

moz://a

The
Alan Turing
Institute

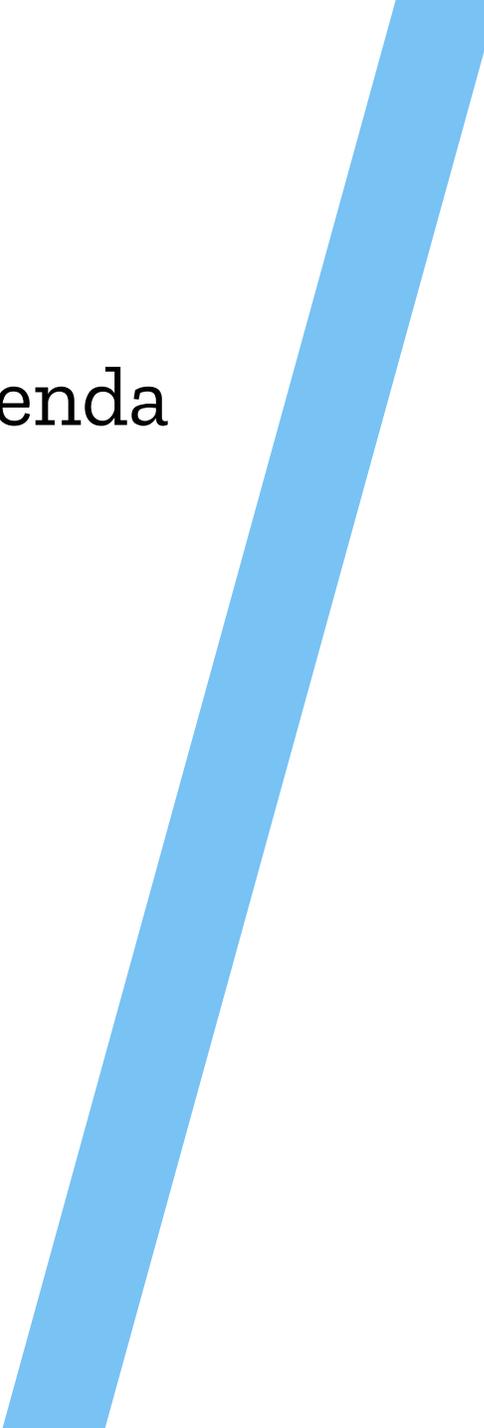
Introduction to Binder(Hub)

and why YOU should build a BinderHub

2019.03.18

Alex Morley

Mozilla Fellow, contributor to The Turing Way



Agenda

1. Binder

What it is and why it's super awesome.

2. Binderhub

How it works and why it works that way.

3. Workshop Goals

Why YOU should build a BinderHub and what everyone wants to get out of the workshop.

Market Research for Binder

(done reproducibly)

Indicate how much you agree with the following statements
using traffic lights.

Practice:

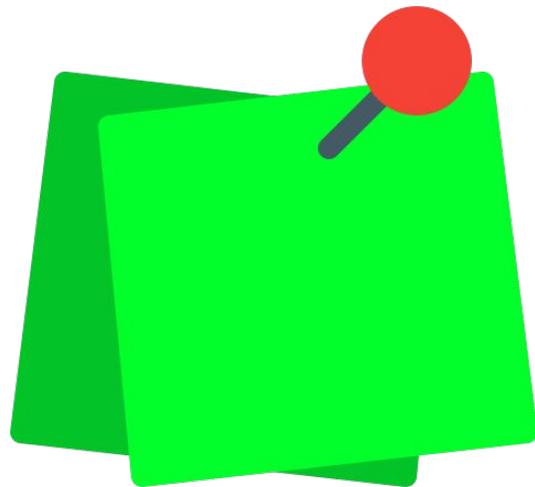
Market Research for Binder

(done reproducibly)

Indicate how much you agree with the following statements
using traffic lights.

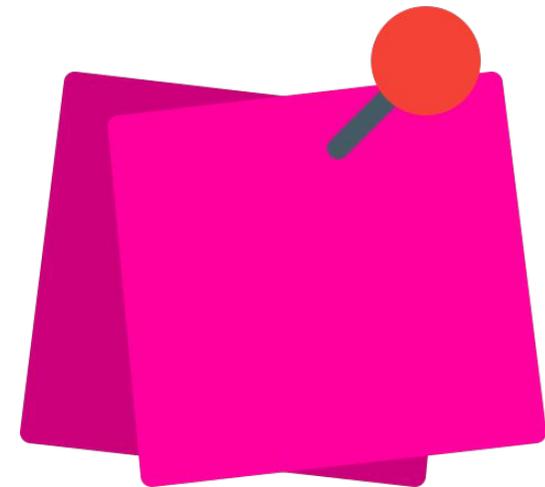
Practice:

I feel fully caffeinated this morning.



Agree

Disagree

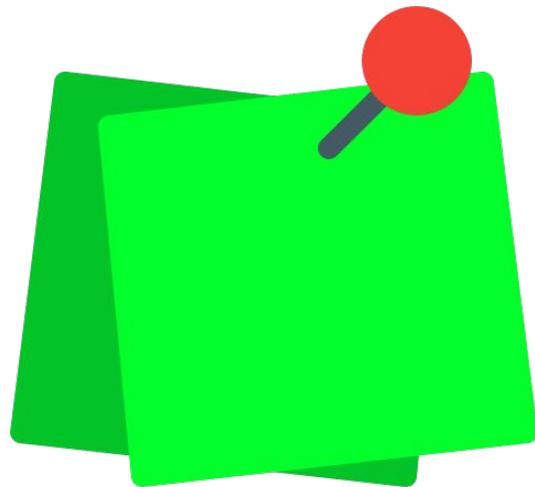


Market Research for Binder

(done reproducibly)

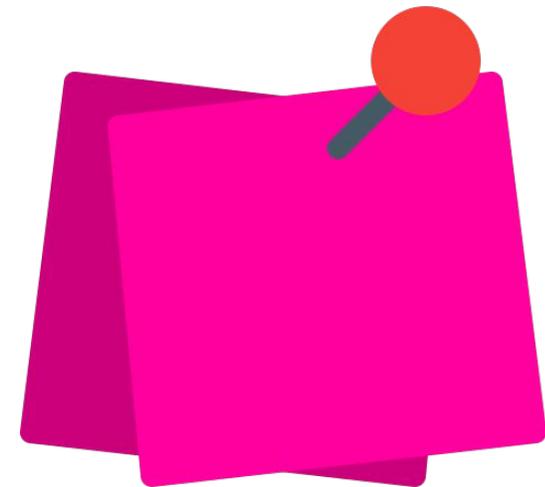
Indicate how much you agree with the following statements
using traffic lights.

Today I am using: Windows



Agree

Disagree

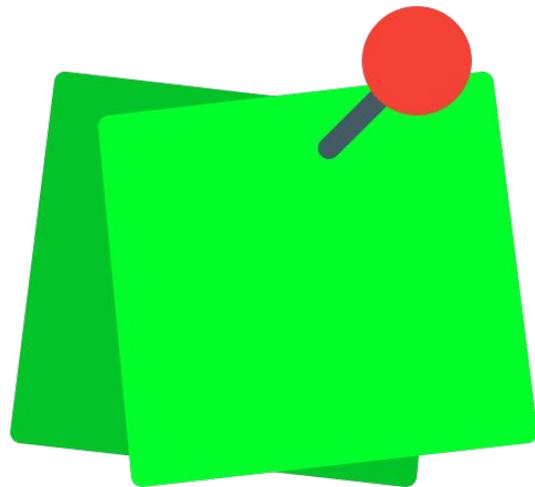


Market Research for Binder

(done reproducibly)

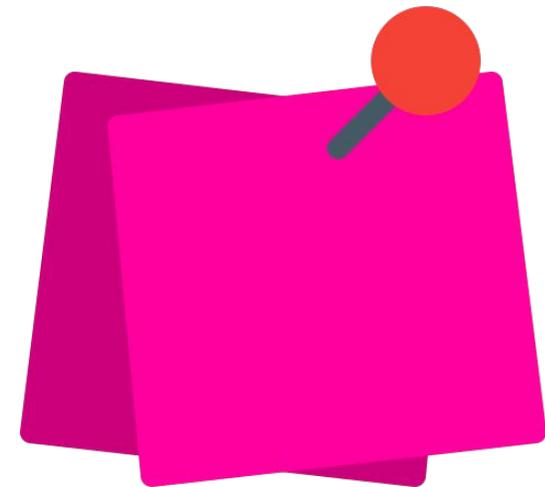
Indicate how much you agree with the following statements
using traffic lights.

