
repo2docker Documentation

Release 0.9.0

Project Jupyter

May 05, 2019

Getting started with repo2docker

1	Installing repo2docker	3
1.1	Prerequisite: Docker	3
1.2	Installing with <code>pip</code>	3
1.3	Installing from source code	3
1.4	Windows support	4
2	Using repo2docker	5
2.1	Calling repo2docker	5
2.2	Building a specific branch, commit or tag	6
2.3	Where to put configuration files	6
2.4	Debugging repo2docker with <code>--debug</code> and <code>--no-build</code>	6
2.5	Command line API	7
3	Frequently Asked Questions (FAQ)	9
3.1	How should I specify another version of Python?	9
3.2	What versions of Python (or R or Julia...) are supported?	9
3.3	Why is my repository failing to build with <code>ResolvePackageNotFound</code> ?	10
3.4	Can I add executable files to the user's <code>PATH</code> ?	10
3.5	How do I set environment variables?	10
3.6	Can I use repo2docker to bootstrap my own Dockerfile?	11
3.7	Can I use repo2docker to edit a local host repository within a Docker environment?	11
3.8	Why is my R shiny app not launching?	11
3.9	Why does repo2docker need to exist? Why not use tool like <code>source2image</code> ?	11
4	Configure the user interface	13
4.1	JupyterLab	13
4.2	nteract	13
4.3	RStudio	14
4.4	Shiny	14
4.5	Stencila	14
5	Choose languages for your environment	15
5.1	Python	15
5.2	The R Language	16
5.3	Julia	16
5.4	Languages not covered here	16
5.5	Using multiple languages at once	16

6	How to automatically create a <code>environment.yml</code> that works with <code>repo2docker</code>	17
6.1	The challenge	17
6.2	The solution	17
7	Share JupyterLab Workspaces with a repository	19
8	Build JupyterHub-ready images	21
9	Using <code>repo2docker</code> as part of your Continuous Integration	23
9.1	Getting Started	23
10	Configuration Files	27
10.1	<code>environment.yml</code> - Install a Python environment	28
10.2	<code>requirements.txt</code> - Install a Python environment	28
10.3	<code>setup.py</code> - Install Python packages	28
10.4	<code>Project.toml</code> - Install a Julia environment	28
10.5	<code>REQUIRE</code> - Install a Julia environment (legacy)	28
10.6	<code>install.R</code> - Install an R/RStudio environment	29
10.7	<code>apt.txt</code> - Install packages with <code>apt-get</code>	29
10.8	<code>DESCRIPTION</code> - Install an R package	29
10.9	<code>manifest.xml</code> - Install Stencila	29
10.10	<code>postBuild</code> - Run code after installing the environment	29
10.11	<code>start</code> - Run code before the user sessions starts	29
10.12	<code>runtime.txt</code> - Specifying runtimes	30
10.13	<code>default.nix</code> - the nix package manager	30
10.14	<code>Dockerfile</code> - Advanced environments	30
11	The Reproducible Execution Environment Specification	31
12	Contributing to <code>repo2docker</code> development	33
12.1	Types of contribution	33
12.2	Process for making a contribution	34
12.3	Guidelines to getting a Pull Request merged	34
12.4	Setting up for Local Development	35
13	The <code>repo2docker</code> roadmap	37
13.1	Using the roadmap	37
13.2	The roadmap proper	38
14	Architecture of <code>repo2docker</code>	39
14.1	Buildpack	39
14.2	Build base environment	40
14.3	Copy repository contents	40
14.4	Assemble repository environment	40
14.5	Push	41
14.6	Run	41
15	Design of <code>repo2docker</code>	43
15.1	Deterministic output	43
15.2	Reproducibility and version stability	44
15.3	Unix principles “do one thing well”	44
15.4	Composability	44
15.5	Pareto principle (The 80-20 Rule)	44
16	Common tasks	45
16.1	Running tests	45

16.2	Update and Freeze BuildPack Dependencies	45
16.3	Creating a Release	46
17	Adding a new buildpack to repo2docker	49
17.1	Criteria to balance and consider	49
17.2	Adding libraries or UI to existing buildpacks	49
18	Changelog	51
18.1	Version x.x.x	51
18.2	Version 0.9.0	51
18.3	Version 0.8.0	52
18.4	Version 0.7.0	53
18.5	Version 0.6	54
18.6	Version 0.5	54
18.7	Version 0.4.1	54
18.8	Version 0.2	54
18.9	Version 0.1.1	54
18.10	Version 0.1	54

`jupyter-repo2docker` is a tool to **build, run, and push Docker images from source code repositories** that run via a Jupyter server.

`repo2docker` fetches a repository (from GitHub, GitLab or other locations) and builds a container image based on the configuration files found in the repository. It can be used to explore a repository locally by building and executing the constructed image of the repository, or as a means of building images that are pushed to a Docker registry.

`repo2docker` is the tool used by [BinderHub](#) to build images on demand.

Please report [Bugs](#), [ask questions](#) or [contribute to the project](#).

Installing `repo2docker`

`repo2docker` requires Python 3.4 and above on Linux and macOS. See [below](#) for more information about Windows support.

1.1 Prerequisite: Docker

Install [Docker](#) as it is required to build Docker images. The [Community Edition](#), is available for free.

Recent versions of Docker are recommended. The latest version of Docker, 18.03, successfully builds repositories from [binder-examples](#). The [BinderHub](#) helm chart uses version 17.11.0-ce-dind. See the [helm chart](#) for more details.

1.2 Installing with `pip`

We recommend installing `repo2docker` with the `pip` tool:

```
python3 -m pip install jupyter-repo2docker
```

for the latest release. To install the most recent code from the upstream repository, run:

```
python3 -m pip install https://github.com/jupyter/repo2docker/archive/master.zip
```

For information on using `repo2docker`, see [Using `repo2docker`](#).

1.3 Installing from source code

Alternatively, you can install `repo2docker` from a local source tree, e.g. in case you are contributing back to this project:

```
git clone https://github.com/jupyter/repo2docker.git
cd repo2docker
python3 -m pip install -e .
```

That's it! For information on using `repo2docker`, see [Using *repo2docker*](#).

1.4 Windows support

Windows support for `repo2docker` is still in the experimental stage.

An article about [using Windows and the WSL](#) (Windows Subsystem for Linux or Bash on Windows) provides additional information about Windows and docker.

CHAPTER 2

Using repo2docker

Note: [Docker must be running](#) in order to run `repo2docker`. For more information on installing `repo2docker`, see [Installing repo2docker](#).

`repo2docker` can build a reproducible computational environment for any repository that follows [The Reproducible Execution Environment Specification](#). `repo2docker` is called with a URL/path to a repository. It then performs these steps:

1. Inspects the repository for [configuration files](#). These will be used to build the environment needed to run the repository.
2. Builds a Docker image with an environment specified in these [configuration files](#).
3. Runs a Jupyter server within the image that lets you explore the repository interactively (optional)
4. Pushes the images to a Docker registry so that it may be accessed remotely (optional)

2.1 Calling repo2docker

`repo2docker` is called with this command:

```
jupyter-repo2docker <URL-or-path to repository>
```

where `<URL-or-path to repository>` is a URL or path to the source repository for which you'd like to build an image.

For example, the following command will build an image of Peter Norvig's [Pytudes](#) repository:

```
jupyter-repo2docker https://github.com/norvig/pytudes
```

Building the image may take a few minutes.

Pytudes uses a `requirements.txt` file to specify its Python environment. Because of this, `repo2docker` will use `pip` to install dependencies listed in this `requirements.txt` file, and these will be present in the generated Docker image. To learn more about configuration files in `repo2docker` visit [Configuration Files](#).

When the image is built, a message will be output to your terminal:

```
Copy/paste this URL into your browser when you connect for the first time,  
to login with a token:  
http://0.0.0.0:36511/?token=f94f8fabb92e22f5bfab116c382b4707fc2cade56ad1ace0
```

Pasting the URL into your browser will open Jupyter Notebook with the dependencies and contents of the source repository in the built image.

2.2 Building a specific branch, commit or tag

To build a particular branch and commit, use the argument `--ref` and specify the `branch-name` or `commit-hash`. For example:

```
jupyter-repo2docker --ref 9ced85dd9a84859d0767369e58f33912a214a3cf https://github.com/  
↪norvig/pytudes
```

Tip: For reproducible builds, we recommend specifying a commit-hash to deterministically build a fixed version of a repository. Not specifying a commit-hash will result in the latest commit of the repository being built.

2.3 Where to put configuration files

`repo2docker` will look for configuration files in:

- A folder named `binder/` in the root of the repository.
- A folder named `.binder/` in the root of the repository.
- The root directory of the repository.

`repo2docker` searches for these folders in order (`binder/`, `.binder/`, `root`). Only configuration files in the first identified folder are considered.

Check the complete list of [configuration files](#) supported by `repo2docker` to see how to configure the build process.

