



Marge-bot: better Git'ing with python

Europython 2018

Alexander Schmolck
Mika Bostrom

README.md

build passing

Marge-bot

Marge-bot is a merge-bot for GitLab that, beside other goodies, implements [the Not Rocket Science Rule Of Software Engineering](#):

automatically maintain a repository of code that always passes all the tests.

— Graydon Hoare, main author of Rust

This simple rule of thumb is still nowadays surprisingly difficult to implement with the state-of-the-art tools, and more a way that scales with team size (also see our [blog post](#)).



What's wrong with this picture?

README.md

build passing

Marge-bot

Marge-bot is a merge-bot for GitLab that, beside other goodies, implements [the Not Rocket Science Rule Of Software Engineering](#):

automatically maintain a repository of code that always passes all the tests.

— Graydon Hoare, main author of Rust

This simple rule of thumb is still nowadays surprisingly difficult to implement with the state-of-the-art tools, and more a way that scales with team size (also see our [blog post](#)).






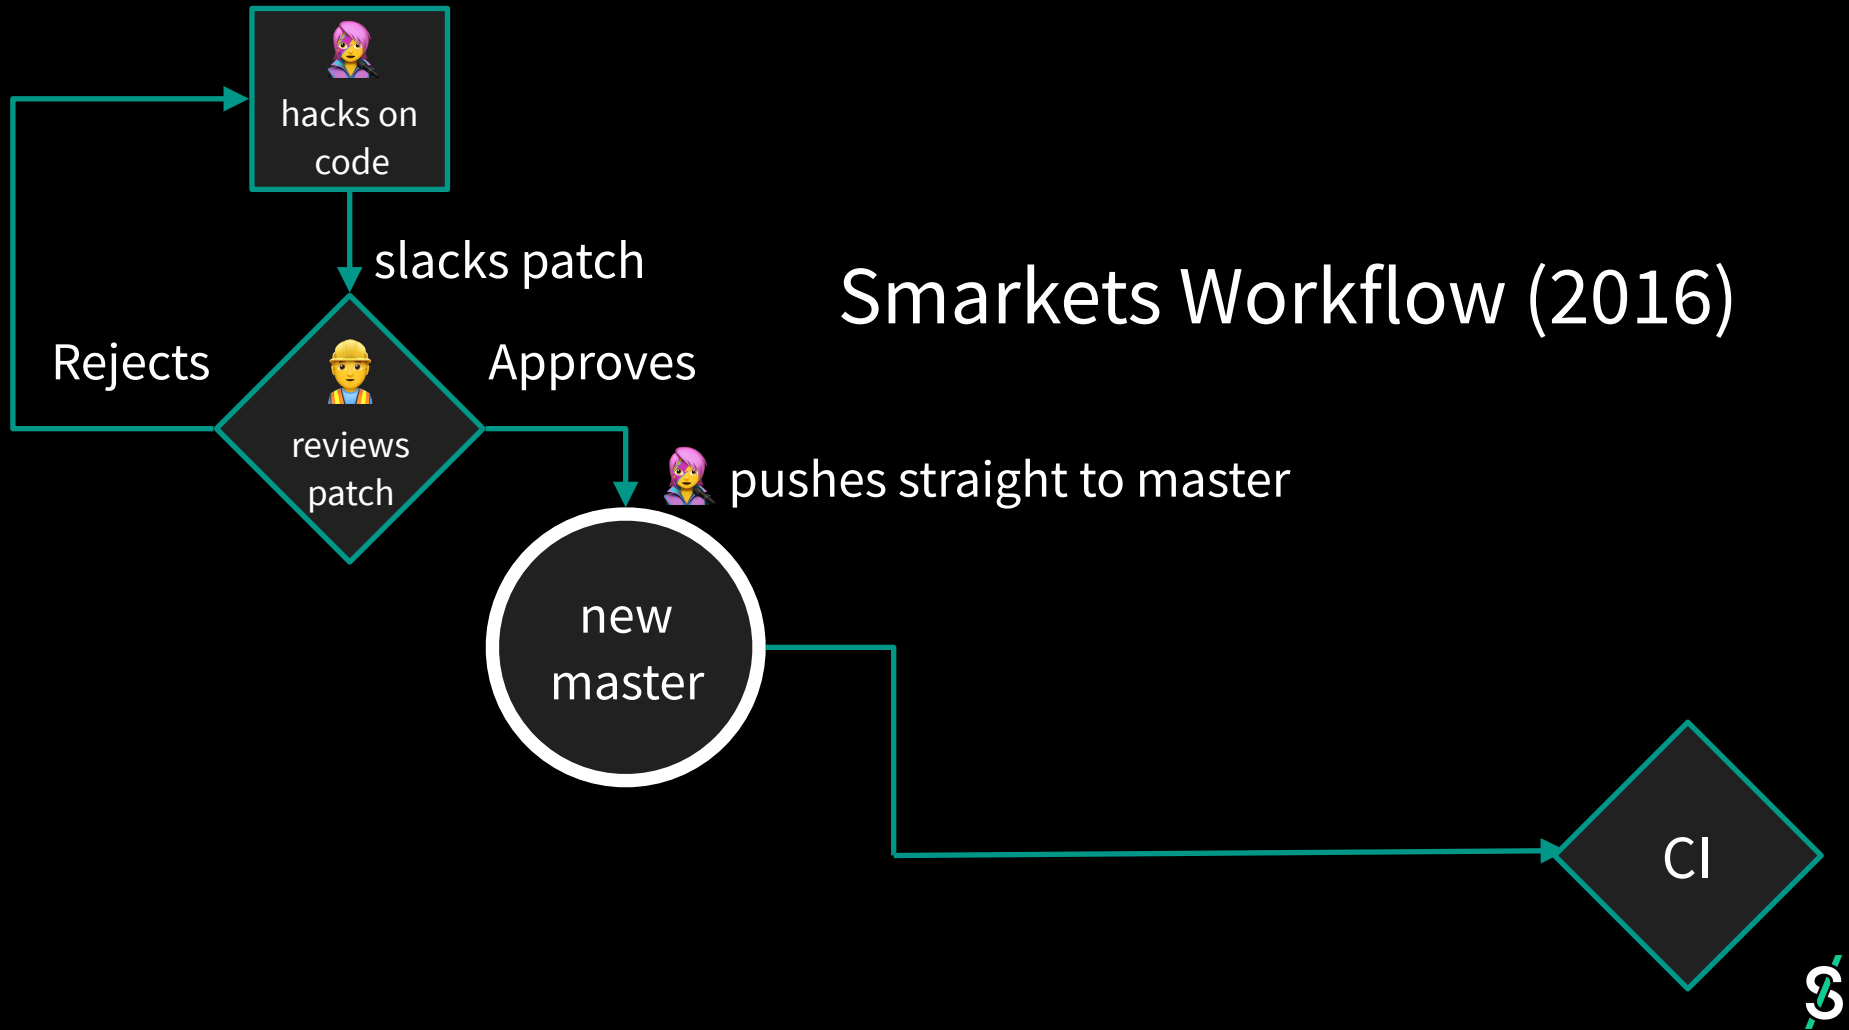
Outline

- The typical CI setup or going from green PRs to red master
- The fix, conceptually
- The fix, in practice: achieving familiarity, usability and scalability
- More nice stuff:
 - auditing, usable git bisect & git revert
- **Audience takeaways:**
 - Flexible & free out-of-the box solution for Gitlab (~3 mins setup)
 - Solutions, DIY or otherwise if you're not on Gitlab (e.g. Github)
 - Some useful Git workflow and Python patterns