Today I am using: Mac (OSX)



Agree

Disagree

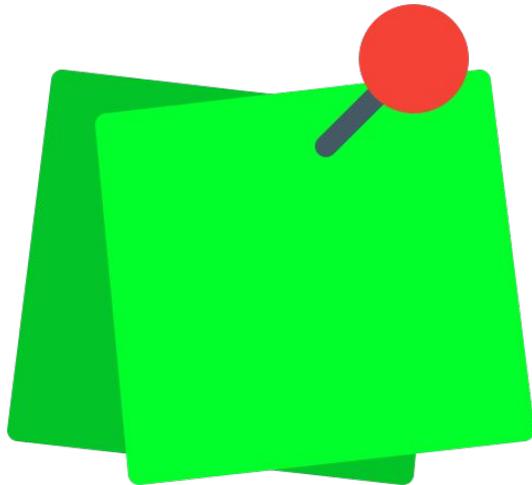


Market Research for Binder

(done reproducibly)

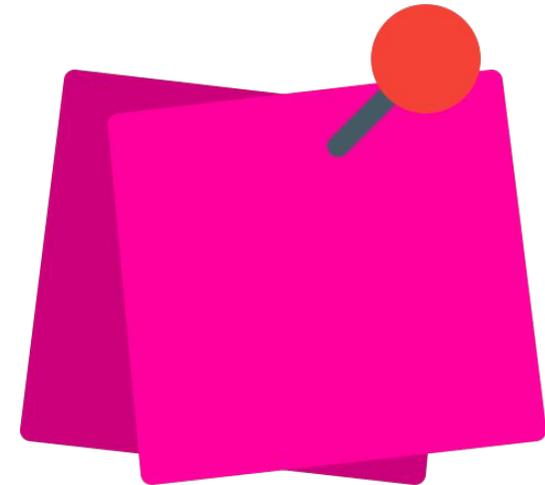
Indicate how much you agree with the following statements
using traffic lights.

Today I am using: Linux (any flavour)



Agree

Disagree

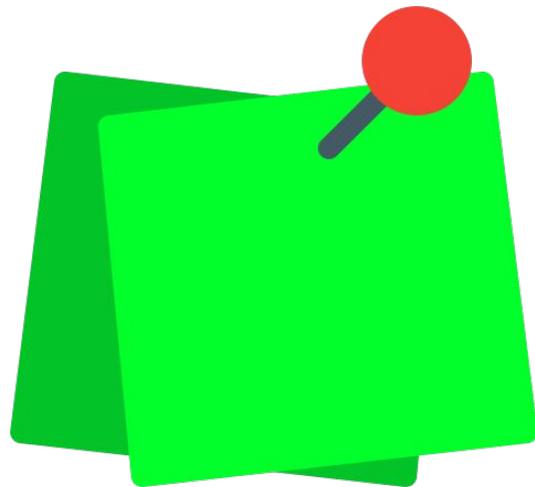


Market Research for Binder

(done reproducibly)

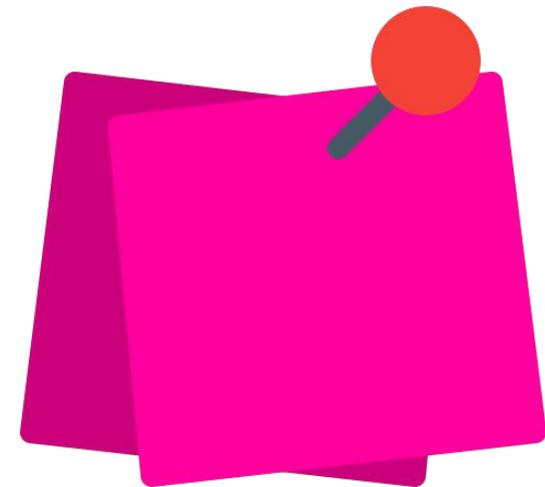
Indicate how much you agree with the following statements
using traffic lights.

I have heard of Binder before.



Agree

Disagree

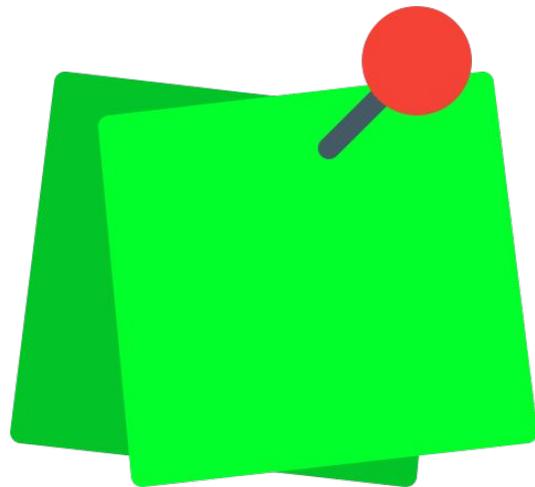


Market Research for Binder

(done reproducibly)

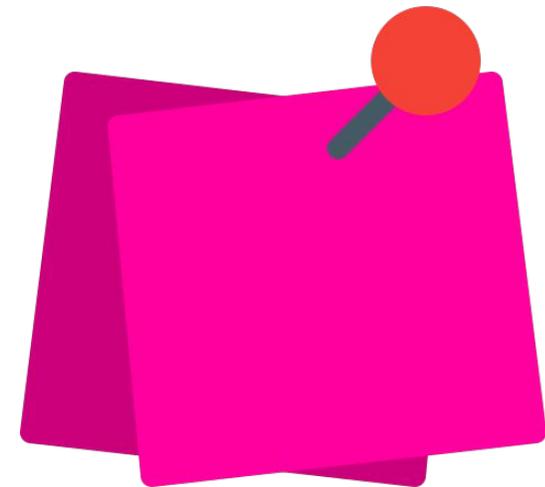
Indicate how much you agree with the following statements using traffic lights.

I have used binder to look at someone else's research.



Agree

Disagree

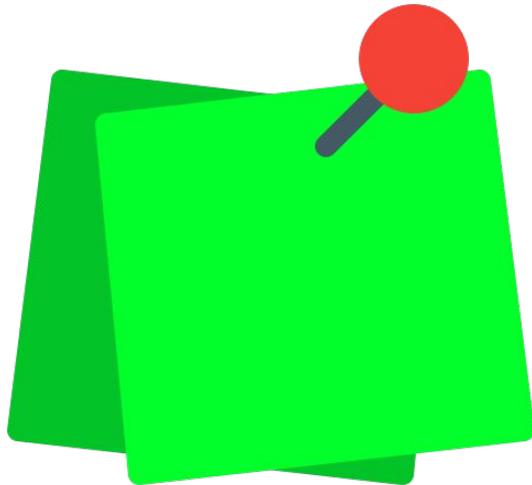


Market Research for Binder

(done reproducibly)

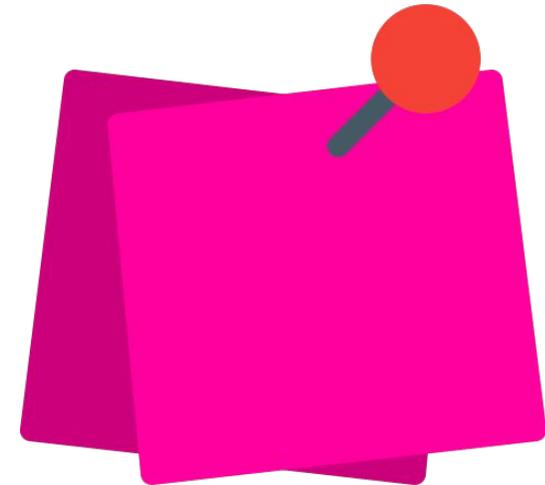
Indicate how much you agree with the following statements using traffic lights.