Note: `repo2docker` builds an environment with Python 3.7 by default. If you'd like a different version, you can specify this in your [configuration files](#).

2.4 Debugging repo2docker with `--debug` and `--no-build`

To debug the docker image being built, pass the `--debug` parameter:

```
jupyter-repo2docker --debug https://github.com/norvig/pytudes
```

This will print the generated Dockerfile, build it, and run it.

To see the generated Dockerfile without actually building it, pass `--no-build` to the commandline. This Dockerfile output is for **debugging purposes** of repo2docker only - it can not be used by docker directly.

```
jupyter-repo2docker --no-build --debug https://github.com/norvig/pytudes
```

2.5 Command line API

2.5.1 jupyter-repo2docker

Fetch a repository and build a container image

```
usage: jupyter-repo2docker [-h] [--config CONFIG] [--json-logs]
                           [--image-name IMAGE_NAME] [--ref REF] [--debug]
                           [--no-build]
                           [--build-memory-limit BUILD_MEMORY_LIMIT]
                           [--no-run] [--publish PORTS] [--publish-all]
                           [--no-clean] [--push] [--volume VOLUMES]
                           [--user-id USER_ID] [--user-name USER_NAME]
                           [--env ENVIRONMENT] [--editable]
                           [--target-repo-dir TARGET_REPO_DIR]
                           [--appendix APPENDIX] [--subdir SUBDIR] [--version]
                           [--cache-from CACHE_FROM]
                           repo ...
```

repo

Path to repository that should be built. Could be local path or a git URL.

cmd

Custom command to run after building container

-h, --help

show this help message and exit

--config <config>

Path to config file for repo2docker

--json-logs

Emit JSON logs instead of human readable logs

--image-name <image_name>

Name of image to be built. If unspecified will be autogenerated

--ref <ref>

If building a git url, which reference to check out. E.g., *master*.

--debug

Turn on debug logging

--no-build

Do not actually build the image. Useful in conjunction with `--debug`.

--build-memory-limit <build_memory_limit>

Total Memory that can be used by the docker build process

--no-run

Do not run container after it has been built

--publish <ports>, **-p** <ports>
Specify port mappings for the image. Needs a command to run in the container.

--publish-all, **-P**
Publish all exposed ports to random host ports.

--no-clean
Don't clean up remote checkouts after we are done

--push
Push docker image to repository

--volume <volumes>, **-v** <volumes>
Volumes to mount inside the container, in form src:dest

--user-id <user_id>
User ID of the primary user in the image

--user-name <user_name>
Username of the primary user in the image

--env <environment>, **-e** <environment>
Environment variables to define at container run time

--editable, **-E**
Use the local repository in edit mode

--target-repo-dir <target_repo_dir>
Path inside the image where contents of the repositories are copied to.
Defaults to \${HOME} if not set

--appendix <appendix>

--subdir <subdir>

--version
Print the repo2docker version and exit.

--cache-from <cache_from>

Frequently Asked Questions (FAQ)

A collection of frequently asked questions with answers. If you have a question and have found an answer, send a PR to add it here!

3.1 How should I specify another version of Python?

One can specify a Python version in the `environment.yml` file of a repository or `runtime.txt` file if using `requirements.txt` instead of `environment.yml`.

3.2 What versions of Python (or R or Julia...) are supported?

3.2.1 Python

Repo2docker officially supports the following versions of Python (specified in your *environment.yml* or *runtime.txt* file):

- 3.7 (added in 0.7, default in 0.8)
- 3.6 (default in 0.7 and earlier)
- 3.5
- 2.7

Additional versions may work, as long as the `base environment` can be installed for your version of Python. The most likely source of incompatibility is if one of the packages in the base environment is not packaged for your Python, either because the version of the package is too new and your chosen Python is too old, or vice versa.

If Python 2.7 is specified, a separate environment for the kernel will be installed with Python 2. The notebook server will run in the default Python 3.7 environment.

3.2.2 Julia

All Julia versions since Julia 0.7.0 are supported via a *Project.toml* file, and this is the recommended way to install Julia environments. Julia versions 0.6.x and earlier are supported via a *REQUIRE* file.

3.2.3 R

Only R 3.4.4 is currently supported, which is installed via apt from the [ubuntu bionic repository](#).

3.3 Why is my repository is failing to build with ResolvePackageNotFound?

If you used `conda env export` to generate your `environment.yml` it will generate a list of packages and versions of packages that is pinned to platform specific versions. These very specific versions are not available in the linux docker image used by repo2docker. A typical error message will look like the following:

```
Step 39/44 : RUN conda env update -n root -f "environment.yml" && conda clean -tipsy &
->& conda list -n root
---> Running in ebe9a67762e4
Solving environment: ...working... failed

ResolvePackageNotFound:
- jsonschema==2.6.0=py36hb385e00_0
- libedit==3.1.20181209=hb402a30_0
- tornado==5.1.1=py36h1de35cc_0
...
```

We recommend to use `conda env export --no-builds -f environment.yml` to export your environment and then edit the file by hand to remove platform specific packages like `appnope`.

See [How to automatically create a environment.yml that works with repo2docker](#) for a recipe on how to create strict exports of your environment that will work with repo2docker.

3.4 Can I add executable files to the user's PATH?

Yes! With a *postBuild - Run code after installing the environment* file, you can place any files that should be called from the command line in the folder `~/ .local/`. This folder will be available in a user's PATH, and can be run from the command line (or as a subsequent build step.)

3.5 How do I set environment variables?

To configure environment variables for all users of a repository use the *start* configuration file.

When running repo2docker locally you can use the `-e` or `--env` command-line flag for each variable that you want to define.

For example `jupyter-repo2docker -e VAR1=val1 -e VAR2=val2 ...`

3.6 Can I use repo2docker to bootstrap my own Dockerfile?

No, you can't.

If you pass the `--debug` flag to `repo2docker`, it outputs the intermediate Dockerfile that is used to build the docker image. While it is tempting to copy this as a base for your own Dockerfile, that is not supported & in most cases will not work. The `--debug` output is just our intermediate generated Dockerfile, and is meant to be built in a very specific way. Hence the output of `--debug` can not be built with a normal `docker build -t .` or similar traditional docker command.

Check out the [binder-examples](#) GitHub organization for example repositories you can copy & modify for your own use!

3.7 Can I use repo2docker to edit a local host repository within a Docker environment?

Yes: use the `--editable` or `-E` flag (don't confuse this with the `-e` flag for environment variables), and run `repo2docker` on a local repository:

```
repo2docker -E my-repository/
```

This builds a Docker container from the files in that repository (using, for example, a `requirements.txt` or `install.R` file), then runs that container, while connecting the working directory inside the container to the local repository outside the container. For example, in case there is a notebook file (`.ipynb`), this will open in a local webbrowser, and one can edit it and save it. The resulting notebook is updated in both the Docker container and the local repository. Once the container is exited, the changed file will still be in the local repository.

This allows for easy testing of the container while debugging some items, as well as using a fully customizable container to edit notebooks (among others).

Note: Editable mode is a convenience option that will bind the repository to the container working directory (usually `$HOME`). If you need to mount to a different location in the container, use the `--volumes` option instead. Similarly, for a fully customized user Dockerfile, this option is not guaranteed to work.

3.8 Why is my R shiny app not launching?

If you are trying to run an R shiny app using the `/shiny/folder_containing_shiny` url option, but the launch returns “The application exited during initialization.”, there might be something wrong with the specification of the app. One way of debugging the app in the container is by running the `rstudio` url, open either the ui or server file for the app, and run the app in the container `rstudio`. This way you can see the `rstudio` logs as it tries to initialise the shiny app. If you are missing a package or other dependency for the container, this will be obvious at this stage.

3.9 Why does repo2docker need to exist? Why not use tool like source2image?

The Jupyter community believes strongly in building on top of pre-existing tools whenever possible (this is why `repo2docker` buildpacks largely build off of patterns that already exist in the data analytics community). We try to perform due-diligence and search for other communities to leverage and help, but sometimes it makes the most

sense to build our own new tool. In the case of repo2docker, we spent time integrating with a pre-existing tool called [source2image](<https://github.com/openshift/source-to-image>). This is an excellent open tool for containerization, but we ultimately decided that it did not fit the use-case we wanted to address. For more information, [here's a short blog post about the decision and the reasoning behind it](<https://github.com/yuvipanda/words/blob/fd096dd49d87e624acd8bdf6d13c0cecb930bb3f/content/post/why-not-s2i.md>).

Configure the user interface

You can build several user interfaces into the resulting Docker image. This is controlled with various *configuration files*.

4.1 JupyterLab

You do not need any extra configuration in order to allow the use of the JupyterLab interface. You can launch JupyterLab from within a user session by opening the Jupyter Notebook and appending `/lab` to the end of the URL like so:

```
http(s)://<server:port>/lab
```

To switch back to the classic notebook, add `/tree` to the URL like so:

```
http(s)://<server:port>/tree
```

For example, the following Binder URL will open the [pyTudes repository](https://mybinder.org/v2/gh/norvig/pytudes/master?urlpath=lab/tree/ipynb) and begin a JupyterLab session in the `ipynb` folder:

<https://mybinder.org/v2/gh/norvig/pytudes/master?urlpath=lab/tree/ipynb>

The `/tree/ipynb` above is how JupyterLab directs you to a specific file or folder.

