0.2)
#> knitr         1.31       2021-01-27 [1] CRAN (R 4.0.2)
#> magrittr     * 2.0.1      2020-11-17 [1] CRAN (R 4.0.2)
#> pkgload       1.2.0      2021-02-23 [1] CRAN (R 4.0.2)
#> R6            2.5.0      2020-10-28 [1] CRAN (R 4.0.2)
#> Rcpp          1.0.6      2021-01-15 [1] CRAN (R 4.0.2)
#> rlang         0.4.10     2020-12-30 [1] CRAN (R 4.0.2)
#> rmarkdown     2.7          2021-02-19 [1] CRAN (R 4.0.2)
#> rprojroot     2.0.2      2020-11-15 [1] CRAN (R 4.0.2)
#> rstudioapi    0.13        2020-11-12 [1] CRAN (R 4.0.2)
#> sessioninfo   1.1.1      2018-11-05 [1] CRAN (R 4.0.2)
#> stringi       1.5.3      2020-09-09 [1] CRAN (R 4.0.2)
#> stringr       1.4.0      2019-02-10 [1] CRAN (R 4.0.2)
#> testthat     * 3.0.2      2021-02-14 [1] CRAN (R 4.0.2)
#> triebeard     0.3.0      2016-08-04 [1] CRAN (R 4.0.2)
#> urltools      1.7.3      2019-04-14 [1] CRAN (R 4.0.2)
#> vcr           * 0.6.0      2020-12-12 [1] CRAN (R 4.0.2)
#> webmockr     * 0.8.0      2021-03-14 [1] CRAN (R 4.0.2)
#> whisker      0.4         2019-08-28 [1] CRAN (R 4.0.2)
#> withr         2.4.1      2021-01-26 [1] CRAN (R 4.0.2)
#> xfun          0.22       2021-03-11 [1] CRAN (R 4.0.2)
#> yaml          2.2.1      2020-02-01 [1] CRAN (R 4.0.2)
#>
#> [1] /Users/runner/work/_temp/Library
#> [2] /Library/Frameworks/R.framework/Versions/4.0/Resources/library
```

Page not found. Use the table of contents or the search bar to find your way back.