I have used Binder to share my own work.



Agree

Disagree

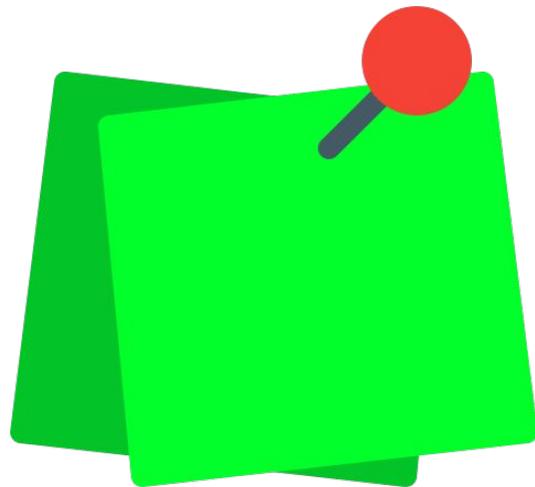


Market Research for Binder

(done reproducibly)

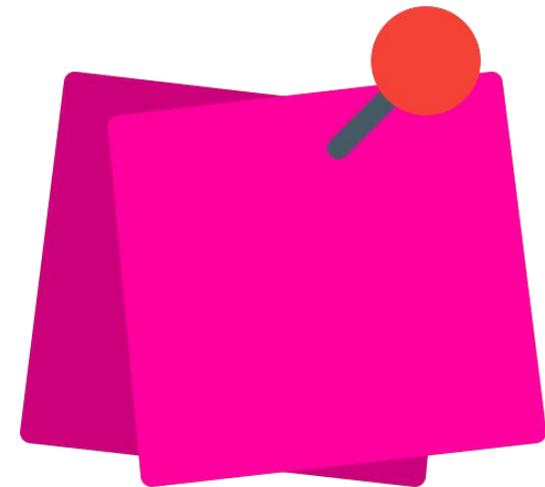
Indicate how much you agree with the following statements using traffic lights.

I have a good understanding of how Binder works.



Agree

Disagree

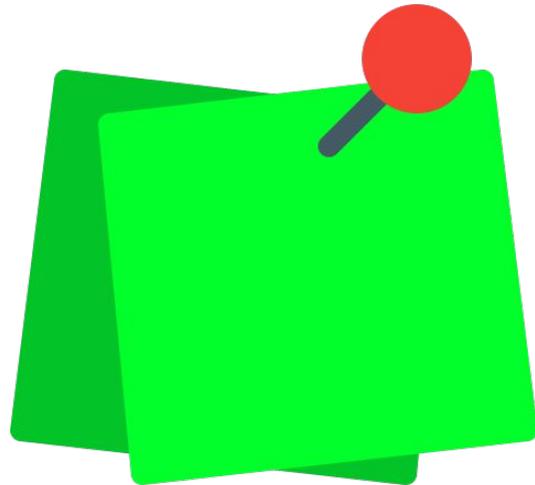


Market Research for Binder

(done reproducibly)

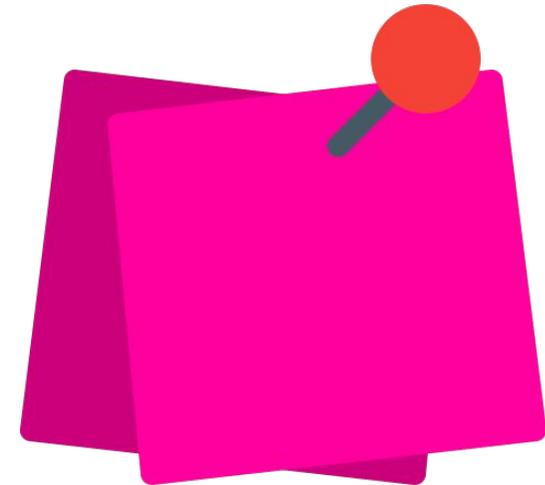
Indicate how much you agree with the following statements using traffic lights.

I have heard (something like) these words before: "Oh, it worked on my computer?"



Agree

Disagree

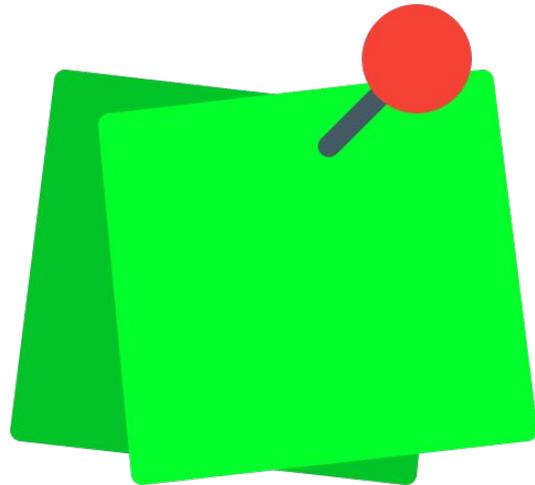


Market Research for Binder

(done reproducibly)

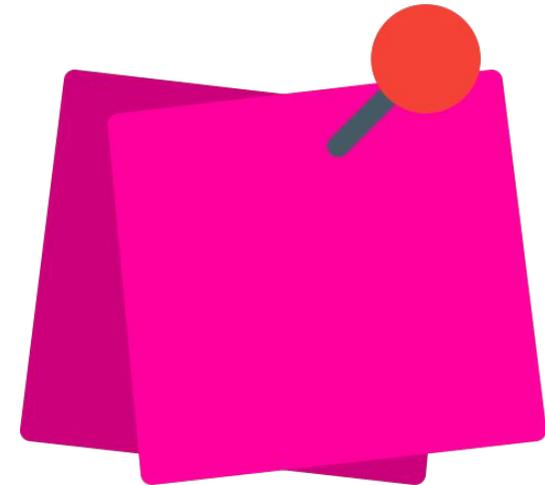
Indicate how much you agree with the following statements using traffic lights.

I have heard (something like) these words before: "Oh, it was working yesterday?"



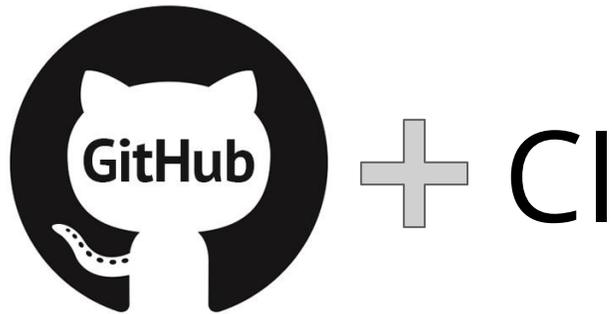
Agree

Disagree



Why Binder?

Oh, it was
working
yesterday?



Oh, it worked
on my
computer?



Binder

A totally fictional user story.

Binder

A *totally* fictional user story.