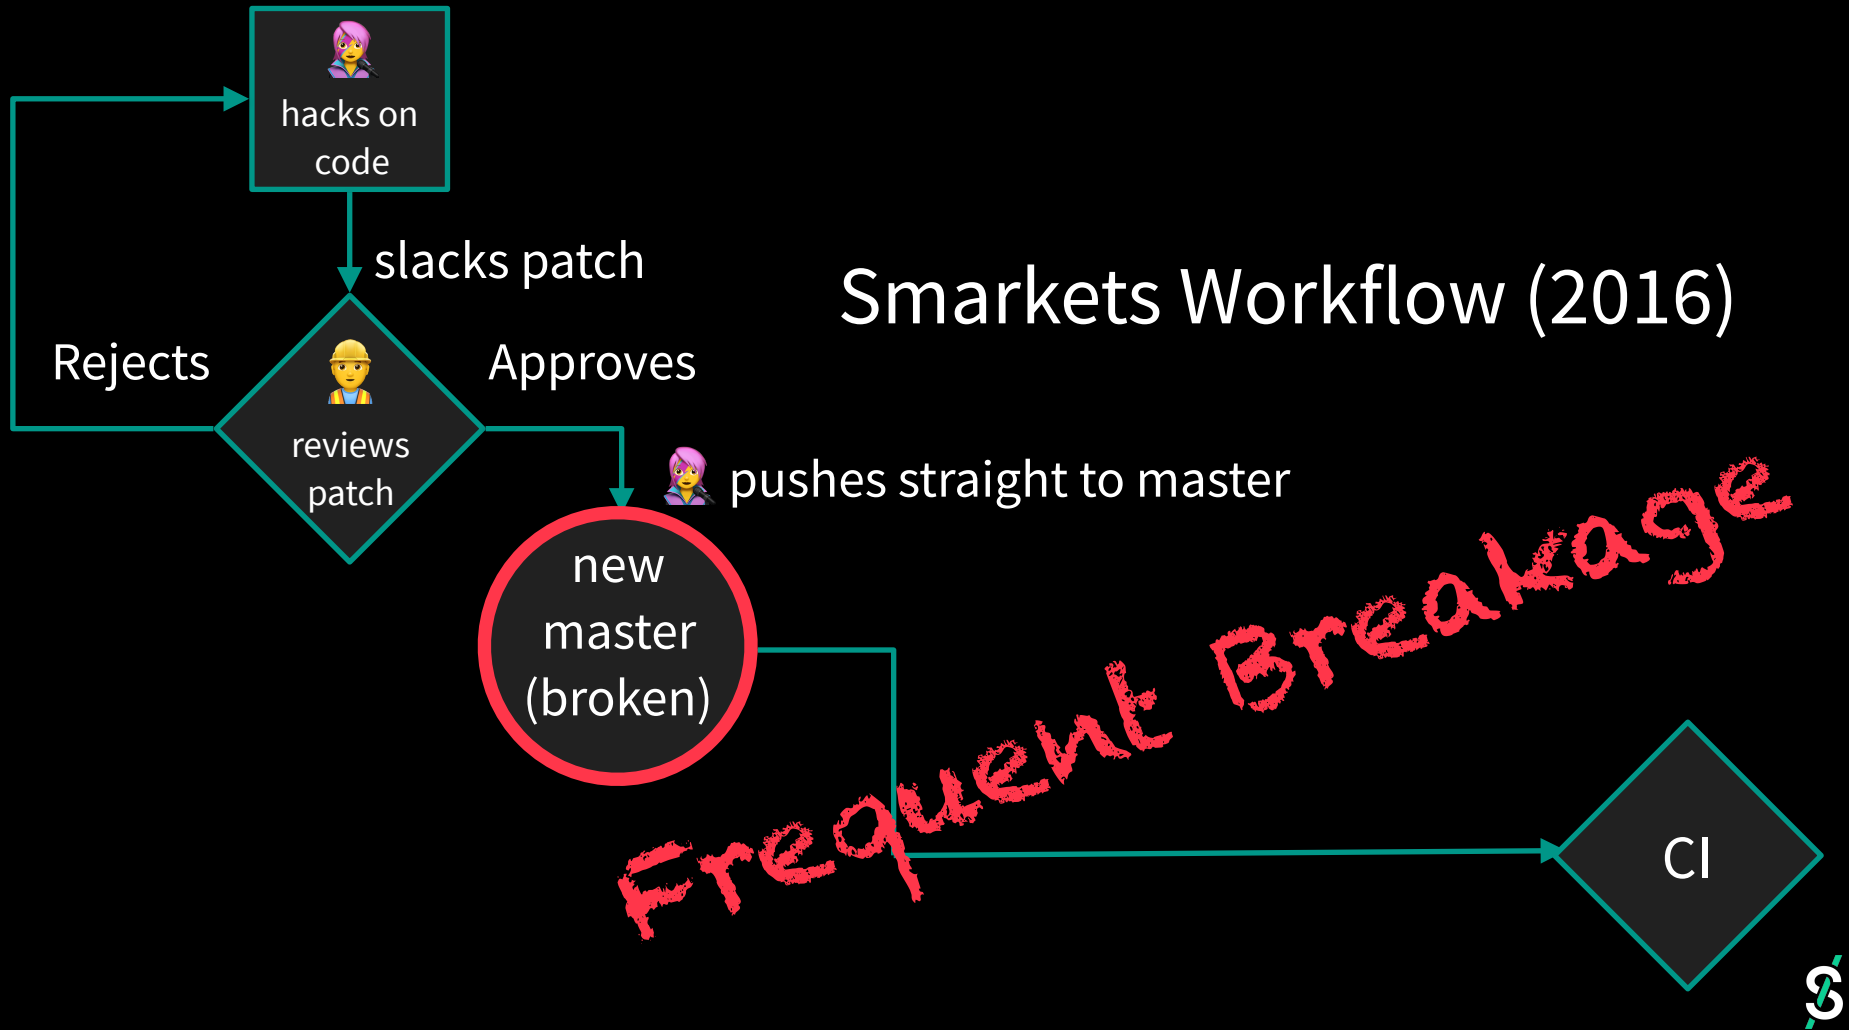


Why **broken** master is bad


-  subsequent PRs are built on sand
-  bad ships to prod more likely
-  retracing your steps becomes hard (e.g. bisect)