To learn more about URLs in JupyterLab and Jupyter Notebook, visit [starting JupyterLab](#).

4.2 nteract

[nteract](#) is a [notebook interface](#) built with React. It is similar to a more feature-filled version of the traditional Jupyter Notebook interface.

nteract comes pre-installed in any session that has been built from a Python repository.

You can launch nteract from within a user session by replacing `/tree` with `/nteract` at the end of a notebook server's URL like so:

```
http(s)://<server:port>/nteract
```

For example, the following Binder URL will open the [pyTudes repository](#) and begin an nteract session in the `ipynb` folder:

```
https://mybinder.org/v2/gh/norvig/pytudes/master?urlpath=nteract/tree/ipynb
```

The `/tree/ipynb` above is how nteract directs you to a specific file or folder.

To learn more about nteract, visit [the nteract website](#).

4.3 RStudio

The RStudio user interface is automatically enabled if a configuration file for R is detected (i.e. an R version specified in `runtime.txt`). If this is detected, RStudio will be accessible by appending `/rstudio` to the URL, like so:

```
http(s)://<server:port>/rstudio
```

For example, the following Binder link will open an RStudio session in the [R demo repository](#).

```
http://mybinder.org/v2/gh/binder-examples/r/master?urlpath=rstudio
```

4.4 Shiny

Shiny lets you create interactive visualizaions with R. Shiny is automatically enabled if a configuration file for R is detected (i.e. an R version specified in `runtime.txt`). If this is detected, Shiny will be accessible by appending `/shiny/<folder-w-shiny-files>` to the URL, like so:

```
http(s)://<server:port>/shiny/bus-dashboard
```

This assumes that a folder called `bus-dashboard` exists in the root of the repository, and that it contains all of the files needed to run a Shiny app.

For example, the following Binder link will open a Shiny session in the [R demo repository](#).

```
http://mybinder.org/v2/gh/binder-examples/r/master?urlpath=shiny/bus-dashboard/
```

4.5 Stencila

The Stencila user interface is automatically enabled if a Stencila document (i.e. a file `manifest.xml`) is detected. Stencila will be accessible by appending `/stencila` to the URL, like so:

```
http(s)://<server:port>/stencila
```

The editor will open the Stencila document corresponding to the last `manifest.xml` found in the file tree. If you want to open a different document, you can configure the path in the URL parameter `archive`:

```
http(s)://<server:port>/stencila/?archive=other-dir
```

Choose languages for your environment

You can define many different languages in your configuration files. This page describes how to use some of the more common ones.

5.1 Python

Your environment will have Python (and specified dependencies) installed when you use one of the following configuration files:

- `requirements.txt`
- `environment.yml`

Note: By default, the environment will have **Python 3.7**.

Changed in version 0.8: Upgraded default Python from 3.6 to 3.7.

5.1.1 Specifying a version of Python

To specify a specific version of Python, you have two options:

- Use *environment.yml*. Conda environments let you define the Python version in `environment.yml`. To do so, add `python=X.X` to your dependencies section, like so:

```
name: python 2.7
dependencies:
  - python=2.7
  - numpy
```

- Use *runtime.txt* with *requirements.txt*. If you are using `requirements.txt` instead of `environment.yml`, you can specify the Python runtime version in a separate file called `runtime.txt`. This file contains a single line of the following form:

```
python-X.X
```

For example:

```
python-3.6
```

5.2 The R Language

To ensure that R is installed, you must specify a version of R in a `runtime.txt` file. This takes the following form:

```
r-YYYY-MM-DD
```

The date corresponds to the state of the MRAN repository at this day. Make sure that you choose a day with the desired version of your packages. For example, to use the MRAN repository on January 1st, 2018, add this line to `runtime.txt`:

```
r-2018-01-01
```

Note that to install specific packages with the R environment, you should use the `install.R` configuration file.

5.3 Julia

To build an environment with Julia, include a configuration file called `Project.toml`. The format of this file is documented at [the Julia Pkg.jl documentation](#). To specify a specific version of Julia to install, put a Julia version in the `[compat]` section of the `Project.toml` file, as described here: <https://julialang.github.io/Pkg.jl/v1/compatibility/>.

5.4 Languages not covered here

If a language is not “officially” supported by a build pack, it can often be installed with a `postBuild` script. This will run arbitrary `bash` commands, and can be used to download / install a language.

5.5 Using multiple languages at once

It may also be possible to combine multiple languages in a single environment. The details on how to accomplish this with all possible combinations are outside the scope of this guide. However we recommend that you take a look at the [Multi-Language Demo](#) repository for some inspiration.

How to automatically create a `environment.yml` that works with `repo2docker`

This how-to explains how to create a `environment.yml` that specifies all installed packages and their precise versions from your environment.

6.1 The challenge

`conda env export -f environment.yml` creates a strict export of all packages. This is the most robust for reproducibility, but it does bake in potential platform-specific packages, so you can only use an exported environment on the same platform.

`repo2docker` uses a linux based image as the starting point for every docker image it creates. However a lot of people use OSX or Windows as their day to day operating system. This means that the `environment.yml` created by a strict export will not work with error messages saying that certain packages can not be resolved (`ResolvePackageNotFound`).

6.2 The solution

Follow this procedure to create a strict export of your environment that will work with `repo2docker` and sites like mybinder.org.

We will launch a terminal inside a basic docker image, install the packages you need and then perform a strict export of the environment.

1. install `repo2docker` on your computer by following [Installing repo2docker](#)
2. in a terminal launch a basic repository `repo2docker https://github.com/binder-examples/conda-freeze` inside `repo2docker`
3. open the URL printed at the end in a browser, the URL should look like `http://127.0.0.1:61037/?token=30e61ec80bda6dd0d14805ea76bb59e7b0cd78b5d6b436f0`

4. open a terminal by clicking “New -> Terminal” next to the “Upload” button on the right hand side of the webpage
5. install the packages your project requires with `conda install <yourpackages>`
6. use `conda env export -n root` to print the environment
7. copy and paste the environment you just printed into a `environment.yml` in your projects repository
8. close your browser tabs and exit the repo2docker session by pressing Ctrl-C.

This will give you a strict export of your environment that precisely pins the versions of packages in your environment based on a linux environment.

Share JupyterLab Workspaces with a repository

JupyterLab uses `workspaces` to save the current state of windows, settings, and documents that are open in a JupyterLab session. It is a way to persist the general configuration over time.

It is possible to export JupyterLab workspaces and load them in to another JupyterLab installation in order to share a workspace with someone else.

In order to package your workspace with a repository, we recommend following the steps in this example repository:

<https://github.com/ian-r-rose/binder-workspace-demo/>

Build JupyterHub-ready images

JupyterHub allows multiple users to collaborate on a shared Jupyter server. `repo2docker` can build Docker images that can be shared within a JupyterHub deployment. For example, mybinder.org uses JupyterHub and `repo2docker` to allow anyone to build a Docker image of a git repository online and share an executable version of the repository with a URL to the built image.

To build **JupyterHub**-ready Docker images with `repo2docker`, the version of your JupyterHub deployment must be included in the `environment.yml` or `requirements.txt` of the git repositories you build.

If your instance of JupyterHub uses `DockerSpawner`, you will need to set its command to run `jupyterhub-singleuser` by adding this line in your configuration file:

```
c.DockerSpawner.cmd = ['jupyterhub-singleuser']
```

Using `repo2docker` as part of your Continuous Integration

We've created for you the `continuous-build` repository so that you can push a `Docker` container to `Docker Hub` directly from a GitHub repository that has a Jupyter notebook. Here are instructions to do this.

9.1 Getting Started

Today you will be doing the following:

1. Fork and clone the `continuous-build` GitHub repository to obtain the hidden `.circleci` folder.
2. Creating an image repository on Docker Hub
3. Connecting your repository to CircleCI
4. Push, commit, or create a pull request to trigger a build.

You don't need to install any dependencies on your host to build the container, it will be done on a continuous integration server, and the container built and available to you to pull from Docker Hub.

9.1.1 Step 1. Clone the Repository

First, fork the `continuous-build` GitHub repository to your account, and clone the branch via either:

```
git clone https://www.github.com/<username>/continuous-build
```

or

```
git clone git@github.com:<username>/continuous-build.git
```

9.1.2 Step 2. Choose your Configuration

The hidden folder `.circleci/config.yml` has instructions for [CircleCI](#) to automatically discover and build your repo2docker Jupyter notebook container. The default template provided in the repository in this folder will do the most basic steps, including:

1. Clone the repository with the notebook that you specify
2. Build a Docker image
3. Push the build image to Docker Hub

This repository aims to provide templates for your use. If you have a request for a new template, please [let us know](#). We will add templates as they are requested to do additional tasks like test containers, run nbconvert, etc.