The screenshot shows a GitHub repository page for 'alexmorley / binder-demo'. At the top, there is a navigation bar with the GitHub logo, a search bar, and links for 'Pull requests', 'Issues', 'Marketplace', and 'Explore'. On the right side of the navigation bar, there are notification and user profile icons. Below the navigation bar, the repository name 'alexmorley / binder-demo' is displayed, along with 'Unwatch' (1), 'Star' (0), and 'Fork' (0) buttons. A secondary navigation bar includes 'Code', 'Issues' (0), 'Pull requests' (0), 'Projects' (0), 'Wiki', 'Insights', and 'Settings'. The main content area shows a message: 'No description, website, or topics provided.' with an 'Edit' button and a 'Manage topics' link. Below this, a summary bar indicates '1 commit', '1 branch', '0 releases', and '1 contributor'. A row of buttons includes 'Branch: master', 'New pull request', 'Create new file', 'Upload files', 'Find File', and 'Clone or download'. The commit history shows one commit by 'alexmorley' with the message 'Initial commit' and the hash 'f52964f' from 'a minute ago'. Below the commit list, the 'README.md' file is shown with its content: 'binder-demo'.

alexmorley / binder-demo

Unwatch 1 Star 0 Fork 0

Code Issues 0 Pull requests 0 Projects 0 Wiki Insights Settings

No description, website, or topics provided. Edit

Manage topics

1 commit 1 branch 0 releases 1 contributor

Branch: master New pull request Create new file Upload files Find File Clone or download

alexmorley Initial commit Latest commit f52964f a minute ago

README.md Initial commit a minute ago

README.md

binder-demo

© 2019 GitHub, Inc. Terms Privacy Security Status Help Contact GitHub Pricing API Training Blog About

Binder

Add a file:

The screenshot shows the GitHub interface for the repository 'alexmorley / binder-demo'. At the top, there is a navigation bar with the GitHub logo, a search bar, and links for 'Pull requests', 'Issues', 'Marketplace', and 'Explore'. Below this, the repository name is displayed along with statistics: 'Unwatch 1', 'Star 0', and 'Fork 0'. A secondary navigation bar includes 'Code', 'Issues 0', 'Pull requests 0', 'Projects 0', 'Wiki', 'Insights', and 'Settings'. The main content area shows a message: 'No description, website, or topics provided.' with an 'Edit' button. Below this, there are statistics for '1 commit', '1 branch', '0 releases', and '1 contributor'. A row of buttons includes 'Branch: master', 'New pull request', 'Create new file' (circled in red), 'Upload files', 'Find File', and 'Clone or download'. Below the buttons, a commit history entry is shown for 'alexmorley Initial commit' with the latest commit hash 'f52964f' and a timestamp of 'a minute ago'. A file named 'README.md' is listed as an 'Initial commit' from 'a minute ago'. The file content area shows the text 'binder-demo'.

alexmorley / binder-demo

Unwatch 1 Star 0 Fork 0

Code Issues 0 Pull requests 0 Projects 0 Wiki Insights Settings

No description, website, or topics provided. Edit

Manage topics

1 commit 1 branch 0 releases 1 contributor

Branch: master New pull request Create new file Upload files Find File Clone or download

alexmorley Initial commit Latest commit f52964f a minute ago

README.md Initial commit a minute ago

README.md

binder-demo

© 2019 GitHub, Inc. Terms Privacy Security Status Help Contact GitHub Pricing API Training Blog About

Binder

Add some data to the file

The screenshot shows a GitHub repository page for 'alexmorley / binder-demo'. The repository has 1 watch, 0 stars, and 0 forks. The file 'data.csv' is being edited in the 'Edit new file' mode. The file content is as follows:

```
1 question_tag,number_agree,number_disagree
2 heard_of_binder,12,0
3 used_binder,4,8
4 share_binder,3,9
5 understand_binder,6,6
```

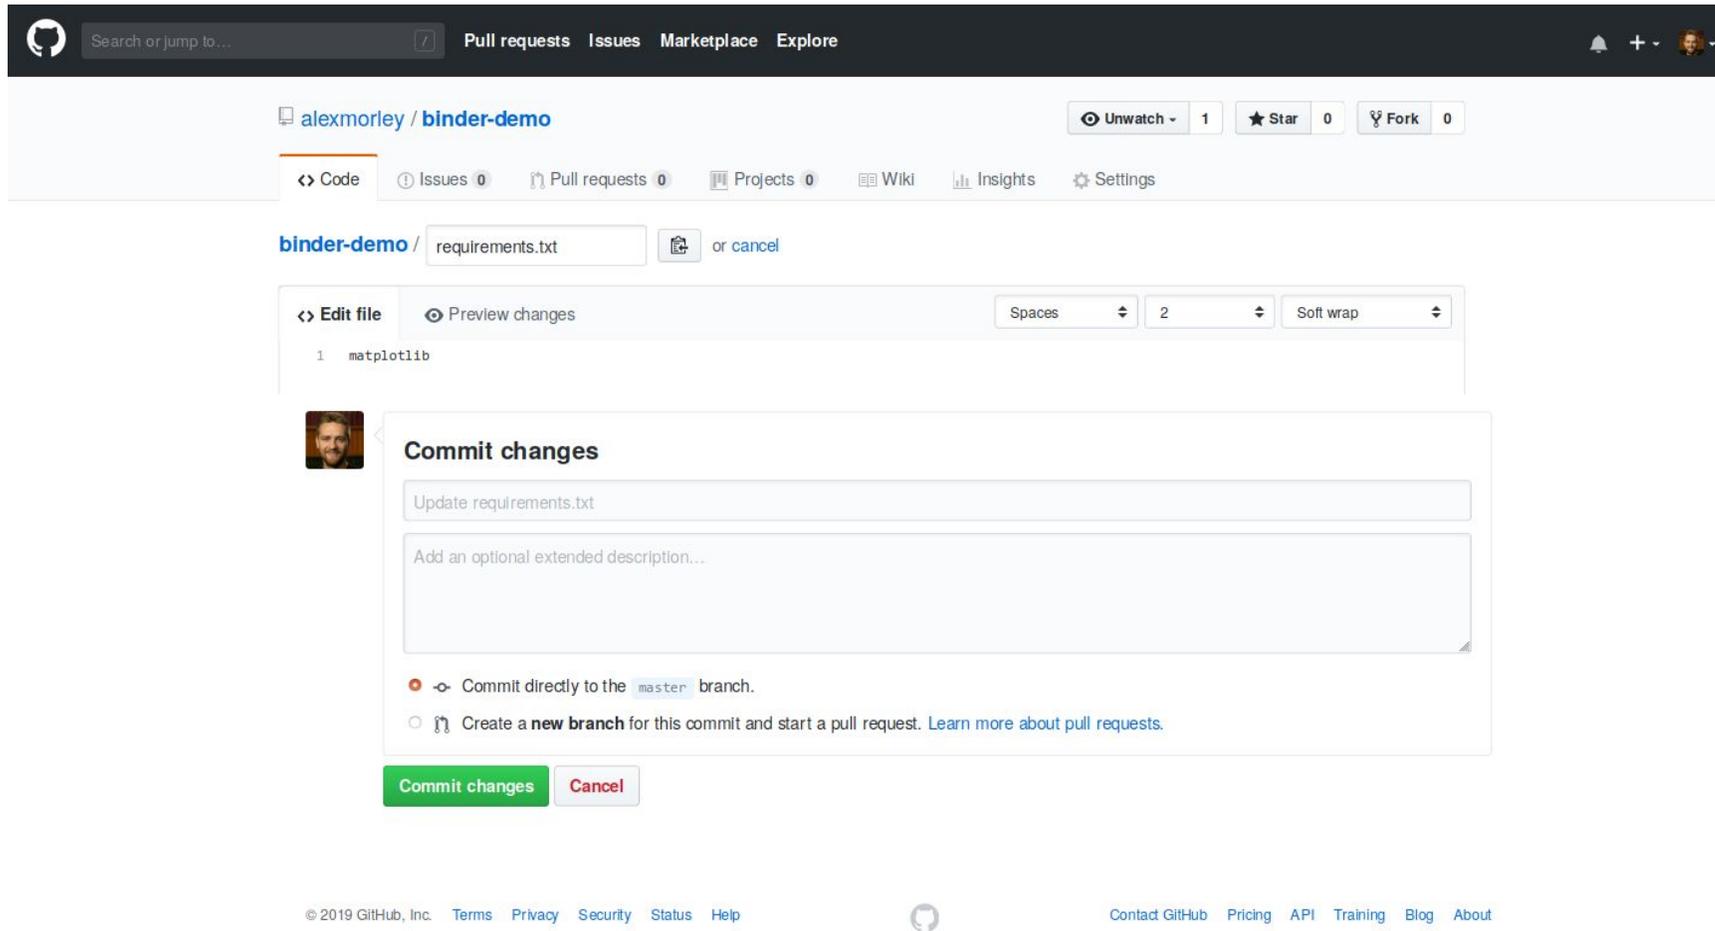
A 'Commit new file' dialog is open, showing the following options:

- Commit directly to the `master` branch.
- Create a **new branch** for this commit and start a pull request. [Learn more about pull requests.](#)

Buttons for 'Commit new file' and 'Cancel' are visible at the bottom of the dialog.