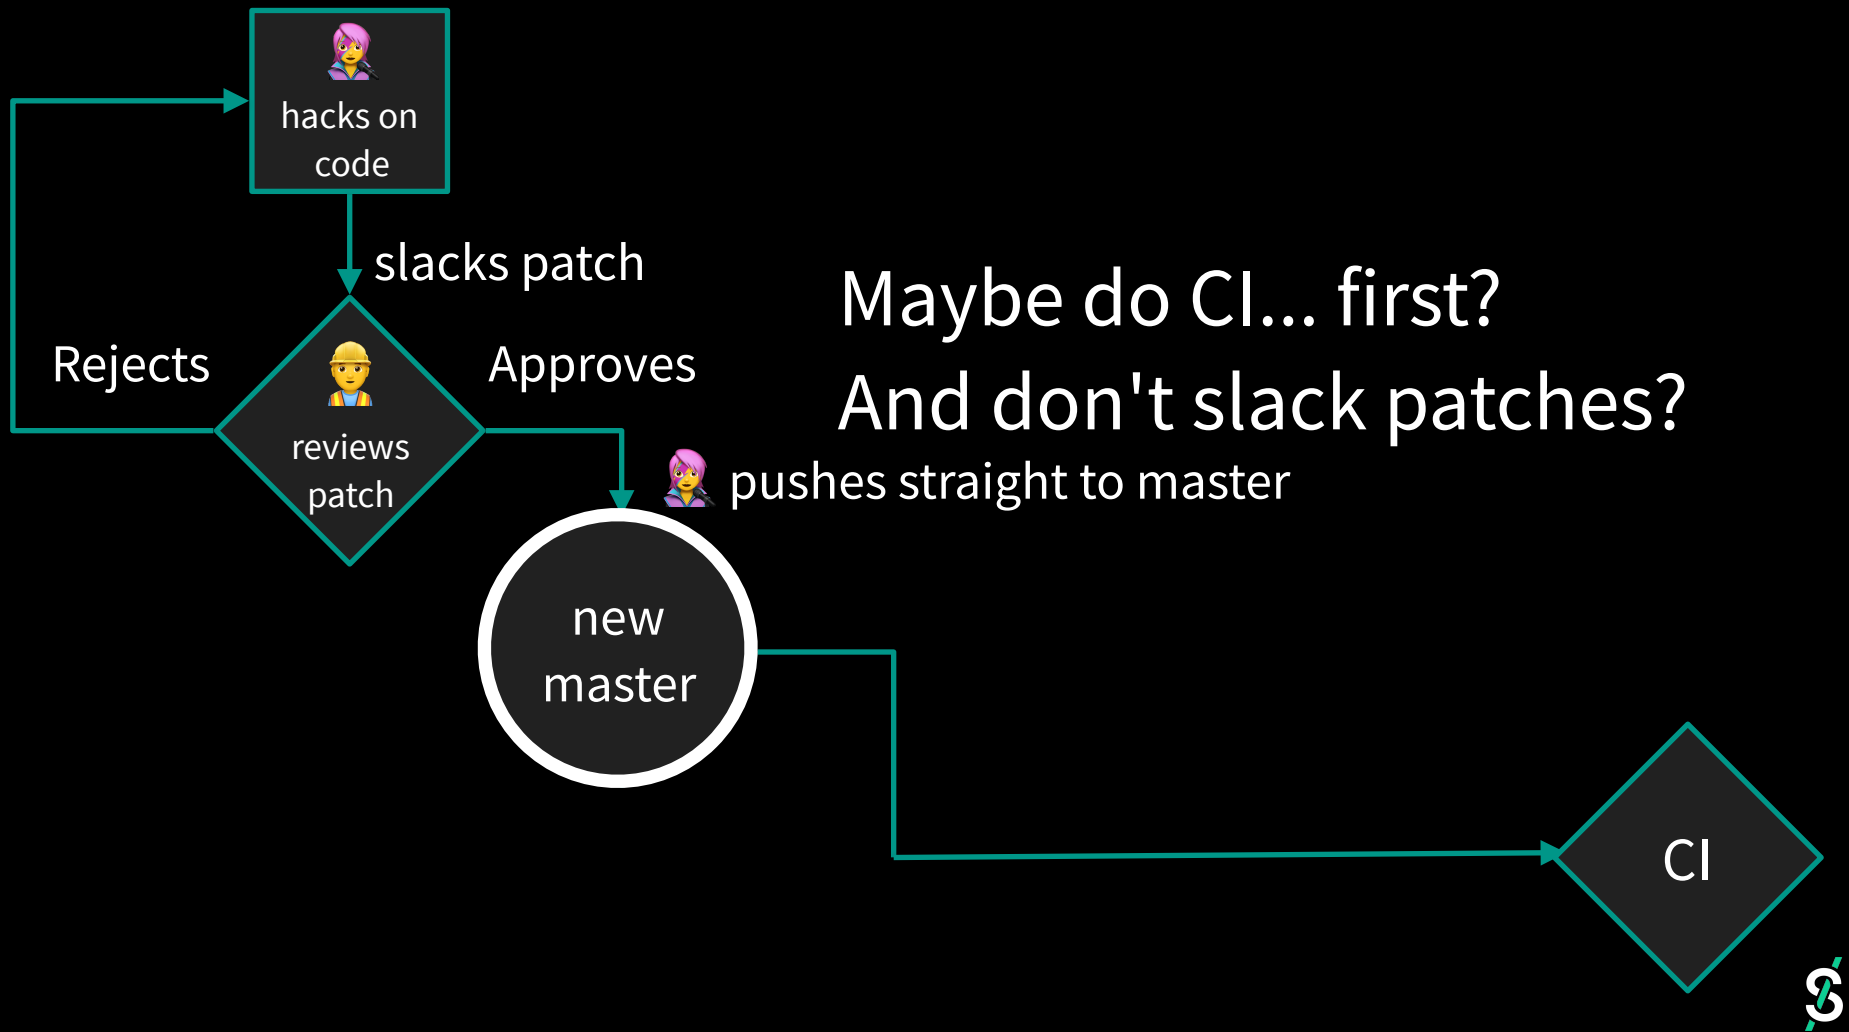


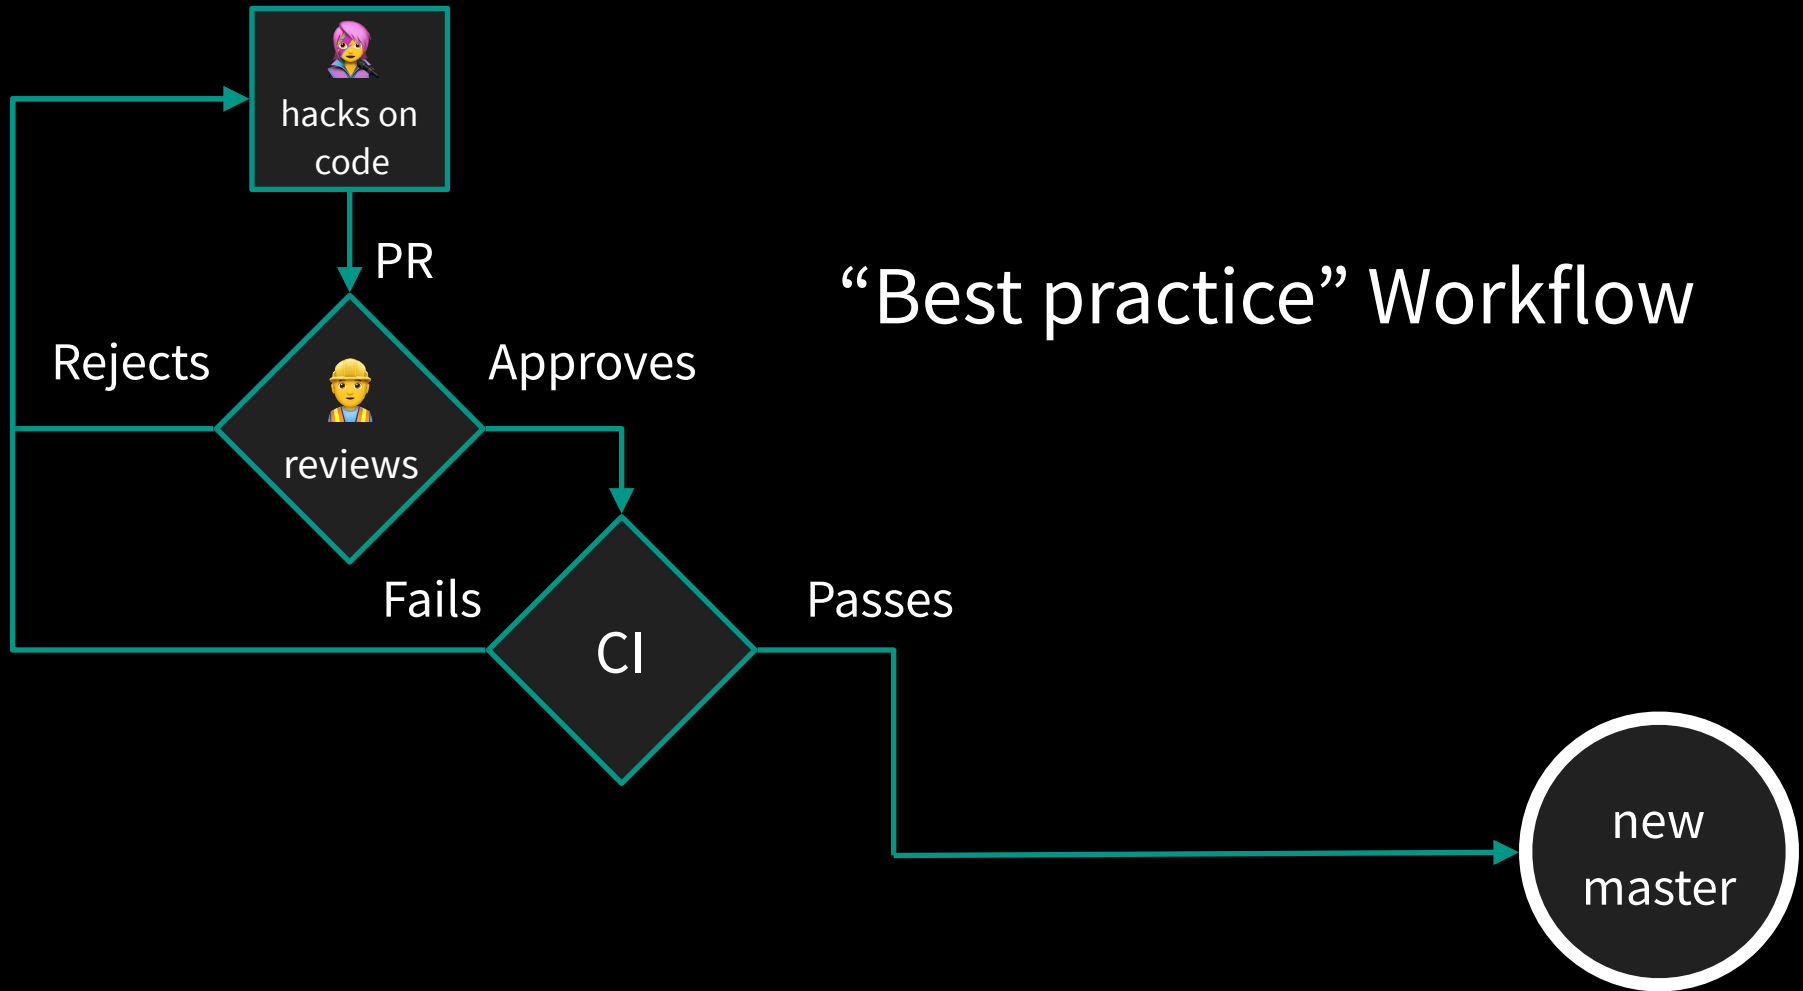
Smarkets Workflow (2016)