Thus, if I have a repository named `myrepo` and I want to use the default configuration on circleCI, I would copy it there from the `continuous-build` folder. In the example below, I'm creating a new folder called "myrepo" and then copying the entire folder there:

```
mkdir -p myrepo
cp -R continuous-build/.circleci myrepo/
```

You would then logically create a GitHub repository in the "myrepo" folder, add the circleci configuration folder, and continue on to the next steps.

```
cd myrepo
git init
git add .circleci
```

9.1.3 Step 3. Docker Hub

Go to [Docker Hub](#), log in, and click the big blue button that says "create repository" (not an automated build). Choose an organization and name that you like (in the traditional format `<ORG>/<NAME>`), and remember it! We will be adding it, along with your Docker credentials, to be encrypted CircleCI environment variables.

9.1.4 Step 4. Connect to CircleCI

If you navigate to the main [app page](#) you should be able to click "Add Projects" and then select your repository. If you don't see it on the list, then select a different organization in the top left. Once you find the repository, you can click the button to "Start Building" and accept the defaults.

Before you push or trigger a build, let's set up the following environment variables. Also in the project interface on CircleCi, click the gears icon next to the project name to get to your project settings. Under settings, click on the "Environment Variables" tab. In this section, you want to define the following:

1. `CONTAINER_NAME` should be the name of the Docker Hub repository you just created.
2. `DOCKER_TAG` is the tag you want to use. If not defined, will use first 10 characters of commit.
3. `DOCKER_USER` and `DOCKER_PASS` should be your credentials (to allowing pushing)
4. `REPO_NAME` should be the full GitHub url (or other) of the repository with the notebook. This doesn't have to coincide with the repository you are using to do the build (e.g., "myrepo" in our example).

If you don't define the `CONTAINER_NAME` it will default to be the repository where it is building from, which you should only do if the Docker Hub repository is named equivalently. If you don't define either of the variables from step 3. for the Docker credentials, your image will build but not be pushed to Docker Hub. Finally, if you don't define the `REPO_NAME` it will again use the name of the repository defined for the `CONTAINER_NAME`.

9.1.5 Step 5. Push Away, Merrill!

Once the environment variables are set up, you can push or issue a pull request to see circle build the workflow. Remember that you only need the `.circleci/config.yml` and not any other files in the repository. If your notebook is hosted in the same repository, you might want to add these, along with your `requirements.txt`, etc.

Tip: By default, new builds on CircleCI will not build for pull requests and you can change this default in the settings. You can easily add filters (or other criteria and actions) to be performed during or after the build by editing the `.circleci/config.yml` file in your repository.

9.1.6 Step 5. Use Your Container!

You should then be able to pull your new container, and run it! Here is an example:

```
docker pull <ORG>/<NAME>
docker run -it --name repo2docker -p 8888:8888 <ORG>/<NAME> jupyter notebook --ip 0.0.
↪0.0
```

For a pre-built working example, try the following:

```
docker pull vanessa/repo2docker
docker run -it --name repo2docker -p 8888:8888 vanessa/repo2docker jupyter notebook --
↪ip 0.0.0.0
```

You can then enter the url and token provided in the browser to access your notebook. When you are done and need to stop and remove the container:

```
docker stop repo2docker
docker rm repo2docker
```


CHAPTER 10

Configuration Files

`repo2docker` looks for configuration files in the repository being built to determine how to build it. In general, `repo2docker` uses the same configuration files as other software installation tools, rather than creating new custom configuration files.

A number of `repo2docker` configuration files can be combined to compose more complex setups.

The [binder examples](#) organization on GitHub contains a list of sample repositories for common configurations that `repo2docker` can build with various configuration files such as Python and R installation in a repository.

Below is a list of supported configuration files (roughly in the order of build priority):

- `environment.yml` - Install a Python environment
- `requirements.txt` - Install a Python environment
- `setup.py` - Install Python packages
- `Project.toml` - Install a Julia environment
- `REQUIRE` - Install a Julia environment (legacy)
- `install.R` - Install an R/RStudio environment
- `apt.txt` - Install packages with `apt-get`
- `DESCRIPTION` - Install an R package
- `manifest.xml` - Install Stencila
- `postBuild` - Run code after installing the environment
- `start` - Run code before the user sessions starts
- `runtime.txt` - Specifying runtimes
- `default.nix` - the nix package manager
- `Dockerfile` - Advanced environments

10.1 `environment.yml` - Install a Python environment

`environment.yml` is the standard configuration file used by [conda](#) that lets you install any kind of package, including Python, R, and C/C++ packages. `repo2docker` does not use `environment.yml` to create and activate a new conda environment. Rather, it updates a base conda environment with the packages listed in `environment.yml`. This means that the environment will always have the same default name, not the name specified in `environment.yml`.

Note: You can install files from pip in your `environment.yml` as well. For example, see the [binder-examples/environment.yml](#) file.

You can also specify which Python version to install in your built environment with `environment.yml`. By default, `repo2docker` installs **Python 3.7** with your `environment.yml` unless you include the version of Python in this file. `conda` supports all versions of Python, though `repo2docker` support is best with Python 3.7, 3.6, 3.5 and 2.7.

Warning: If you include a Python version in a `runtime.txt` file in addition to your `environment.yml`, your `runtime.txt` will be ignored.

10.2 `requirements.txt` - Install a Python environment

This specifies a list of Python packages that should be installed in your environment. Our [requirements.txt example](#) on GitHub shows a typical requirements file.

10.3 `setup.py` - Install Python packages

To install your repository like a Python package, you may include a `setup.py` file. `repo2docker` installs `setup.py` files by running `pip install -e ..`

10.4 `Project.toml` - Install a Julia environment

A `Project.toml` (or `JuliaProject.toml`) file can specify both the version of Julia to be used and a list of Julia packages to be installed. If a `Manifest.toml` is present, it will determine the exact versions of the Julia packages that are installed.

10.5 `REQUIRE` - Install a Julia environment (legacy)

A `REQUIRE` file can specify both the version of Julia to be used and which Julia packages should be used. The use of `REQUIRE` is only recommended for pre 1.0 Julia versions. The recommended way of installing a Julia environment that uses Julia 1.0 or newer is to use a `Project.toml` file. If both a `REQUIRE` and a `Project.toml` file are detected, the `REQUIRE` file is ignored. To see an example of a Julia repository with `REQUIRE` and `environment.yml`, visit [binder-examples/julia-python](#).

10.6 `install.R` - Install an R/RStudio environment

This is used to install R libraries pinned to a specific snapshot on [MRAN](#). To set the date of the snapshot add a `runtime.txt`. For an example `install.R` file, visit our [example install.R file](#).

10.7 `apt.txt` - Install packages with apt-get

A list of Debian packages that should be installed. The base image used is usually the latest released version of Ubuntu. We use `apt.txt`, for example, to install LaTeX in our [example apt.txt for LaTeX](#).

10.8 `DESCRIPTION` - Install an R package

To install your repository like an R package, you may include a `DESCRIPTION` file. `repo2docker` installs the package and dependencies from the `DESCRIPTION` by running `devtools::install_git(".")`.

You also need to have a `runtime.txt` file that is formatted as `r-<YYYY>-<MM>-<DD>`, where YYYY-MM-DD is a snapshot of MRAN that will be used for your R installation.

10.9 `manifest.xml` - Install Stencila

[Stencila](#) is an open source office suite for reproducible research. It is powered by the open file format [Dar](#).

If your repository contains a Stencila document, `repo2docker` detects it based on the file `manifest.xml`. The required [execution contexts](#) are extracted from a Dar article (i.e. files named `*.jats.xml`).

You may also have a `runtime.txt` and/or an `install.R` to manually configure your R installation.

To see example repositories, visit our [Stencila with R](#) and [Stencila with Python](#) examples.

10.10 `postBuild` - Run code after installing the environment

A script that can contain arbitrary commands to be run after the whole repository has been built. If you want this to be a shell script, make sure the first line is `#!/bin/bash`.

An example use-case of `postBuild` file is JupyterLab's demo on mybinder.org. It uses a `postBuild` file in a folder called `binder` to [prepare their demo for binder](#).

10.11 `start` - Run code before the user sessions starts

A script that can contain simple commands to be run at runtime (as an [ENTRYPOINT](#) to the docker container). If you want this to be a shell script, make sure the first line is `#!/bin/bash`. The last line must be `exec "$@"` or equivalent.

Use this to set environment variables that software installed in your container expects to be set. This script is executed each time your binder is started and should at most take a few seconds to run.

If you only need to run things once during the build phase use [postBuild - Run code after installing the environment](#).

10.12 `runtime.txt` - Specifying runtimes

Sometimes you want to specify the version of the runtime (e.g. the version of Python or R), but the environment specification format don't let you specify this information (e.g. `requirements.txt` or `install.R`). For these cases, we have a special file, `runtime.txt`.