Binder

Specify a few requirements



The screenshot shows a GitHub repository page for 'alexmorley/binder-demo'. The repository has 1 Unwatch, 0 Stars, and 0 Forks. The file 'requirements.txt' is selected for editing. The commit dialog is open, showing the commit message 'Update requirements.txt' and an optional extended description field. The commit options are:

- Commit directly to the `master` branch.
- Create a **new branch** for this commit and start a pull request. [Learn more about pull requests.](#)

Buttons for 'Commit changes' and 'Cancel' are visible at the bottom of the dialog.

© 2019 GitHub, Inc. [Terms](#) [Privacy](#) [Security](#) [Status](#) [Help](#) [Contact GitHub](#) [Pricing](#) [API](#) [Training](#) [Blog](#) [About](#)

Binder

Copy-paste the github repo name into mybinder.org. Click Launch.



Turn a Git repo into a collection of interactive notebooks

Have a repository full of Jupyter notebooks? With Binder, open those notebooks in an executable environment, making your code immediately reproducible by anyone, anywhere.

A screenshot of the Binder web interface. The main heading is "Build and launch a repository". Below it is a form with several fields: "GitHub repository name or URL" containing "https://github.com/alexmorley/binder-demo", "Git branch, tag, or commit" with a dropdown menu, and "Path to a notebook file (optional)" with a text input and a "File" button. A prominent orange "launch" button is to the right of the path field. Below the form, there is a section "Copy the URL below and share your Binder with others:" with a text box containing "https://mybinder.org/v2/gh/alexmorley/binder-demo/master". At the bottom, there is a section "Copy the text below, then paste into your README to show a binder badge:" with a small badge icon and a right-pointing arrow. Two red circles are drawn on the image: one around the "GitHub repository name or URL" field and another around the "launch" button.

How it works

1 Enter your repository information

Binder

Open a notebook

Files Running Clusters

Select items to perform actions on them.

0 /

| | | |
|--------------------------|------------------|--|
| <input type="checkbox"/> | data.csv | |
| <input type="checkbox"/> | README.md | |
| <input type="checkbox"/> | requirements.txt | |

Upload New

Notebook:
Python 3

Other:
Text File 3 B
Folder 4 B



Binder

The possibilities are endless

 jupyter Untitled Last Checkpoint: 5 minutes ago (unsaved changes)



File Edit View Insert Cell Kernel Widgets Help

Trusted

Python 3 

         Run    Code  

The Possibilities are endless

```
In [ ]: import anything_from_requirements.txt
```

Binder

Submitting Changes Back to the Repo

How did this all happen?!

BinderHub. Not quite magic.

Build and launch a repository

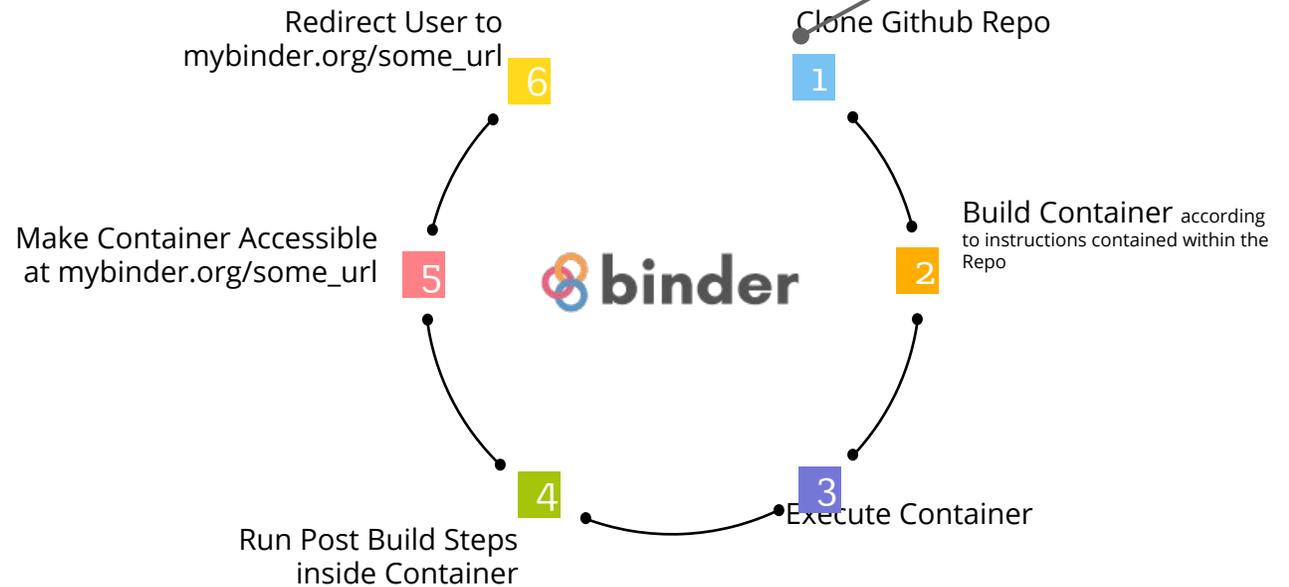
GitHub repository name or URL

 GitHub ▾

Git branch, tag, or commit

 ⓘ

Path to a notebook file (optional)

 File ▾

BinderHub

Build and launch a repository

GitHub repository name or URL

 GitHub ▾

Git branch, tag, or commit

 📄

Path to a notebook file (optional)

 File ▾

Clone Github Repo

1



BinderHub

Build and launch a repository

GitHub repository name or URL

 GitHub ▾

Git branch, tag, or commit

 📄

Path to a notebook file (optional)