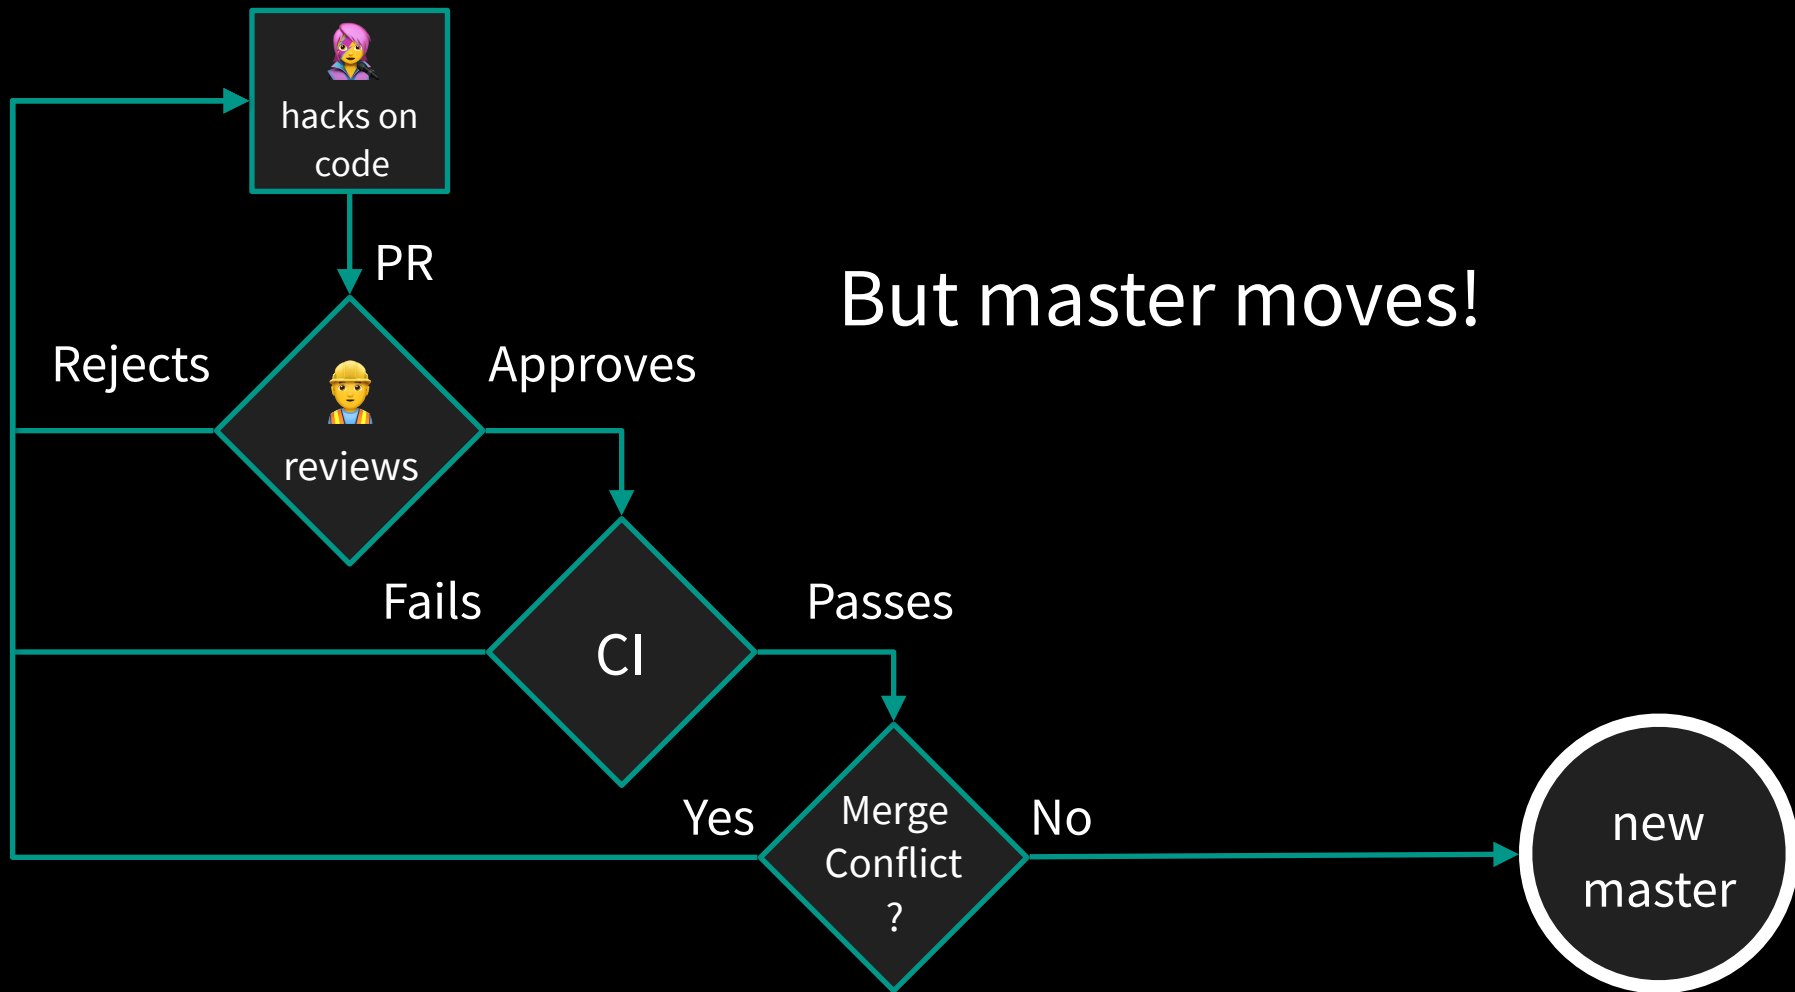


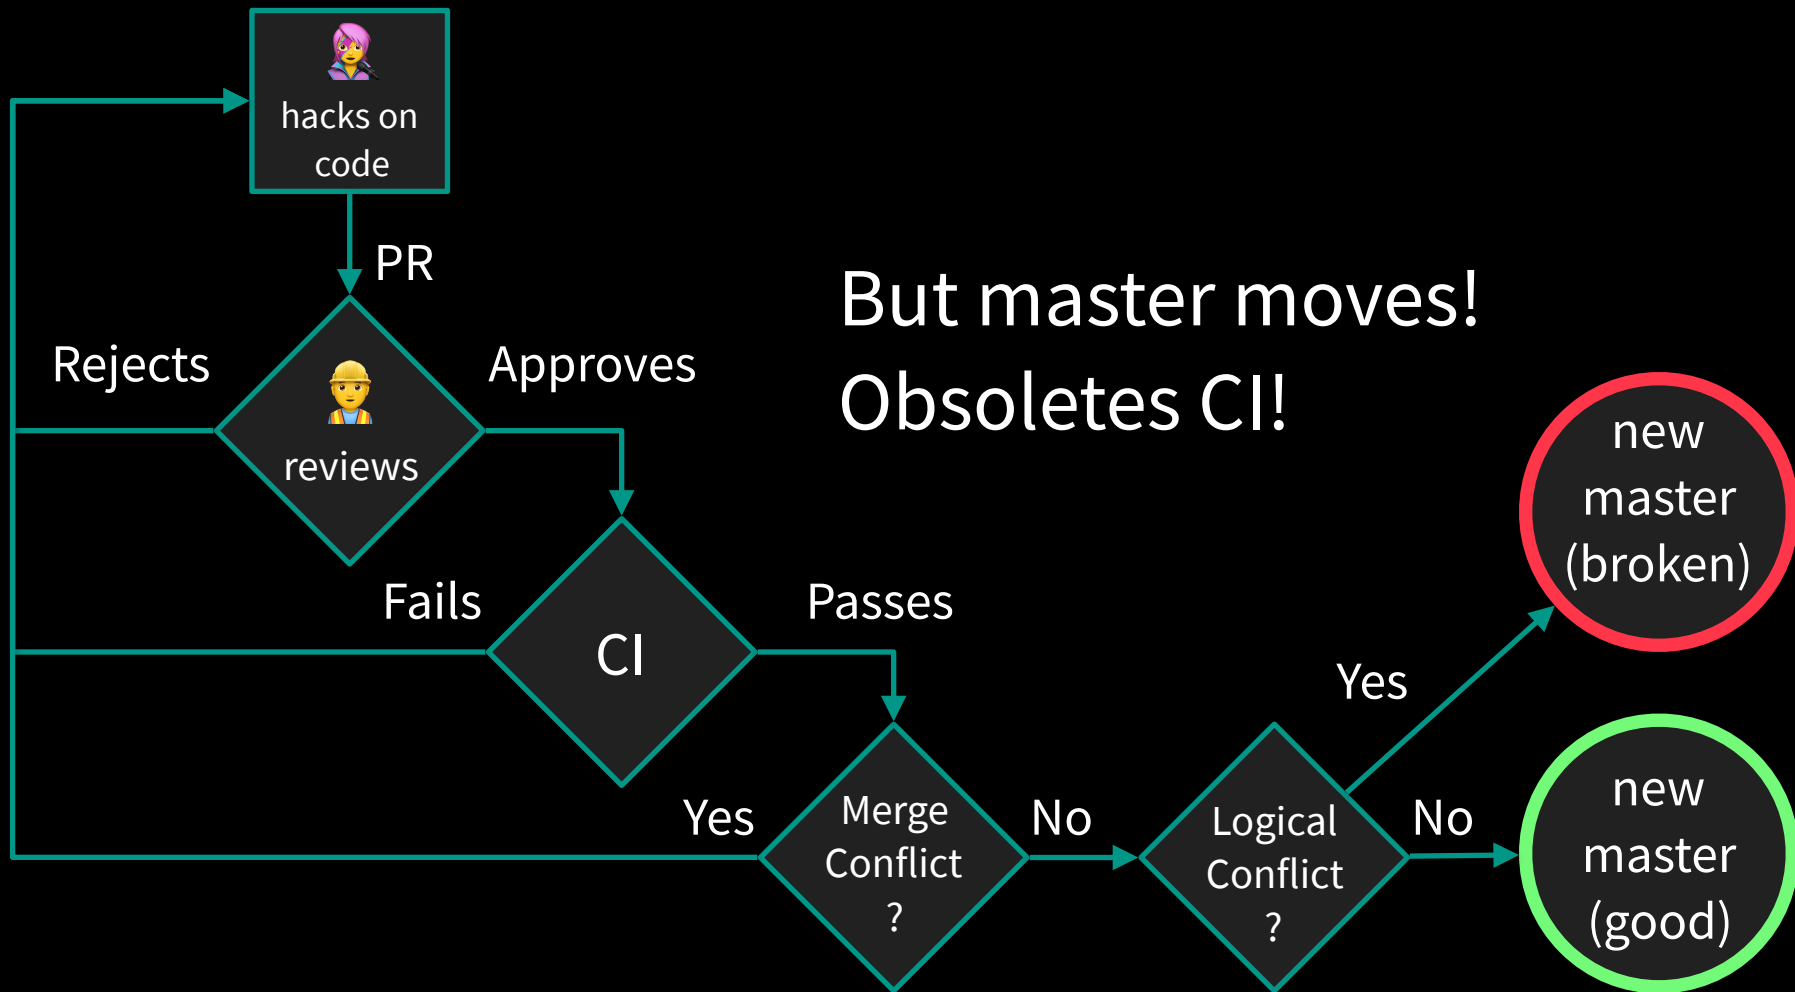
Master can be **broken** because

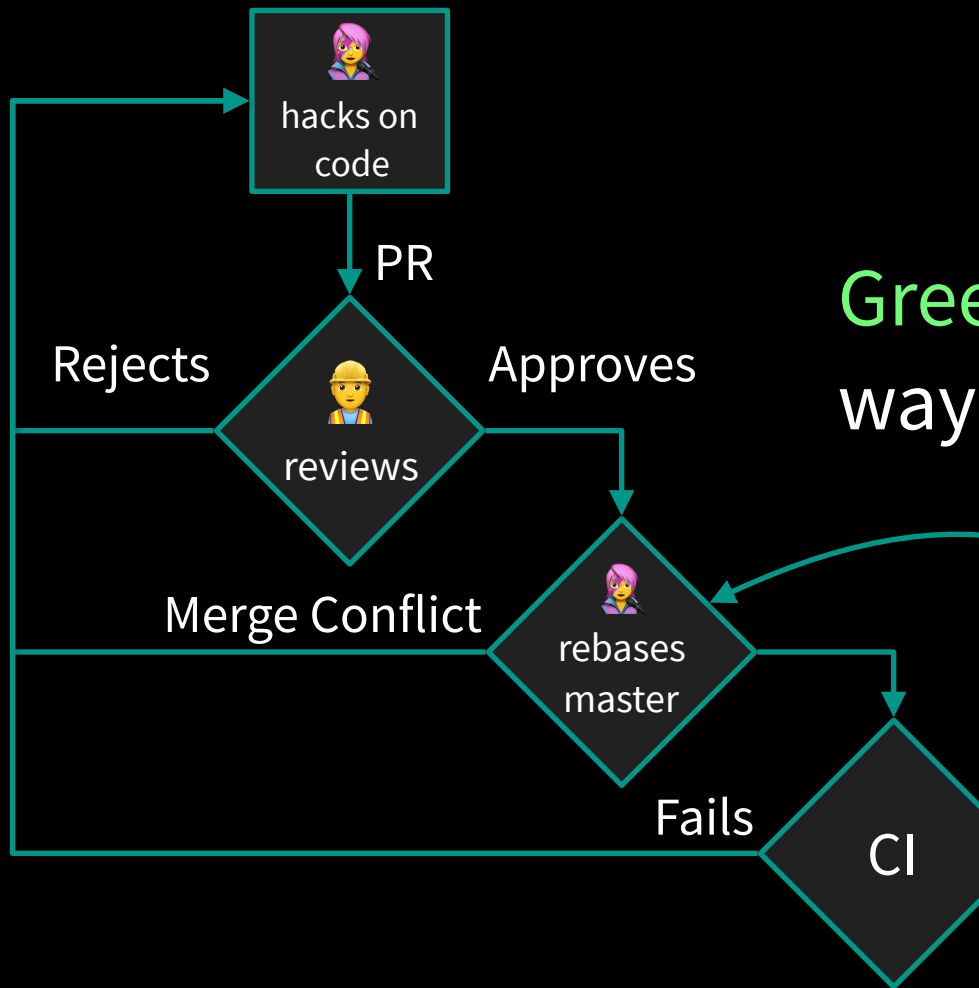
- bad workflows
(spoiler:  marge-bot can fix that for you!)
- “flaky builds”
(non-determinism is harder to fix)