Note: `runtime.txt` is only supported when used with environment specifications that do not already support specifying the runtime (e.g. when using `environment.yml` for conda or `Project.toml` for Julia, `runtime.txt` will be ignored).

To use `python-2.7`: add `python-2.7` in `runtime.txt` file. The repository will run in an env with Python 2 installed. To see a full example repository, visit our [Python2 example](#).

repo2docker uses R libraries pinned to a specific snapshot on [MRAN](#). You need to have a `runtime.txt` file that is formatted as `r-<YYYY>-<MM>-<DD>`, where `YYYY-MM-DD` is a snapshot at MRAN that will be used for installing libraries.

To see an example R repository, visit our [R example in binder-examples](#).

10.13 `default.nix` - the nix package manager

Specify packages to be installed by the [nix package manager](#). When you use this config file all other configuration files (like `requirements.txt`) that specify packages are ignored. When using `nix` you have to specify all packages and dependencies explicitly, including the Jupyter notebook package that repo2docker expects to be installed. If you do not install Jupyter explicitly repo2docker will no be able to start your container.

`nix-shell` is used to evaluate a `nix` expression written in a `default.nix` file. Make sure to [pin your nixpkgs](#) to produce a reproducible environment.

To see an example repository visit [nix binder example](#).

10.14 `Dockerfile` - Advanced environments

In the majority of cases, providing your own `Dockerfile` is not necessary as the base images provide core functionality, compact image sizes, and efficient builds. We recommend trying the other configuration files before deciding to use your own `Dockerfile`.

With `Dockerfiles`, a regular Docker build will be performed.

Note: If a `Dockerfile` is present, all other configuration files will be ignored.

See the [Advanced Binder Documentation](#) for best-practices with `Dockerfiles`.

The Reproducible Execution Environment Specification

repo2docker scans a repository for particular *Configuration Files*, such as `requirements.txt` or `REQUIRE`. The collection of files, their contents, and the resulting actions that repo2docker takes is known as the **Reproducible Execution Environment Specification** (or REES).

The goal of the REES is to automate and encourage existing community best practices for reproducible computational environments. This includes installing packages using community-standard specification files and their corresponding tools, such as `requirements.txt` (with `pip`), `REQUIRE` (with Julia), or `apt.txt` (with `apt`). While repo2docker automates the creation of the environment, a human should be able to look at a REES-compliant repository and reproduce the environment using common, clear steps without repo2docker software.

Currently, the definition of the REE Specification is the following:

Any directory containing zero or more files from the *Configuration Files* list is a valid reproducible execution environment as defined by the REES. The configuration files have to all be placed either in the root of the directory, in a `binder/` sub-directory or a `.binder/` sub-directory.

For example, the REES recognises `requirements.txt` as a valid config file. The file format is as defined by the `requirements.txt` standard of the Python community. A REES-compliant tool will install a Python interpreter (of unspecified version) and perform the equivalent action of `pip install -r requirements.txt` so that the user can afterwards run python and use the packages installed.

Contributing to repo2docker development

Thank you for thinking about contributing to repo2docker! This is an open source project that is developed and maintained entirely by volunteers. *Your contribution* is integral to the future of the project. THANK YOU!

12.1 Types of contribution

There are many ways to contribute to repo2docker:

- **Update the documentation.** If you're reading a page or docstring and it doesn't make sense (or doesn't exist!), please let us know by opening a bug report. It's even more amazing if you can give us a suggested change.
- **Fix bugs or add requested features.** Have a look through the [issue tracker](#) and see if there are any tagged as "help wanted". As the label suggests, we'd love your help!
- **Report a bug.** If repo2docker isn't doing what you thought it would do then open a [bug report](#). That issue template will ask you a few questions described in more detail below.
- **Suggest a new feature.** We know that there are lots of ways to extend repo2docker! If you're interested in adding a feature then please open a [feature request](#). That issue template will ask you a few questions described in detail below.
- **Review someone's Pull Request.** Whenever somebody proposes changes to the repo2docker codebase, the community reviews the changes, and provides feedback, edits, and suggestions. Check out the [open pull requests](#) and provide feedback that helps improve the PR and get it merged. Please keep your feedback positive and constructive!
- **Tell people about repo2docker.** As we said above, repo2docker is built by and for its community. If you know anyone who would like to use repo2docker, please tell them about the project! You could give a talk about it, or run a demonstration. The sky is the limit :rocket::star2:.

If you're not sure where to get started, then please come and say hello in our [Gitter channel](#), or open an discussion thread at the [Jupyter discourse forum](#).

12.2 Process for making a contribution

This outlines the process for getting changes to the repo2docker project merged.

1. Identify the correct issue template: [bug report](#) or [feature request](#).

Bug reports ([examples](#), [new issue](#)) will ask you for a description of the problem, the expected behaviour, the actual behaviour, how to reproduce the problem, and your personal set up. Bugs can include problems with the documentation, or code not running as expected.

It is really important that you make it easy for the maintainers to reproduce the problem you're having. This guide on creating a [minimal, complete and verifiable example](#) is a great place to start.

Feature requests ([examples](#), [new issue](#)) will ask you for the proposed change, any alternatives that you have considered, a description of who would use this feature, and a best-guess of how much work it will take and what skills are required to accomplish.

Very easy feature requests might be updates to the documentation to clarify steps for new users. Harder feature requests may be to add new functionality to the project and will need more in depth discussion about who can complete and maintain the work.

Feature requests are a great opportunity for you to advocate for the use case you're suggesting. They help others understand how much effort it would be to integrate the work, and - if you're successful at convincing them that this effort is worth it - make it more likely that they to choose to work on it with you.

2. Open an issue. Getting consensus with the community is a great way to save time later.
3. Make edits in [your fork](#) of the [repo2docker repository](#).
4. Make a [pull request](#). Read the [next section](#) for guidelines for both reviewers and contributors on merging a PR.
5. Edit [the changelog](#) by appending your feature / bug fix to the development version.
6. Wait for a community member to merge your changes. Remember that **someone else must merge your pull request**. That goes for new contributors and long term maintainers alike.
7. (optional) Deploy a new version of repo2docker to mybinder.org by [following these steps](#)

12.3 Guidelines to getting a Pull Request merged

These are not hard rules to be enforced by but they are suggestions written by the repo2docker maintainers to help complete your contribution as smoothly as possible for both you and for them.

- **Create a PR as early as possible**, marking it with [WIP] while you work on it. This avoids duplicated work, lets you get high level feedback on functionality or API changes, and/or helps find collaborators to work with you.
- **Keep your PR focused**. The best PRs solve one problem. If you end up changing multiple things, please open separate PRs for the different conceptual changes.
- **Add tests to your code**. PRs will not be merged if Travis is failing.
- **Apply PEP8** as much as possible, but not too much. If in doubt, ask.
- **Use merge commits** instead of merge-by-squashing/-rebasing. This makes it easier to find all changes since the last deployment `git log --merges --pretty=format:"%h %<(10,trunc)%an %<(15)%ar %s" <deployed-revision>..` and your PR easier to review.
- **Make it clear when your PR is ready for review**. Prefix the title of your pull request (PR) with [MRG] if the contribution is complete and should be subjected to a detailed review.

- Enter your changes into the *changelog* in docs/source/changelog.rst.
- Use commit messages to describe *why* you are proposing the changes you are proposing.
- Try to not rush changes (the definition of rush depends on how big your changes are). Remember that everyone in the repo2docker team is a volunteer and we can not (nor would we want to) control their time or interests. Wait patiently for a reviewer to merge the PR. (Remember that **someone else** must merge your PR, even if you have the admin rights to do so.)

12.4 Setting up for Local Development

To develop & test repo2docker locally, you need:

1. Familiarity with using a command line terminal
2. A computer running macOS / Linux
3. Some knowledge of git
4. At least python 3.6
5. Your favorite text editor
6. A recent version of [Docker Community Edition](#)

12.4.1 Clone the repository

First, you need to get a copy of the repo2docker git repository on your local disk. Fork the repository on GitHub, then clone it to your computer:

```
git clone https://github.com/<your-username>/repo2docker
```

This will clone repo2docker into a directory called `repo2docker`. You can make that your current directory with `cd repo2docker`.

12.4.2 Set up a local virtual environment

After cloning the repository (or your fork of the repository), you should set up an isolated environment to install libraries required for running / developing repo2docker.

There are many ways to do this but here we present you with two approaches: virtual environment or pipenv.

- Using virtual environment

```
python3 -m venv .
source bin/activate
pip3 install -e .
pip3 install -r dev-requirements.txt
pip3 install -r docs/doc-requirements.txt
```

This should install all the libraries required for testing & running repo2docker!

- Using pipenv

Note that you will need to install pipenv first using `pip3 install pipenv`. Then from the root directory of this project you can use the following commands:

```
pipenv install --dev
```

This should install both the dev and docs requirements at once!