 File ▾

launch

Clone Github Repo

1



2

Build Container
according to instructions
contained within the Repo



BinderHub

Build and launch a repository

GitHub repository name or URL

 GitHub ▾

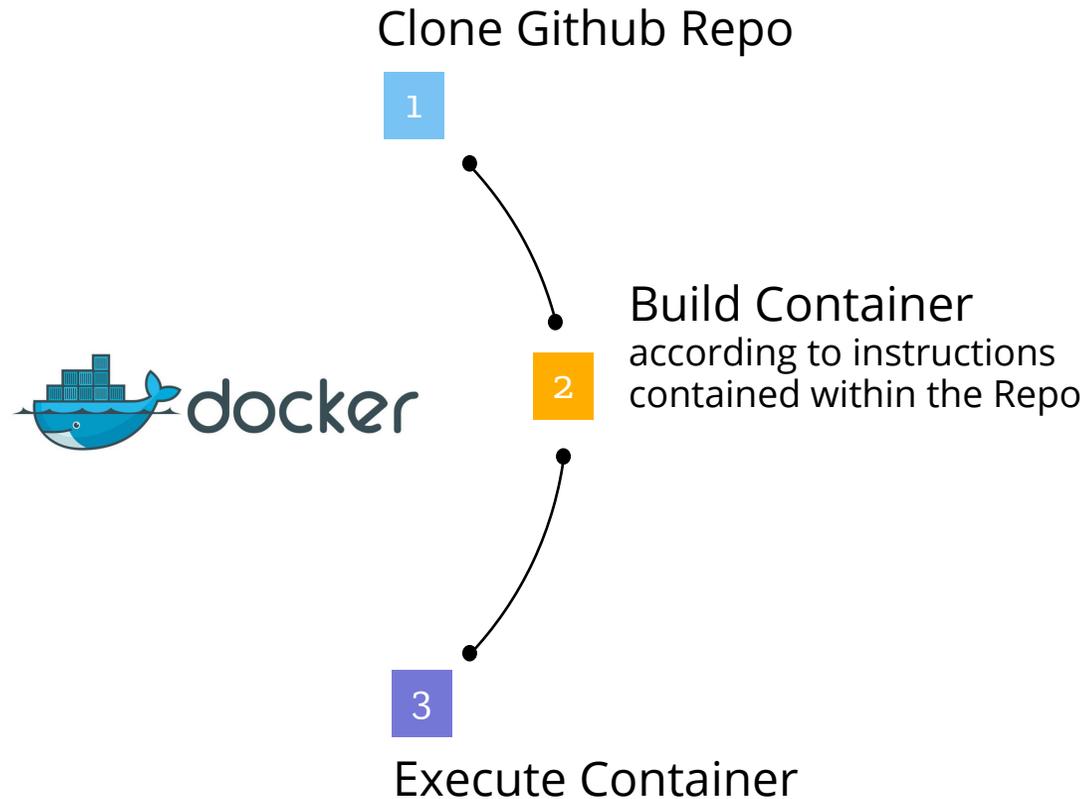
Git branch, tag, or commit

 📄

Path to a notebook file (optional)

 File ▾

launch



BinderHub

Build and launch a repository

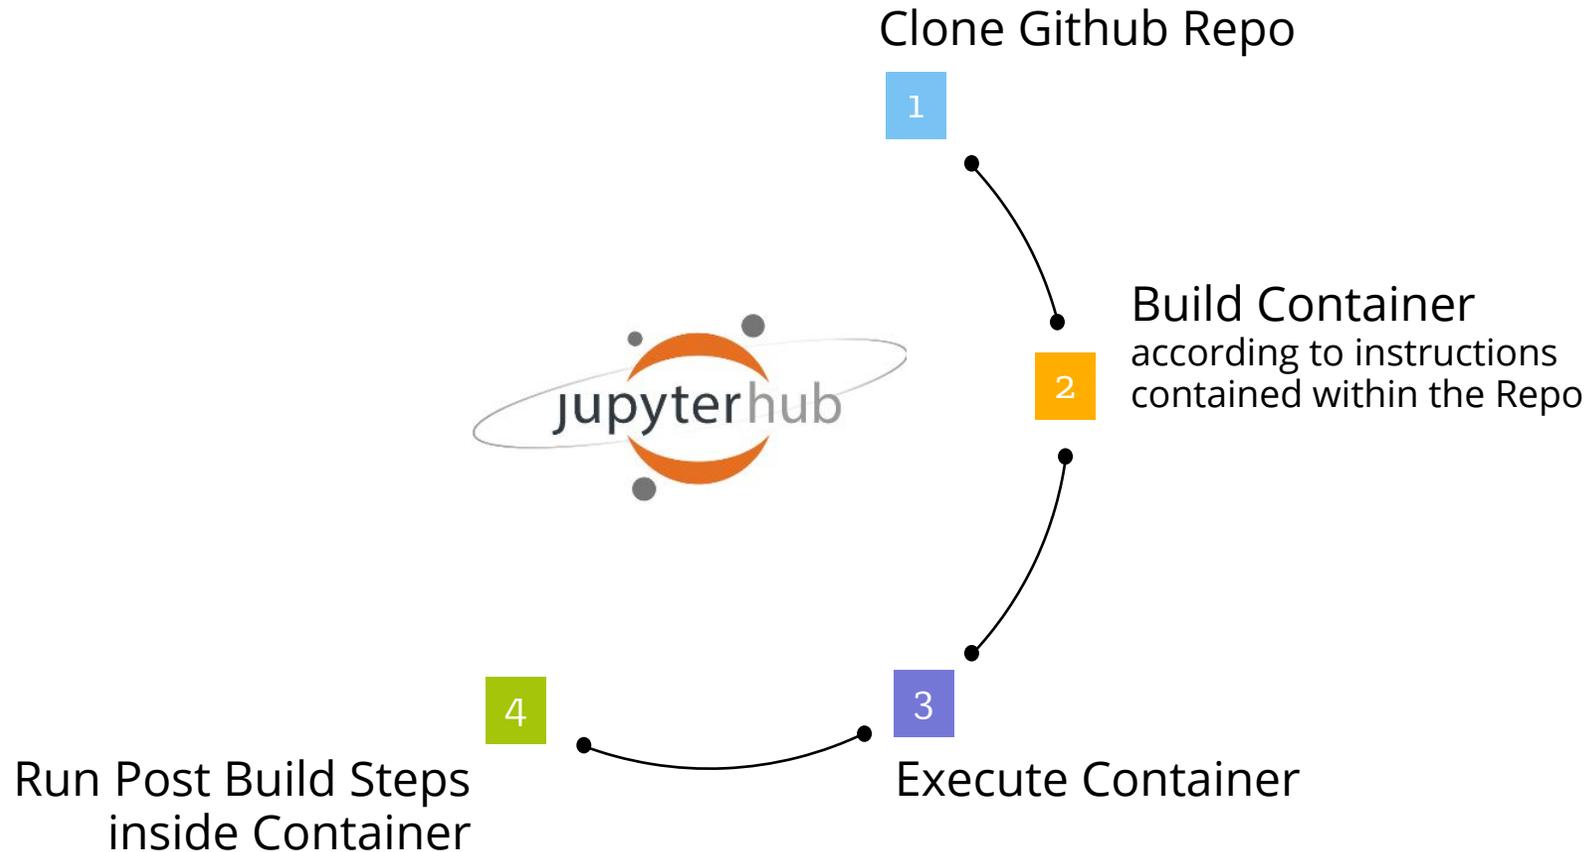
GitHub repository name or URL

 GitHub ▾

Git branch, tag, or commit

 📄

Path to a notebook file (optional)