Green master the hard way

Merge Requests

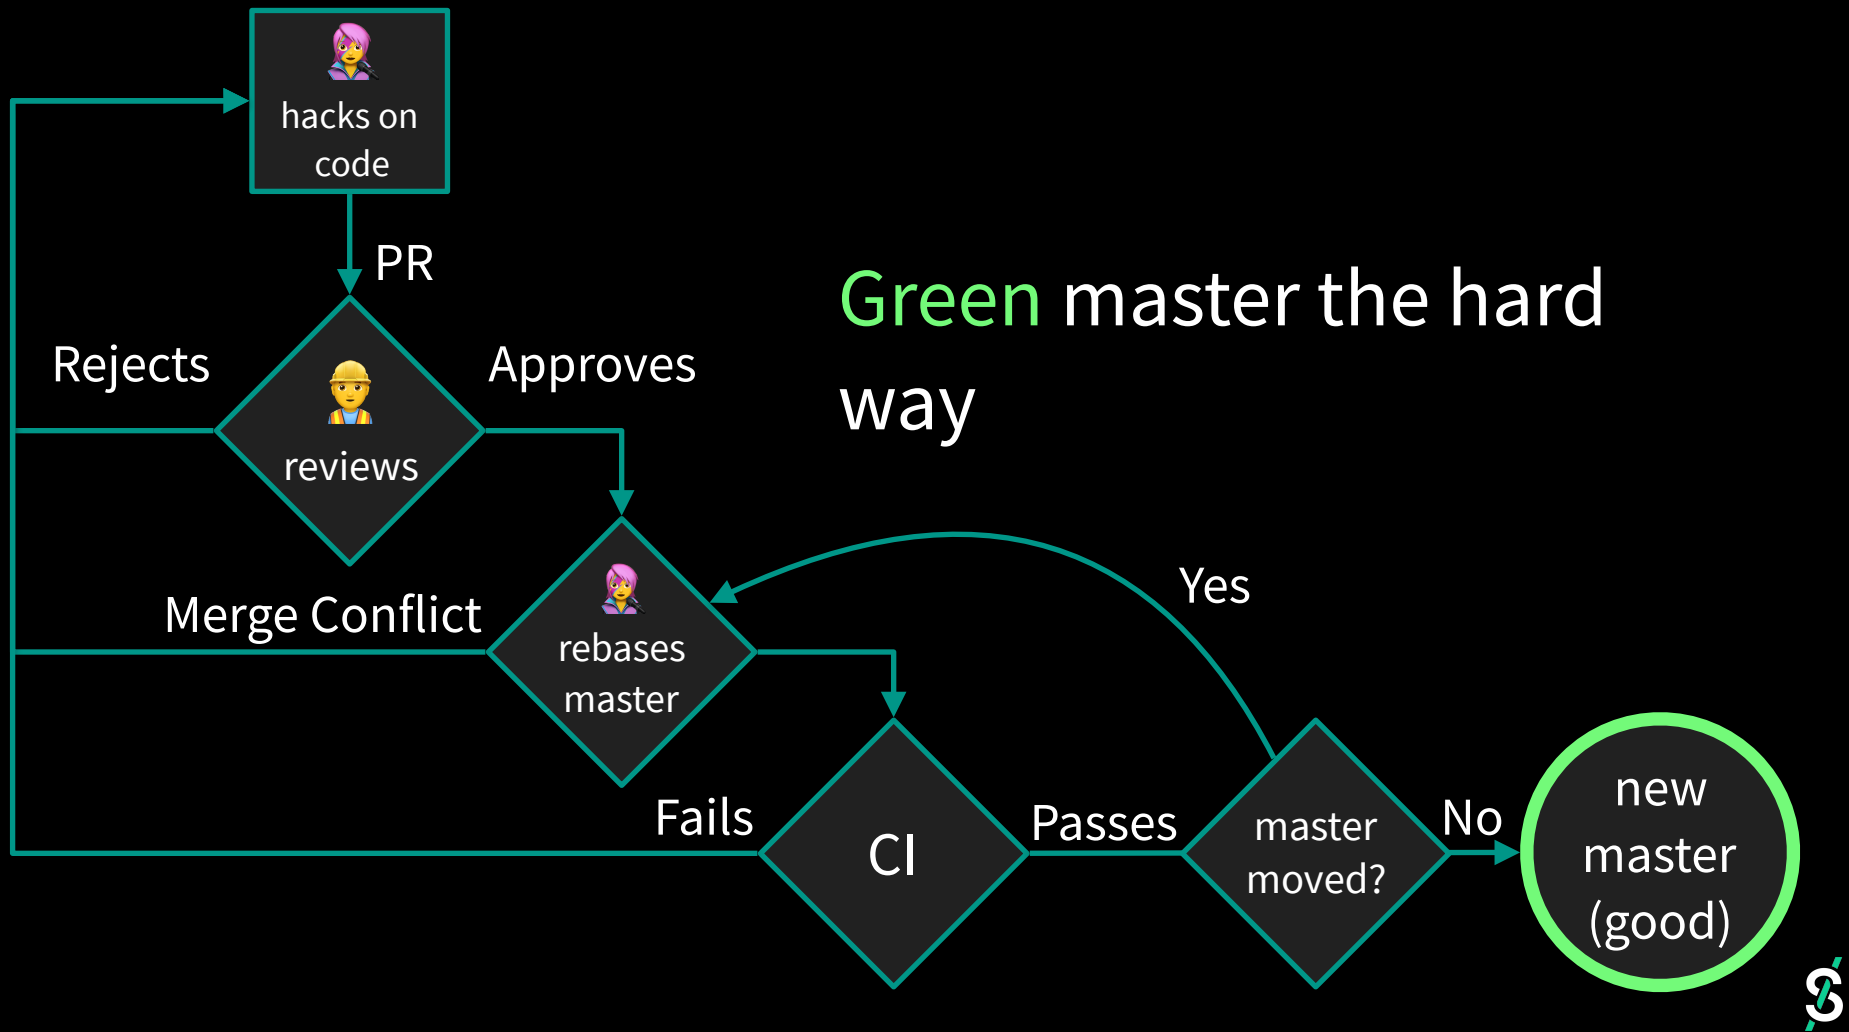
Merge method

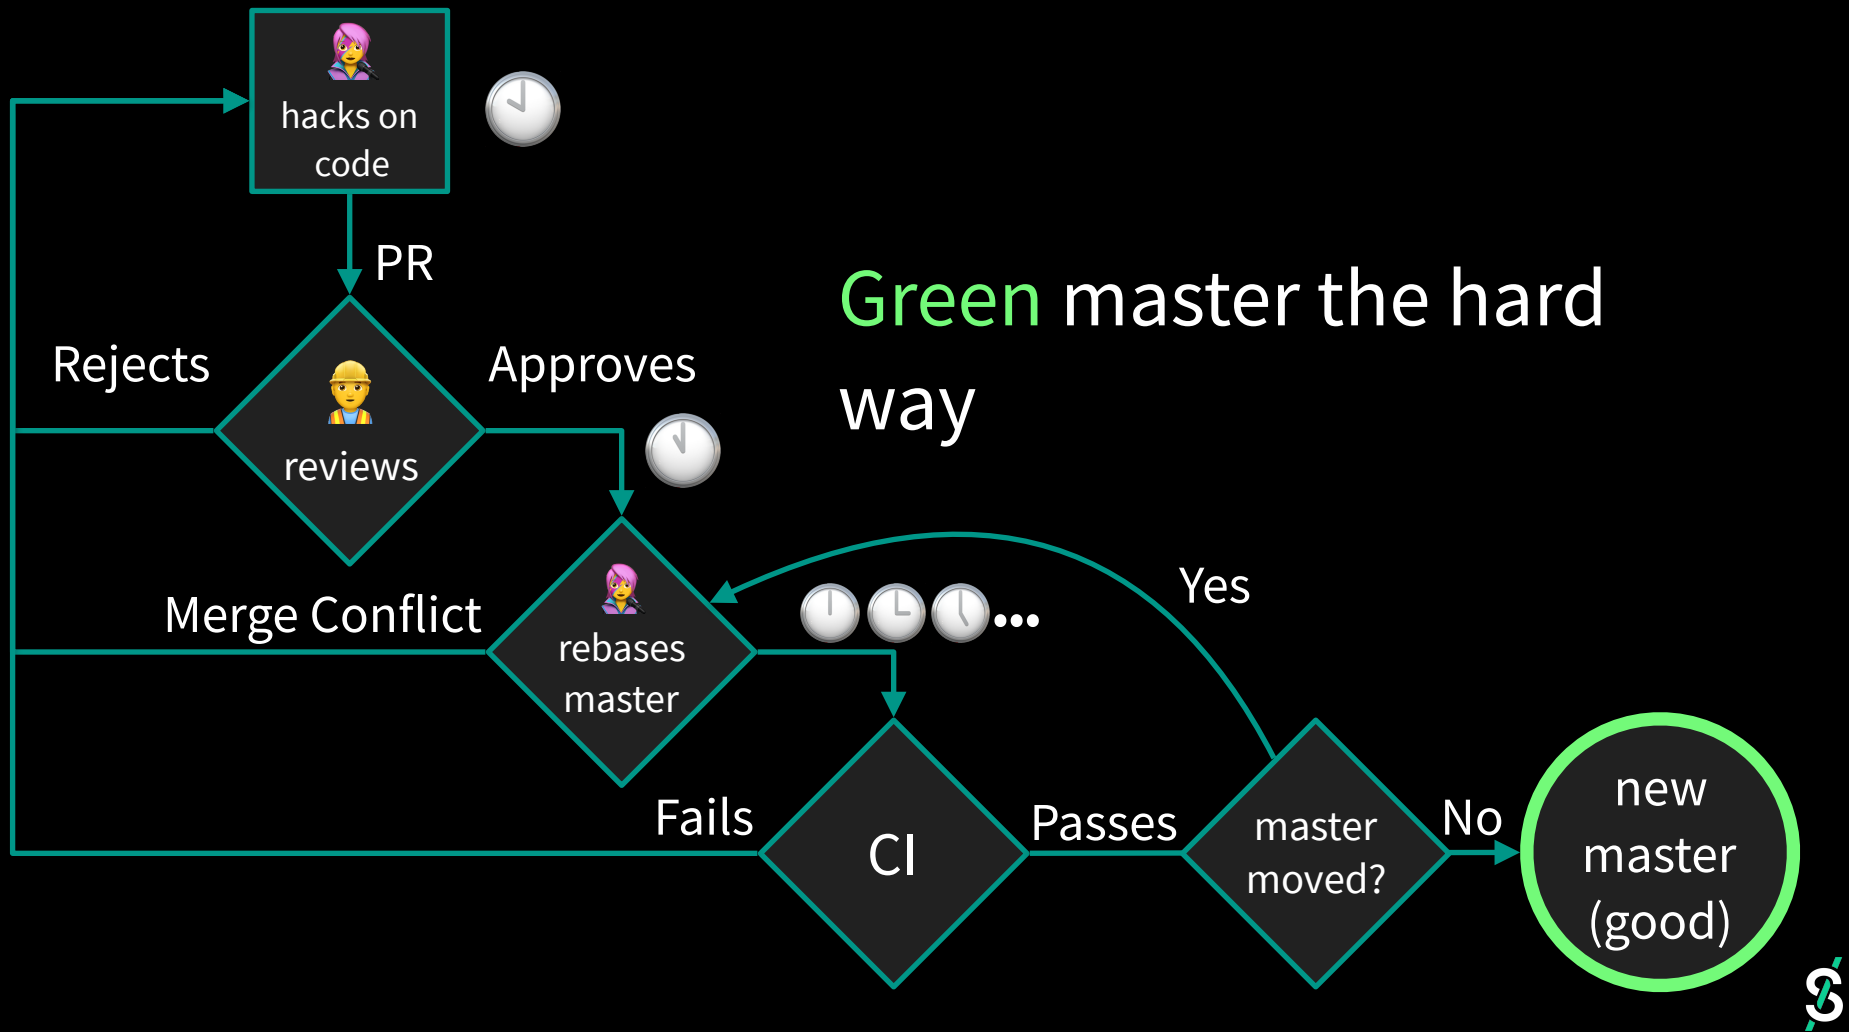
- ☐ **Merge commit**
A merge commit is created for every merge, and merging is allowed as long as there are no conflicts.
- ☐ **Merge commit with semi-linear history**
A merge commit is created for every merge, but merging is only allowed if fast-forward merge is possible. This way you could make sure that if this merge request would build, after merging to target branch it would also build.
When fast-forward merge is not possible, the user is given the option to rebase.
- ☒ **Fast-forward merge**
No merge commits are created and all merges are fast-forwarded, which means that merging is only allowed if the branch could be fast-forwarded.
When fast-forward merge is not possible, the user is given the option to rebase.