12.4.3 Set up

12.4.4 Verify that docker is installed and running

If you do not already have [Docker](#), you should be able to download and install it for your operating system using the links from the [official website](#). After you have installed it, you can verify that it is working by running the following commands:

```
docker version
```

It should output something like:

```
Client:
 Version:      17.09.0-ce
 API version:  1.32
 Go version:   go1.8.3
 Git commit:   afdb6d4
 Built:        Tue Sep 26 22:42:45 2017
 OS/Arch:      linux/amd64

Server:
 Version:      17.09.0-ce
 API version:  1.32 (minimum version 1.12)
 Go version:   go1.8.3
 Git commit:   afdb6d4
 Built:        Tue Sep 26 22:41:24 2017
 OS/Arch:      linux/amd64
 Experimental: false
```

Then you are good to go!

The repo2docker roadmap

This roadmap collects “next steps” for repo2docker. It is about creating a shared understanding of the project’s vision and direction amongst the community of users, contributors, and maintainers. The goal is to communicate priorities and upcoming release plans. It is not aimed at limiting contributions to what is listed here.

13.1 Using the roadmap

13.1.1 Sharing Feedback on the Roadmap

All of the community is encouraged to provide feedback as well as share new ideas with the community. Please do so by submitting an issue. If you want to have an informal conversation first use one of the other communication channels. After submitting the issue, others from the community will probably respond with questions or comments they have to clarify the issue. The maintainers will help identify what a good next step is for the issue.

13.1.2 What do we mean by “next step”?

When submitting an issue, think about what “next step” category best describes your issue:

- **now**, concrete/actionable step that is ready for someone to start work on. These might be items that have a link to an issue or more abstract like “decrease typos and dead links in the documentation”
- **soon**, less concrete/actionable step that is going to happen soon, discussions around the topic are coming close to an end at which point it can move into the “now” category
- **later**, abstract ideas or tasks, need a lot of discussion or experimentation to shape the idea so that it can be executed. Can also contain concrete/actionable steps that have been postponed on purpose (these are steps that could be in “now” but the decision was taken to work on them later)

13.1.3 Reviewing and Updating the Roadmap

The roadmap will get updated as time passes (next review by 31st January 2019) based on discussions and ideas captured as issues. This means this list should not be exhaustive, it should only represent the “top of the stack” of ideas. It should not function as a wish list, collection of feature requests or todo list. For those please create a [new issue](#).

The roadmap should give the reader an idea of what is happening next, what needs input and discussion before it can happen and what has been postponed.

13.2 The roadmap proper

13.2.1 Project vision

Repo2docker is a dependable tool used by humans that reduces the complexity of creating the environment in which a piece of software can be executed.

13.2.2 Now

The “Now” items are being actively worked on by the project:

- reduce documentation typos and syntax errors
- increase test coverage to 80% (see <https://codecov.io/gh/jupyter/repo2docker/tree/master/repo2docker> for low coverage files)
- mounting repository contents in locations that is not `/home/jovyan`
- investigate options for pinning repo2docker versions ([#490](#))

13.2.3 Soon

The “Soon” items are being discussed/a plan of action is being made. Once an item reaches the point of an actionable plan and person who wants to work on it, the item will be moved to the “Now” section. Typically, these will be moved at a future review of the roadmap.

- create the contributor highway, define the route from newcomer to project lead
- add Julia Manifest support (<https://docs.julialang.org/en/v1/stdlib/Pkg/index.html>, [#486](#))
- support different base images/build pack stacks ([#487](#))

13.2.4 Later

The “Later” items are things that are at the back of the project’s mind. At this time there is no active plan for an item. The project would like to find the resources and time to discuss and then execute these ideas.

- support execution on a remote host (with more resources than available locally) via the command-line
- add support for using ZIP files as the repo (repo2docker `https://example.com/an-archive.zip`) this will give us access to several archives (like Zenodo) that expose things as ZIP files.
- add support for Zenodo (repo2docker `10.5281/zenodo.1476680`) so Zenodo software archives can be used as the source in addition to a git repository

Architecture of repo2docker

This is a living document talking about the architecture of repo2docker from various perspectives.

14.1 Buildpack

The **buildpack** concept comes from [Heroku](#) and Ruby on Rails' [Convention over Configuration](#) doctrine.

Instead of the user specifying a complete specification of exactly how they want their environment to be, they can focus only on how their environment differs from a conventional environment. This means instead of deciding 'should I get Python from Apt or pyenv or ?', user can just specify 'I want python-3.6'. Usually, specifying a **runtime** and list of **libraries** with explicit **versions** is all that is needed.

In repo2docker, a Buildpack does the following things:

1. **Detect** if it can handle a given repository
2. **Build** a base language environment in the docker image
3. **Copy** the contents of the repository into the docker image
4. **Assemble** a specific environment in the docker image based on repository contents
5. **Push** the built docker image to a specific docker registry (optional)
6. **Run** the build docker image as a docker container (optional)

14.1.1 Detect

When given a repository, repo2docker first has to determine which buildpack to use. It takes the following steps to determine this:

1. Look at the ordered list of `BuildPack` objects listed in `Repo2Docker.buildpacks` traitlet. This is populated with a default set of buildpacks in most-specific-to-least-specific order. Other applications using this can add / change this using traditional [traitlet](#) configuration mechanisms.

2. Calls the `detect` method of each `BuildPack` object. This method assumes that the repository is present in the current working directory, and should return `True` if the repository is something that it should be used for. For example, a `BuildPack` that uses `conda` to install libraries can check for presence of an `environment.yml` file and say ‘yes, I can handle this repository’ by returning `True`. Usually buildpacks look for presence of specific files (`requirements.txt`, `environment.yml`, `install.R`, `manifest.xml` etc) to determine if they can handle a repository or not. Buildpacks may also look into specific files to determine specifics of the required environment, such as the Stencila integration which extracts the required language-specific executions contexts from an XML file (see base `BuildPack`). More than one buildpack may use such information, as properties can be inherited (e.g. the R buildpack uses the list of required Stencila contexts to see if R must be installed).
3. If no `BuildPack` returns true, then `repo2docker` will use the default `BuildPack` (defined in `Repo2Docker.default_buildpack_traitlet`).

14.2 Build base environment

Once a buildpack is chosen, it builds a **base environment** that is mostly the same for various repositories built with the same buildpack.

For example, in `CondaBuildPack`, the base environment consists of installing `miniconda` and basic notebook packages (from `repo2docker/buildpacks/conda/environment.yml`). This is going to be the same for most repositories built with `CondaBuildPack`, so we want to use `docker layer caching` as much as possible for performance reasons. Next time a repository is built with `CondaBuildPack`, we can skip straight to the **copy** step (since the base environment docker image *layers* have already been built and cached).

The `get_build_scripts` and `get_build_script_files` methods are primarily used for this. `get_build_scripts` can return arbitrary bash script lines that can be run as different users, and `get_build_script_files` is used to copy specific scripts (such as a conda installer) into the image to be run as part of `get_build_scripts`. Code in either has following constraints:

1. You can *not* use the contents of repository in them, since this happens before the repository is copied into the image. For example, `pip install -r requirements.txt` will not work, since there’s no `requirements.txt` inside the image at this point. This is an explicit design decision, to enable better layer caching.
2. You *may*, however, read the contents of the repository and modify the scripts emitted based on that! For example, in `CondaBuildPack`, if there’s Python 2 specified in `environment.yml`, a different kind of environment is set up. The reading of the `environment.yml` is performed in the `BuildPack` itself, and not in the scripts returned by `get_build_scripts`. This is fine. `BuildPack` authors should still try to minimize the variants created in this fashion, to optimize the build cache.

14.3 Copy repository contents

The contents of the repository are copied unconditionally into the Docker image, and made available for all further commands. This is common to most `BuildPacks`, and the code is in the `build` method of the `BuildPack` base class.

14.4 Assemble repository environment

The **assemble** stage builds the specific environment that is requested by the repository. This usually means installing required libraries specified in a format native to the language (`requirements.txt`, `environment.yml`, `REQUIRE`, `install.R`, etc).

Most of this work is done in `get_assemble_scripts` method. It can return arbitrary bash script lines that can be run as different users, and has access to the repository contents (unlike `get_build_scripts`). The docker image layers produced by this usually can not be cached, so less restrictions apply to this than to `get_build_scripts`.

At the end of the assemble step, the docker image is ready to be used in various ways!

14.5 Push

Optionally, repo2docker can **push** a built image to a [docker registry](#). This is done as a convenience only (since you can do the same with a `docker push` after using repo2docker only to build), and implemented in `Repo2Docker.push` method. It is only activated if using the `--push` commandline flag.

14.6 Run

Optionally, repo2docker can **run** the built image and allow the user to access the Jupyter Notebook running inside by default. This is also done as a convenience only (since you can do the same with `docker run` after using repo2docker only to build), and implemented in `Repo2Docker.run`. It is activated by default unless the `--no-run` commandline flag is passed.