 File ▾

BinderHub

Build and launch a repository

GitHub repository name or URL

 GitHub ▾

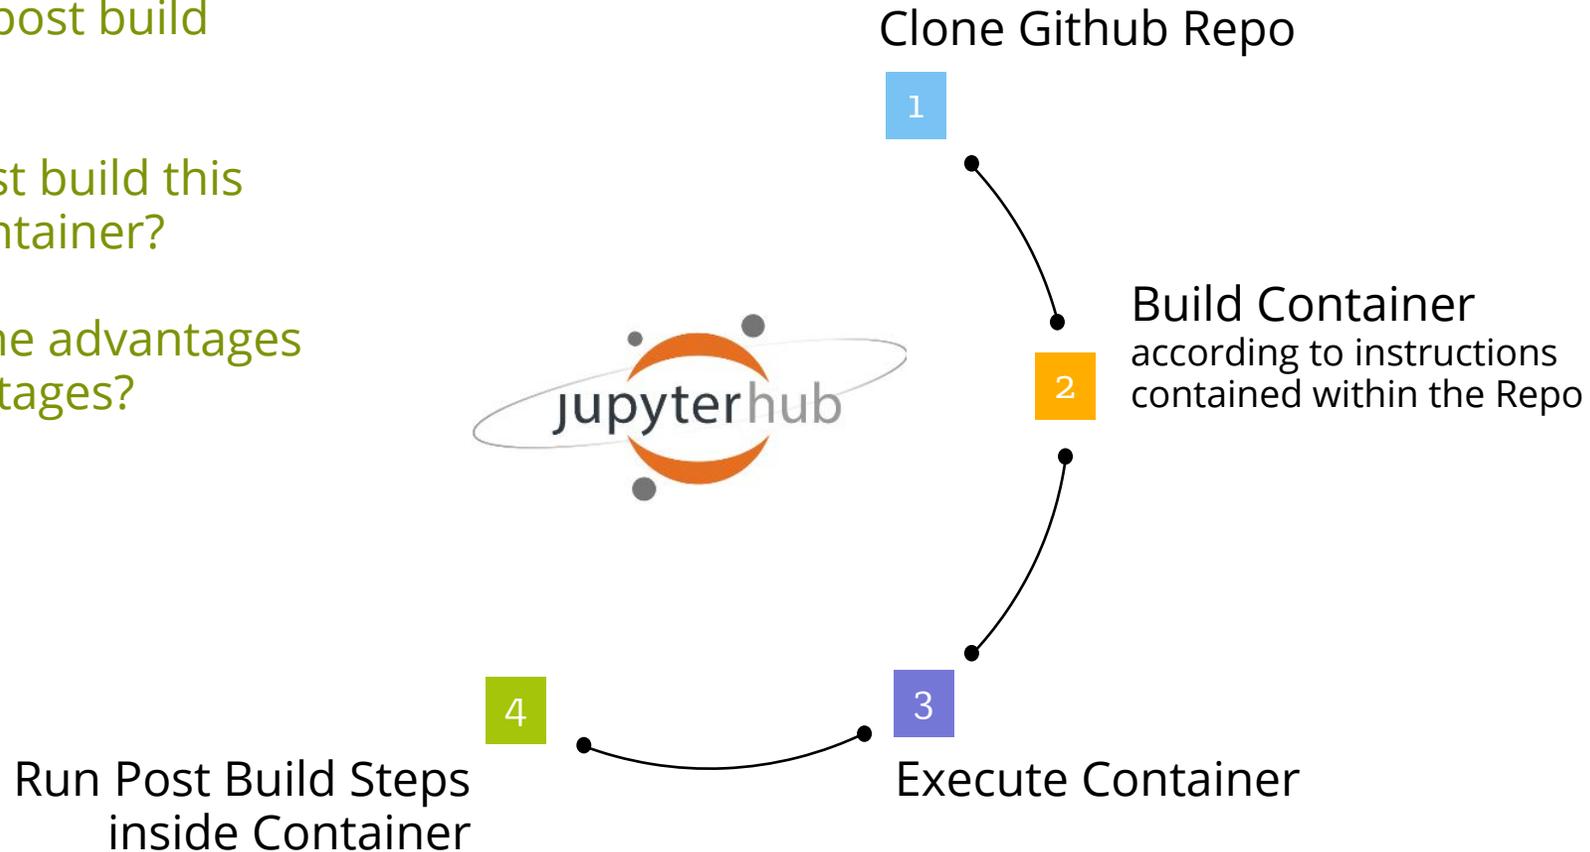
Git branch, tag, or commit

 📄

Path to a notebook file (optional)

 File ▾

- Why run a post build step?
- Why not just build this into the container?
- What are the advantages & disadvantages?



BinderHub

Build and launch a repository

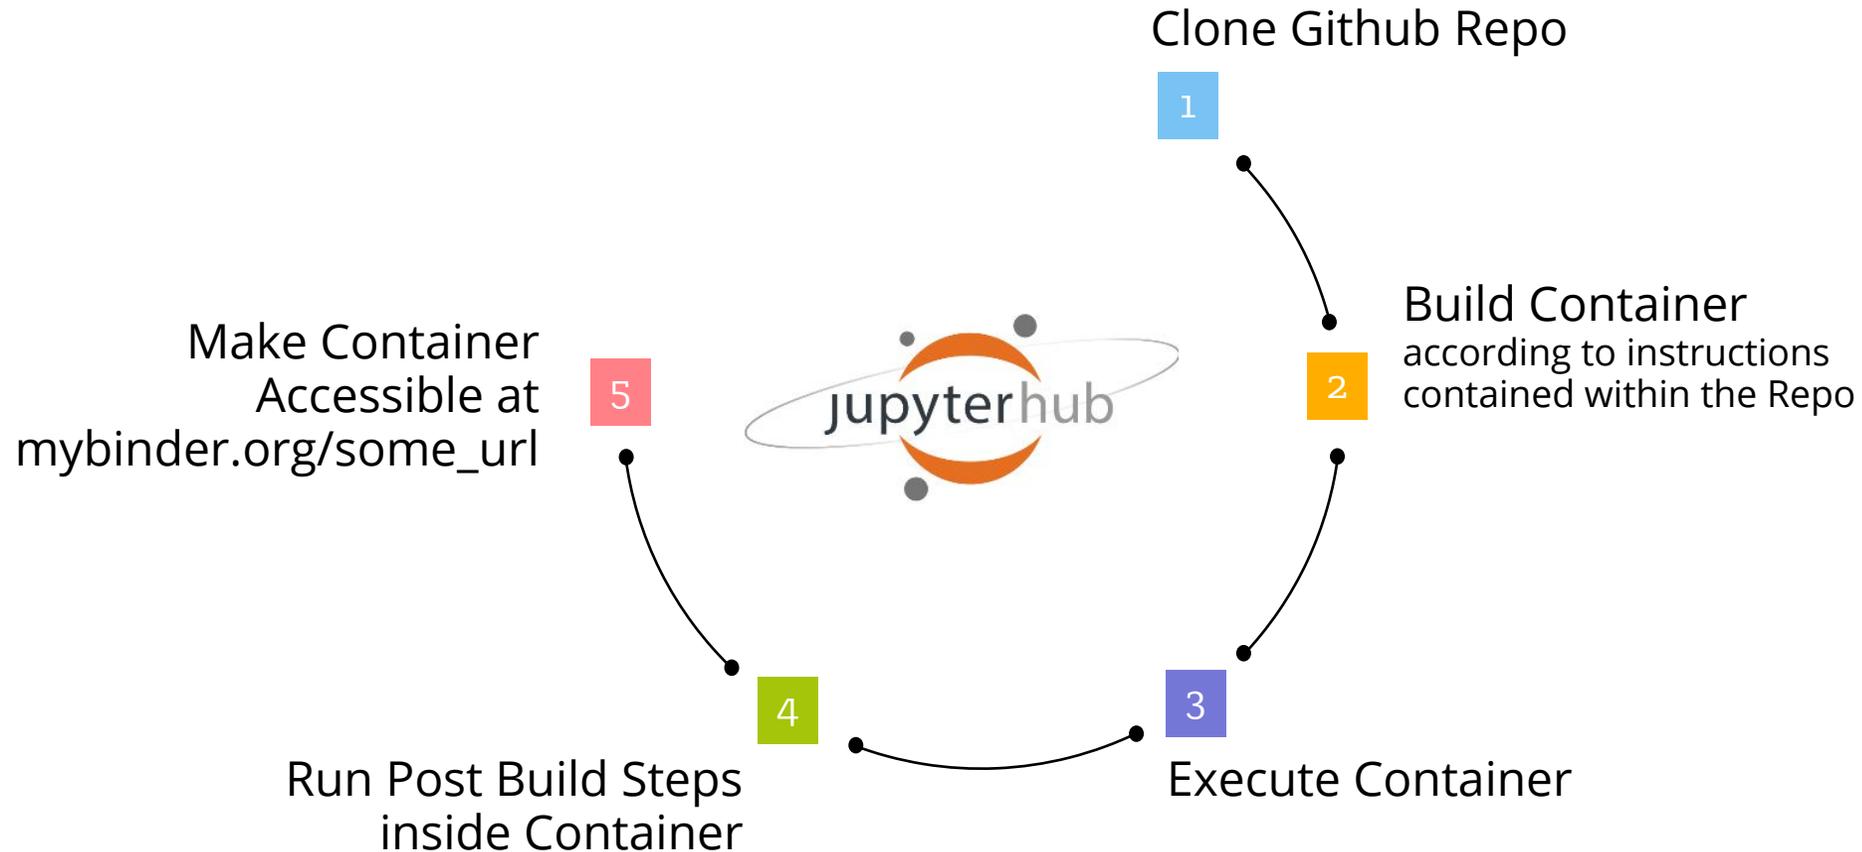
GitHub repository name or URL

 GitHub ▾

Git branch, tag, or commit

 📄

Path to a notebook file (optional)