Default description template for merge requests ?

Description parsed with [GitLab Flavored Markdown](#).

- ☐ **Activate merge request approvals ?**
Merge request approvals allow you to set the number of necessary approvals and predefine a list of approvers that you will need to approve every merge request in a project.
- ☒ **Only allow merge requests to be merged if the pipeline succeeds**
Pipelines need to be configured to enable this feature. ?
- ☐ **Only allow merge requests to be merged if all discussions are resolved**
- ☒ **Show link to create/view merge request when pushing from the command line**





Sad 🙋

Sad

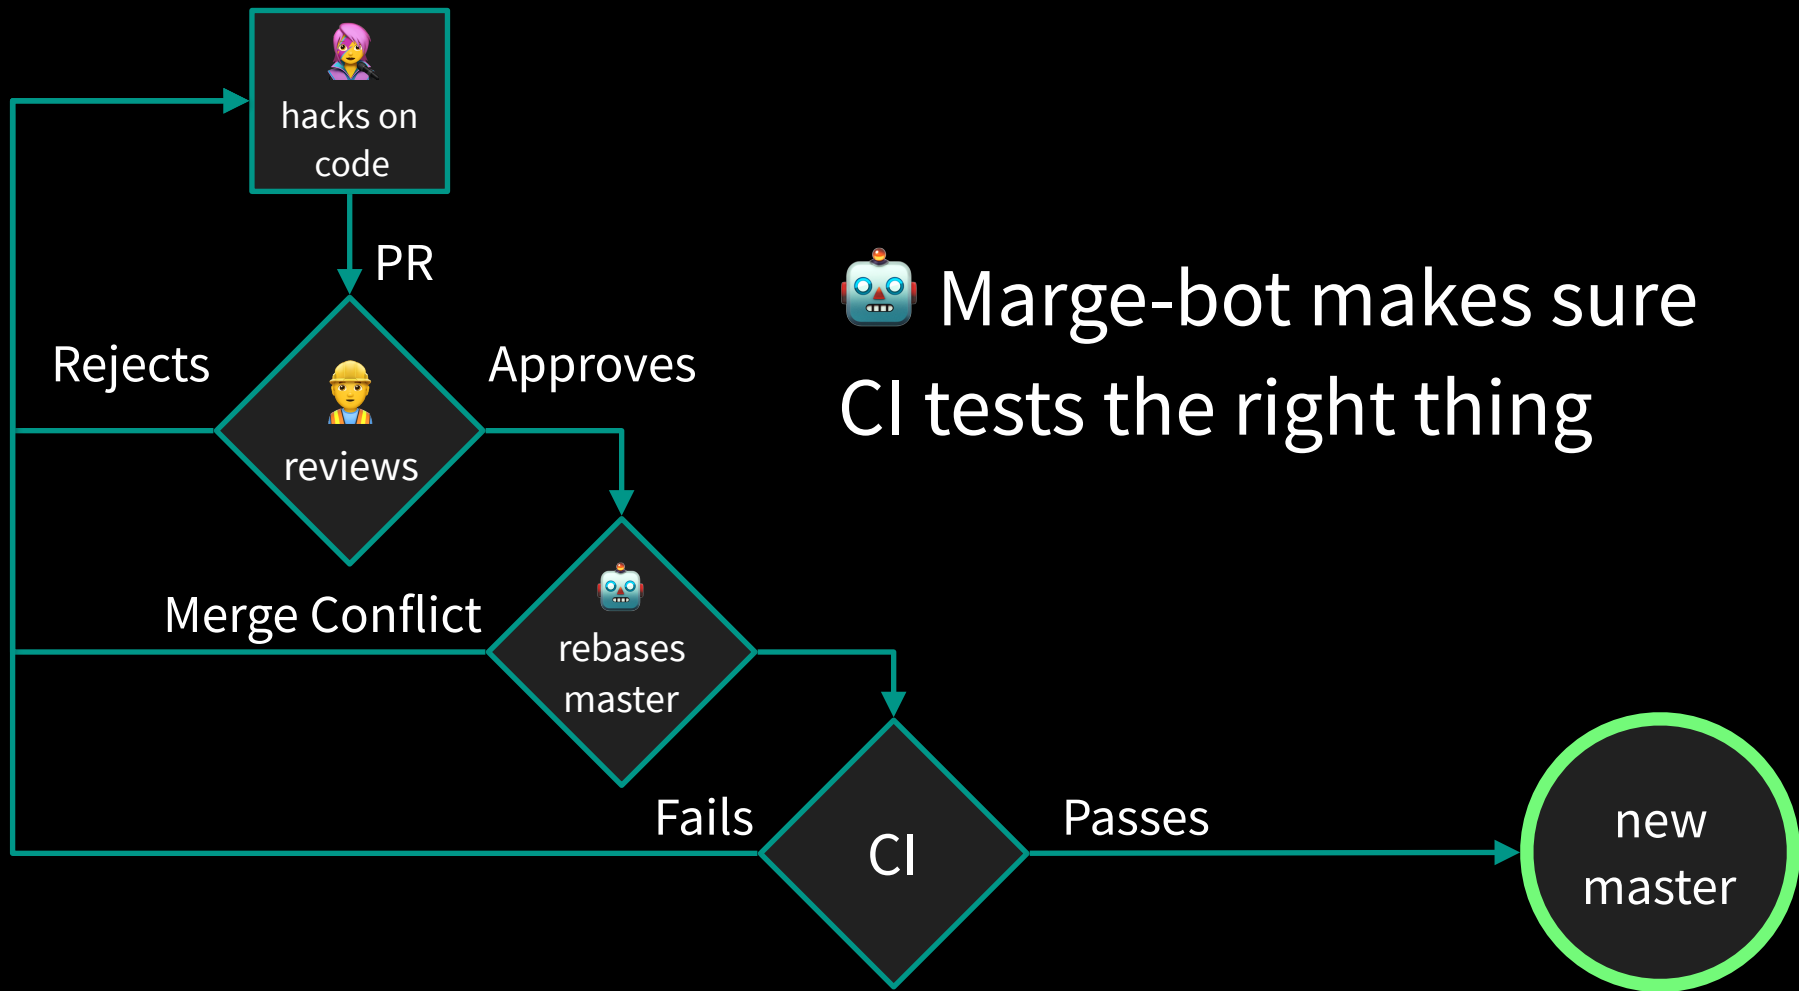



Sad 

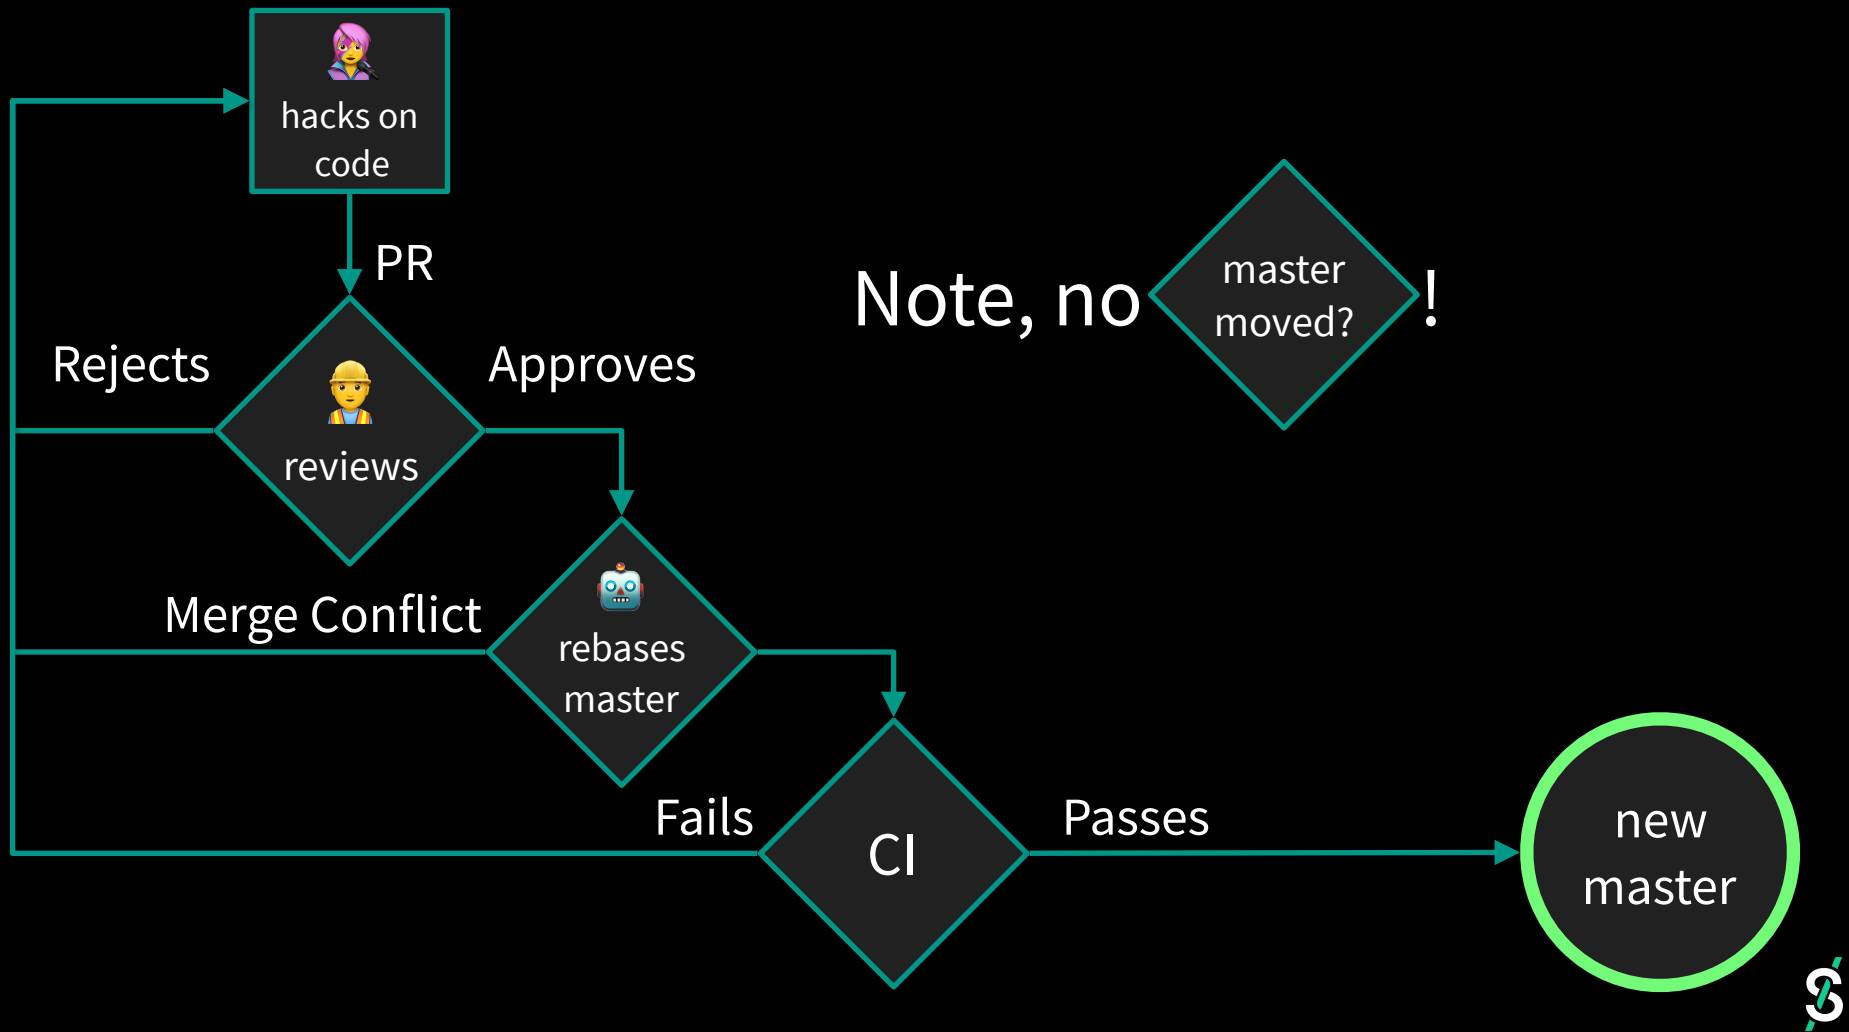
writes 

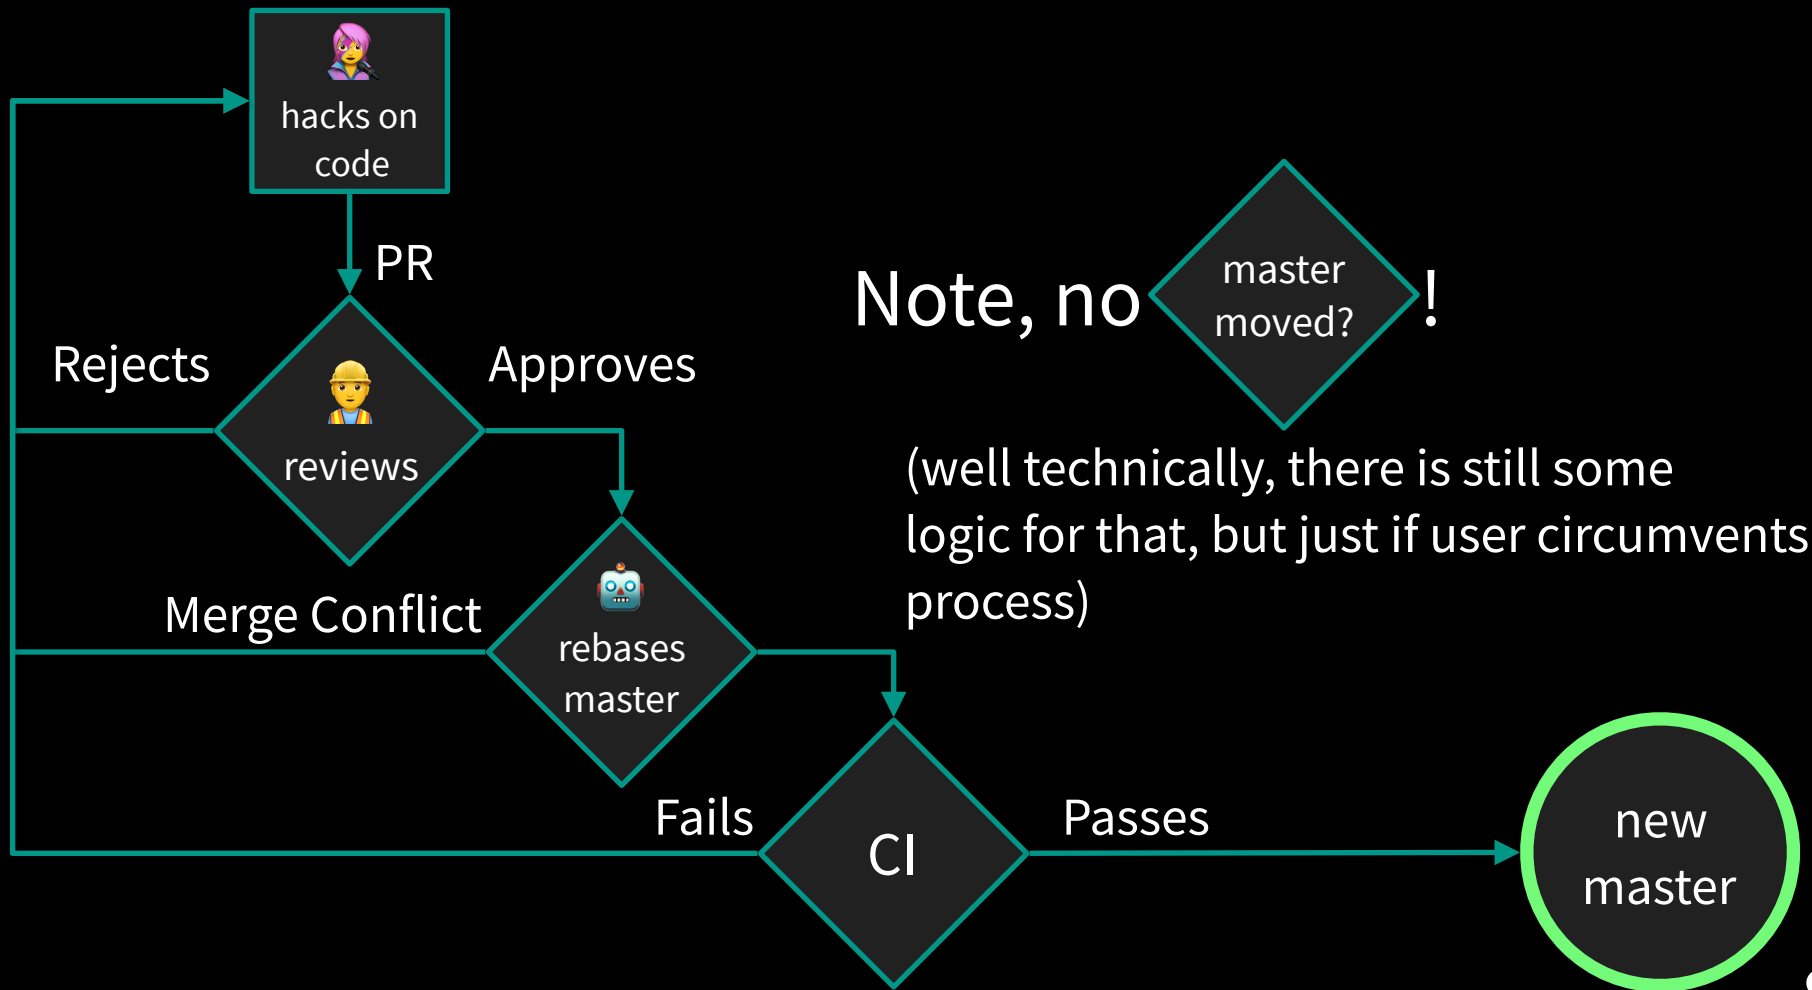
**This is the
story of this** 🤖







 Marge-bot makes sure
CI tests the right thing









How green PRs can break master









Going from **Green** PRs to **Red** master

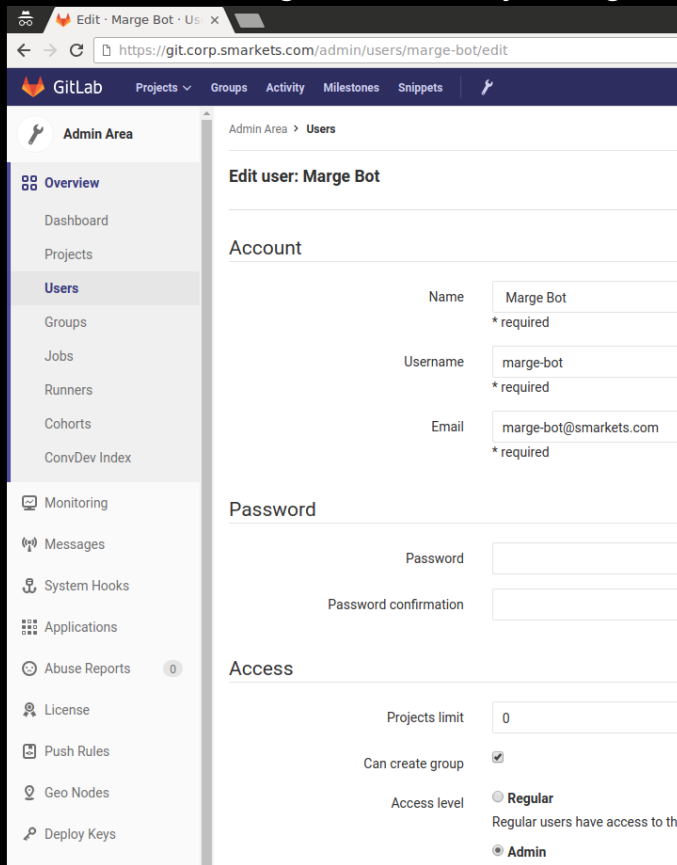
- **New wine into old skins:** PR 1 changes a function's API and all its call-sites, PR 2 introduces a new call site
- **Outdated coverage:** PR1 improves test coverage, PR2 changes the API
- **Fragile Baseclass:** PR1 makes internal changes to how Klass is implemented; PR2 adds some functionality to SubKlass that breaks because e.g. Klass.f no longer calls Klass.g internally.

None of these will cause **merge** conflicts or (feature-branch) CI failures; you'll find out there is a **logical** conflict after successful merge to master.



Running Marge-bot (in 3mins or less)

1. Create marge-bot ssh key and gitlab account/token



GitLab Admin Area

Admin Area > Users

Edit user: Marge Bot