Design of repo2docker

The `repo2docker` buildpacks are inspired by [Heroku's Build Packs](#). The philosophy for the `repo2docker` buildpacks includes:

- using common configuration files for familiar installation and packaging tools
- allowing configuration files to be combined to compose more complex setups
- specifying default locations for configuration files (in the repository's root, `binder` or `.binder` directory)

When designing `repo2docker` and adding to it in the future, the developers are influenced by two primary use cases. The use cases for `repo2docker` which drive most design decisions are:

1. Automated image building used by projects like [BinderHub](#)
2. Manual image building and running the image from the command line client, `jupyter-repo2docker`, by users interactively on their workstations

15.1 Deterministic output

The core of `repo2docker` can be considered a [deterministic algorithm](#). When given an input directory which has a particular repository checked out, it deterministically produces a Dockerfile based on the contents of the directory. So if we run `repo2docker` on the same directory multiple times, we get the exact same Dockerfile output.

This provides a few advantages:

1. Reuse of cached built artifacts based on a repository's identity increases efficiency and reliability. For example, if we had already run `repo2docker` on a git repository at a particular commit hash, we know we can just re-use the old output, since we know it is going to be the same. This provides massive performance & architectural advantages when building additional tools (like [BinderHub](#)) on top of `repo2docker`.
2. We produce Dockerfiles that have as much in common as possible across multiple repositories, enabling better use of the Docker build cache. This also provides massive performance advantages.

15.2 Reproducibility and version stability

Many ingredients go into making an image from a repository:

1. version of the base docker image
2. version of `repo2docker` itself
3. versions of the libraries installed by the repository

`repo2docker` controls the first two, the user controls the third one. The current policy for the version of the base image is that we will keep pace with Ubuntu releases until we reach the next release with Long Term Support (LTS). We currently use Artful Aardvark (17.10) and the next LTS version will be Bionic Beaver (18.04).

The version of `repo2docker` used to build an image can influence which packages are installed by default and which features are supported during the build process. We will periodically update those packages to keep step with releases of Jupyter Notebook, JupyterLab, etc. For packages that are installed by default but where you want to control the version we recommend you specify them explicitly in your dependencies.

15.3 Unix principles “do one thing well”

`repo2docker` should do one thing, and do it well. This one thing is:

Given a repository, deterministically build a docker image from it.

There’s also some convenience code (to run the built image) for users, but that’s separated out cleanly. This allows easy use by other projects (like BinderHub).

There is additional (and very useful) design advice on this in the [Art of Unix Programming](#) which is a highly recommended quick read.

15.4 Composability

Although other projects, like `s2i`, exist to convert source to Docker images, `repo2docker` provides the additional functionality to support *composable* environments. We want to easily have an image with Python3+Julia+R-3.2 environments, rather than just one single language environment. While generally one language environment per container works well, in many scientific / datascience computing environments you need multiple languages working together to get anything done. So all buildpacks are composable, and need to be able to work well with other languages.

15.5 Pareto principle (The 80-20 Rule)

Roughly speaking, we want to support 80% of use cases, and provide an escape hatch (raw Dockerfiles) for the other 20%. We explicitly want to provide support only for the most common use cases - covering every possible use case never ends well.

An easy process for getting support for more languages here is to demonstrate their value with Dockerfiles that other people can use, and then show that this pattern is popular enough to be included inside `repo2docker`. Remember that ‘yes’ is forever (very hard to remove features!), but ‘no’ is only temporary!

CHAPTER 16

Common tasks

These are some common tasks to be done as a part of developing and maintaining `repo2docker`. If you'd like more guidance for how to do these things, reach out in the [JupyterHub Gitter channel](#).

16.1 Running tests

We have a lot of tests for various cases supported by `repo2docker` in the `tests/` subdirectory. If you fix a bug or add new functionality consider adding a new test to prevent the bug from coming back. These use `py.test`.

You can run all the tests with:

```
py.test -s tests/*
```

If you want to run a specific test, you can do so with:

```
py.test -s tests/<path-to-test>
```

16.2 Update and Freeze BuildPack Dependencies

This section covers the process by which `repo2docker` defines and updates the dependencies that are installed by default for several buildpacks.

For both the `conda` and `virtualenv` (pip) base environments in the **Conda BuildPack** and **Python BuildPack**, we install specific pinned versions of all dependencies. We explicitly list the dependencies we want, then *freeze* them at commit time to explicitly list all the transitive dependencies at current versions. This way, we know that all dependencies will have the exact same version installed at all times.

To update one of the dependencies shared across all `repo2docker` builds, you must follow these steps (with more detailed information in the sections below):

1. Make sure you have [Docker](#) running on your computer

2. Bump the version numbers of the dependencies you want to update in the `conda` environment ([link](#))
3. Make a pull request with your changes ([link](#))

See the subsections below for more detailed instructions.

16.2.1 Conda dependencies

1. There are two files related to conda dependencies. Edit as needed.
 - `repo2docker/buildpacks/conda/environment.yml`
Contains list of packages to install in Python3 conda environments, which are the default. **This is where all Notebook versions & notebook extensions (such as JupyterLab / nteract) go.**
 - `repo2docker/buildpacks/conda/environment.py-2.7.yml`
Contains list of packages to install in Python2 conda environments, which can be specifically requested by users. **This only needs IPyKernel and kernel related libraries.** Notebook / Notebook Extension need not be installed here.
2. Once you edit either of these files to add a new package / bump version on an existing package, you should then run:

```
cd ./repo2docker/buildpacks/conda/  
python freeze.py
```

This script will resolve dependencies and write them to the respective `.frozen.yml` files. You will need `docker` installed to run this script.

3. After the freeze script finishes, a number of files will have been created. Commit the following subset of files to git:

```
repo2docker/buildpacks/conda/environment.yml  
repo2docker/buildpacks/conda/environment.frozen.yml  
repo2docker/buildpacks/conda/environment.py-2.7.yml  
repo2docker/buildpacks/conda/environment.py-2.7.frozen.yml  
repo2docker/buildpacks/conda/environment.py-3.5.frozen.yml  
repo2docker/buildpacks/conda/environment.py-3.6.frozen.yml
```

4. Make a pull request; see details below.
5. Once the pull request is approved (but not yet merged), Update the change log (details below) and commit the change log, then update the pull request.

16.2.2 Make a Pull Request

Once you've made the commit, please make a Pull Request to the `jupyterhub/repo2docker` repository, with a description of what versions were bumped / what new packages were added and why. If you fix a bug or add new functionality consider adding a new test to prevent the bug from coming back/the feature breaking in the future.

16.3 Creating a Release

We try to make a release of `repo2docker` every few months if possible.

We follow semantic versioning.

Check that the Change log is ready and then tag a new release locally:

```
V=0.7.0 git tag -am "release $V" $V
git push origin --tags
```

When the travis run completes check that the new release is available on PyPI.

16.3.1 Update the change log

To add your change to the change log, find the relevant Feature/Bug fix/API change section for the next release near the top of the file; then add one or two sentences as a new bullet point about your changes. Include the pull request or issue number between square brackets at the end.

Some details:

- versioning follows the x.y.z, major.minor.bugfix numbering
- bug fixes go into the next bugfix release. If there isn't any, you can create a new section (see point below). Don't worry if you're not sure about that, and think it should go into a next major or minor release: an admin will let you know, or move the change later to the appropriate section
- API changes should preferably go into the next major release, unless they are backward compatible (for example, a deprecated function keyword): then they can go into the next minor release. For release with major release 0, non-backward compatible breaking changes are also fine for the next minor release.
- new features should go into the next minor release.
- if there is no section for the appropriate release, you can add one:

follow the versioning scheme, by simply increasing the relevant number for one of the major /minor/bugfix numbers, appropriate for your change (see the above bullet points); add the release section. Then add three subsections: new features, api changes, and bug fixes. Leave out the sections that are not appropriate for the newly added release section.

Release candidate versions in the change log are only temporary, and should be superseded by either a next release candidate, or the final release for that version (bugfix version 0).

16.3.2 Keeping the Pipfile and requirements files up to date

We now have both a `dev-requirements.txt` and a `Pipfile` for repo2docker, as such it is important to keep these in sync/up-to-date.

Both files use pip identifiers so if you are updating for example the Sphinx version in the `doc-requirements.txt` (currently `Sphinx = ">=1.4, !=1.5.4"`) you can use the same syntax to update the `Pipfile` and viceversa.

At the moment this has to be done manually so please make sure to update both files accordingly.

Adding a new buildpack to repo2docker

A new buildpack is needed when a new language or a new package manager should be supported. Existing buildpacks are a good model for how new buildpacks should be structured.

17.1 Criteria to balance and consider

Criteria to balance are:

1. Maintenance burden on repo2docker.
2. How easy it is to use a given setup without support from repo2docker natively. There are two escape hatches here - `postBuild` and `Dockerfile`.
3. How widely used is this language / package manager? This is the primary tradeoff with point (1). We (the Binder / Jupyter team) want to make new formats as little as possible, so ideally we can just say “X repositories on binder already use this using one of the escape hatches in (2), so let us make it easy and add native support”.