 File ▾

BinderHub

Build and launch a repository

GitHub repository name or URL

 GitHub ▾

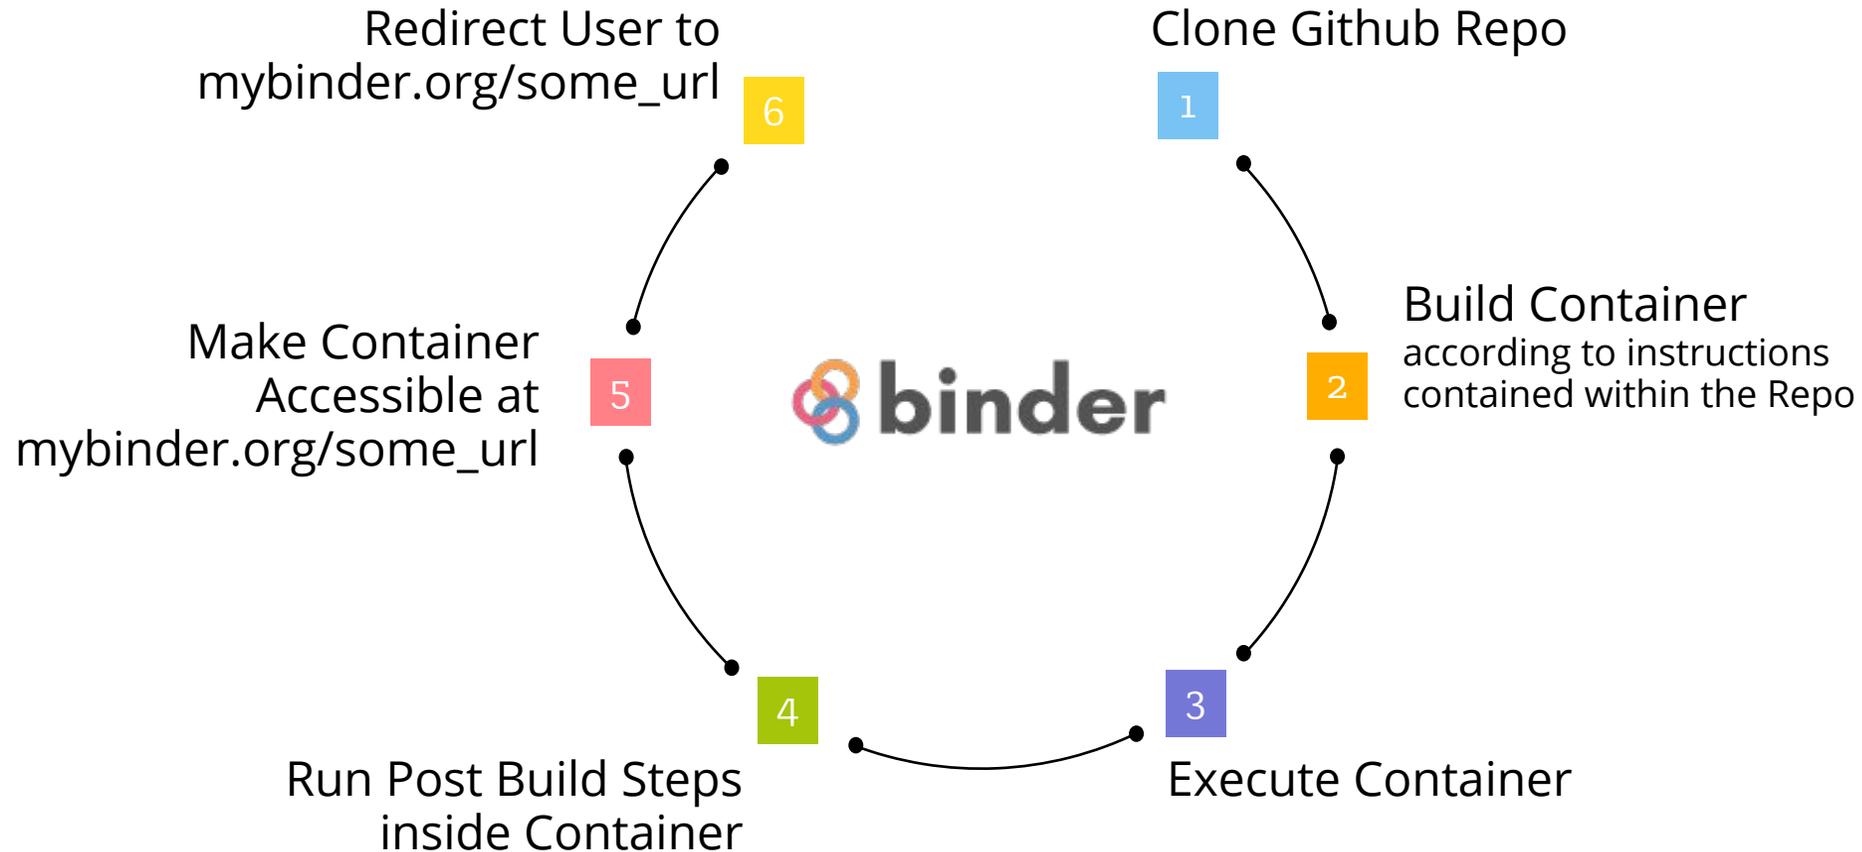
Git branch, tag, or commit



Path to a notebook file (optional)

File ▾

launch



How did this all happen?!

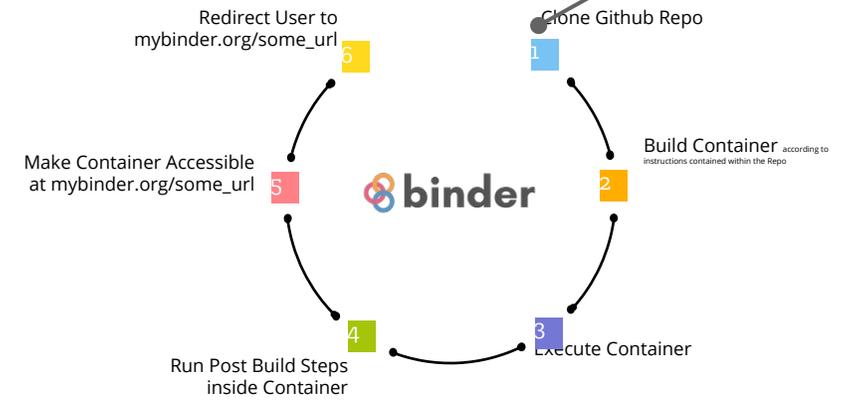
No magic. Just glueing together lots of other software.

Build and launch a repository

GitHub repository name or URL
 GitHub

Git branch, tag, or commit

Path to a notebook file (optional)
 File



jupyterhub

GitHub

docker

+

A few more...

repo2docker

So can I just
``pip install
binderhub?``

Not quite. At least, it won't do
what you think.

**Why can we not just run this all from one
server?**

- 1.
- 2.
- 3.

So can I just
`pip install
binderhub?`

Not quite. At least, it won't do
what you think.

**Why can we not just run this all from one
server?**

1. Resource Intensive
2. Resource Control
3. Security

Solution: Kubernetes



The solution that strikes fear and passion into sysadmins around the world.

Why can we not just run this all from one server?

1. Resource Intensive - **Cluster Management**
2. Resource Control - **Container Management**
3. Security - **Container Isolation**

Problem? Kubernetes



The solution that strikes fear and passion into sysadmins around the world.

Err that sounds great? Why does it strike fear into anyone?

Can be a pain to set up correctly.

Solution: Zero-to-Binderhub

Fear no more!

What we'd like from you:

- Make a note of which bits were tricky.
- Think about what doing this “for real” at home/your workplace? What extra would you need / want?
- Add these thoughts/notes on stickies and put them in the Parking Lot - we're going to try fix them this afternoon!

Thanks for
listening :)