Account

Name: Marge Bot *required

Username: marge-bot *required

Email: marge-bot@smarkets.com *required

Password

Password:

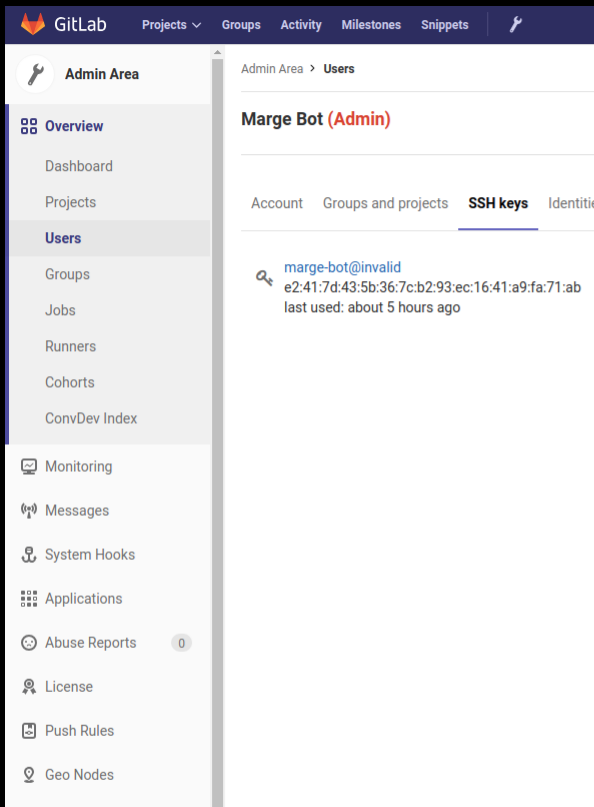
Password confirmation:

Access

Projects limit: 0

Can create group: ☒

Access level: ☒ Regular ☐ Admin



GitLab Admin Area

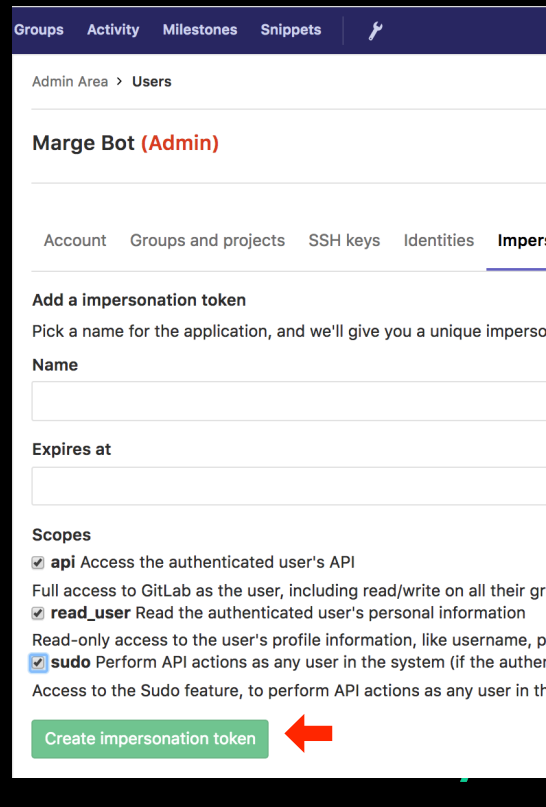
Admin Area > Users

Marge Bot (Admin)

Account Groups and projects SSH keys Identities

SSH keys

marge-bot@invalid
e2:41:7d:43:5b:36:7c:b2:93:ec:16:41:a9:fa:71:ab
last used: about 5 hours ago



GitLab Admin Area

Admin Area > Users

Marge Bot (Admin)

Account Groups and projects SSH keys Identities Impersonation

Add a impersonation token

Pick a name for the application, and we'll give you a unique impersonation token

Name:

Expires at:

Scopes

☒ api Access the authenticated user's API

☒ read_user Read the authenticated user's personal information

☒ sudo Perform API actions as any user in the system (if the authenticated user has access to the Sudo feature, to perform API actions as any user in the system)

Create impersonation token

Running Marge-bot (in 3mins or less)

1. Create marge-bot ssh key and gitlab account/token

2.

```
docker run --restart=on-failure \  
  --env-file=<(\  
    echo MARGE_AUTH_TOKEN="$(cat marge-bot.token)";\  
    echo MARGE_SSH_KEY="$(cat marge-bot-ssh-key)") \  
  smarkets/marge-bot \  
  --gitlab-url='http://your.gitlab.instance.com'
```



Running Marge-bot (in 3mins or less)

1. Create marge-bot ssh key and gitlab account/token

2.

```
docker run --restart=on-failure \  
  --env-file=<(\  
    echo MARGE_AUTH_TOKEN="$(cat marge-bot.token)";\  
    echo MARGE_SSH_KEY="$(cat marge-bot-ssh-key)") \  