17.2 Adding libraries or UI to existing buildpacks

Note that this doesn’t apply to adding additional libraries / UI to existing buildpacks. For example, if we had an R buildpack and it supported IRKernel, it is much easier to just support RStudio / Shiny with it, since those are library additions instead of entirely new buildpacks.

18.1 Version x.x.x

Release date: TBD

18.1.1 New features

18.1.2 API changes

18.1.3 Bug fixes

18.2 Version 0.9.0

Release date: 2019-05-05

18.2.1 New features

- Support for julia *Project.toml*, *JuliaProject.toml* and *Manifest.toml* files in [PR #595](#) by [@davidanthoff](#)
- Set JULIA_PROJECT globally, so that every julia instance starts with the julia environment activated in [PR #612](#) by [@davidanthoff](#).
- Update Miniconda version to 4.6.14 and Conda version to 4.6.14 in [PR #637](#) by [@jhamman](#)
- Install notebook into *notebook* env instead of *root*. Activate conda environments and shell integration via ENTRYPOINT in [PR #651](#) by [@minrk](#)
- Support for *.binder* directory in addition to *binder* directory for location of configuration files, in [PR #653](#) by [@jhamman](#).
- Updated contributor guide and issue templates for bugs, feature requests, and support questions in [PR #654](#) and [PR #655](#) by [@KirstieJane](#) and [@betatim](#).

- Create a page naming and describing the “Reproducible Execution Environment Specification” (the specification used by repo2docker) in [PR #662](#) by [@choldgraf](#).

18.2.2 API changes

18.2.3 Bug fixes

- Install IJulia kernel into `${NB_PYTHON_PREFIX}/share/jupyter` in [PR #622](#) by [@davidanthoff](#).
- Ensure git submodules are updated and initialized correctly in [PR #639](#) by [@djhoese](#).
- Use [archive.debian.org](#) as source for the debian jessie based legacy buildpack in [PR #633](#) by [@betatim](#).
- Update to version 5.7.6 of the *notebook* package used in all environments in [PR #628](#) by [@betatim](#).
- Update to version 5.7.8 of the *notebook* package and version 2.0.12 of *nteract-on-jupyter* in [PR #650](#) by [@betatim](#).
- Switch to newer version of jupyter-server-proxy to fix websocket handling in [PR #646](#) by [@betatim](#).
- Update to pip version 19.0.3 in [PR #647](#) by [@betatim](#).
- Ensure ENTRYPOINT is an absolute path in [PR #657](#) by [@yuvipanda](#).
- Fix handling of `-build-memory-limit` values without a postfix in [PR #652](#) by [@betatim](#).

18.3 Version 0.8.0

Release date: 2019-02-21

18.3.1 New features

- Add additional metadata to docker images about how they were built [PR #500](#) by [@jrbourbeau](#).
- Allow users to install global NPM packages: [PR #573](#) by [@GladysNalvarte](#).
- Add documentation on switching the user interface presented by a container. [PR #568](#) by user:*choldgraf*.
- Increased test coverage to ~87% by [@betatim](#) and [@yuvipanda](#).
- Documentation improvements and additions by [@lheagy](#), [@choldgraf](#).
- Remove f-strings from code base, repo2docker is compatible with Python 3.4+ again by [@jrbourbeau](#) in [PR #520](#).
- Local caching of previously built repostories to speed up launch times by [@betatim](#) in [PR #511](#).
- Make destination of repository content in the container image configurable on the CLI via `--target-repo-dir`. By [@yuvipanda](#) in [PR #507](#).
- Expose CPU limit settings for building and running containers. By [@GladysNalvarte](#) in [PR #579](#).
- Make Python 3.7 the default version. By [@yuvipanda](#) and [@minrk](#) in [PR #539](#).

18.3.2 API changes

18.3.3 Bug fixes

- In some cases the version of conda installed in images was not pinned and got upgraded by user actions. Fixed in [PR #576](#) by [@minrk](#).
- Fix an error related to checking if debug output was enabled or not: [PR #575](#) by [@yuvipanda](#).
- Update interact frontend to version 2.0.0 by [@yuvipanda](#) in [PR #571](#).
- Fix quoting issue in `GIT_CREDENTIAL_ENV` environment variable by [@minrk](#) in [PR #572](#).
- Change to using the first 8 characters of each Git commit, not the last 8, to tag each built docker image of repo2docker itself. [@minrk](#) in [PR #562](#).
- Allow users to select the Julia when using a `requirements.txt` by [@yuvipanda](#) in [PR #557](#).
- Set `JULIA_DEPOT_PATH` to install packages outside the home directory by [@yuvipanda](#) in [PR #555](#).
- Update to Jupyter notebook 5.7.4 [PR #519](#) by [@minrk](#).

18.4 Version 0.7.0

Release date: 2018-12-12

18.4.1 New features

- Build from sub-directory: build the image based on a sub-directory of a repository [PR #413](#) by [@dsludwig](#).
- Editable mode: allows editing a local repository from a live container [PR #421](#) by [@evertrol](#).
- Change log added [PR #426](#) by [@evertrol](#).
- Documentation: improved the documentation for contributors [PR #453](#) by [@choldgraf](#).
- Buildpack: added support for the nix package manager [PR #407](#) by [@costrouc](#).
- Log a ‘success’ message when push is complete [PR #482](#) by [@yuvipanda](#).
- Allow specifying images to reuse cache from [PR #478](#) by [@yuvipanda](#).
- Add JupyterHub back to base environment [PR #476](#) by [@yuvipanda](#).
- Repo2docker has a logo! by [@agahkarakuzu](#) and [@blairhudson](#).
- Improve support for Stencila, including identifying stencila runtime from document context [PR #457](#) by [@nuest](#).

18.4.2 API changes

- Add content provider abstraction [PR #421](#) by [@betatim](#).

18.4.3 Bug fixes

- Update to Jupyter notebook 5.7 [PR #475](#) by [@betatim](#) and [@minrk](#).

18.5 Version 0.6

Released 2018-09-09

18.6 Version 0.5

Released 2018-02-07

18.7 Version 0.4.1

Released 2018-09-06

18.8 Version 0.2

Released 2018-05-25

18.9 Version 0.1.1

Released 2017-04-19

18.10 Version 0.1

Released 2017-04-14

Symbols

- appendix <appendix>
 - jupyter-repo2docker command line option, 8
 - build-memory-limit
 - <build_memory_limit>
 - jupyter-repo2docker command line option, 7
 - cache-from <cache_from>
 - jupyter-repo2docker command line option, 8
 - config <config>
 - jupyter-repo2docker command line option, 7
 - debug
 - jupyter-repo2docker command line option, 7
 - editable, -E
 - jupyter-repo2docker command line option, 8
 - env <environment>, -e <environment>
 - jupyter-repo2docker command line option, 8
 - image-name <image_name>
 - jupyter-repo2docker command line option, 7
 - json-logs
 - jupyter-repo2docker command line option, 7
 - no-build
 - jupyter-repo2docker command line option, 7
 - no-clean
 - jupyter-repo2docker command line option, 8
 - no-run
 - jupyter-repo2docker command line option, 7
 - publish <ports>, -p <ports>
 - jupyter-repo2docker command line option, 7
 - publish-all, -P
 - jupyter-repo2docker command line option, 8
 - push
 - jupyter-repo2docker command line option, 8
 - ref <ref>
 - jupyter-repo2docker command line option, 7
 - subdir <subdir>
 - jupyter-repo2docker command line option, 8
 - target-repo-dir <target_repo_dir>
 - jupyter-repo2docker command line option, 8
 - user-id <user_id>
 - jupyter-repo2docker command line option, 8
 - user-name <user_name>
 - jupyter-repo2docker command line option, 8
 - version
 - jupyter-repo2docker command line option, 8
 - volume <volumes>, -v <volumes>
 - jupyter-repo2docker command line option, 8
 - h, -help
 - jupyter-repo2docker command line option, 7
-
- ## C
- cmd
 - jupyter-repo2docker command line option, 7
-
- ## J
- jupyter-repo2docker command line

- option
- appendix <appendix>, 8
- build-memory-limit
 - <build_memory_limit>, 7
- cache-from <cache_from>, 8
- config <config>, 7
- debug, 7
- editable, -E, 8
- env <environment>, -e
 - <environment>, 8
- image-name <image_name>, 7
- json-logs, 7
- no-build, 7
- no-clean, 8
- no-run, 7
- publish <ports>, -p <ports>, 7
- publish-all, -P, 8
- push, 8
- ref <ref>, 7
- subdir <subdir>, 8
- target-repo-dir <target_repo_dir>,
8
- user-id <user_id>, 8
- user-name <user_name>, 8
- version, 8
- volume <volumes>, -v <volumes>, 8
- h, -help, 7
- cmd, 7
- repo, 7

R

repo

- jupyter-repo2docker command line
 - option, 7