  smarkets/marge-bot \  
  --gitlab-url='http://your.gitlab.instance.com'
```

3. Add marge-bot as a user (dev or master) to your project(s)



Demo of assigning PR to marge-bot



Conceptual Fix is simple

- Maintain a queue of pull requests
- Before merging a PR
 - rebase latest master into PR branch
 - wait for CI to pass
- Profit!
 - master will be always green (unless some tests are flaky)



Making it work practice (Usability/Familiarity)



- **Familiarity**
 - use normal gitlab flow, but assign to Marge-bot rather than pressing “merge after CI passes” (will make sure CI passed and branch has been reviewed)
 - Fair amount of behind-the-scenes work to bend Gitlab API to our will (not designed with this use case in mind; race conditions, need)
- **Usability**
 - Gitlab user name is “ marge-bot” (initial space, sorts first in list of users, so quick to assign to)
 - Marge leaves comments telling you if there is a problem (CI failed, no approval, conflict...)
 - we have a Slack channel that shows the Merge queue maintained by Marge (so place in queue/ETA easy to find)






Making it work practice (Scalability)

- **Scalability** (via Batching)
 - rebasing all open PRs on merge/rebase to master and re-running CI works fine for up to a dozen devs and CI tests that take a few mins
 - beyond that too much load on gitlab and CI build slaves
 - solution: create a “synthetic” batch merge request of top of queue PRs that have passed branch CI already; make synthetic branch top of queue
 - omit PRs that cause merge conflicts
 - if tests pass, merge all individual PRs (bypassing CI)
 - if tests fail, split
 - in either case, throw away “synthetic” branch










Omit
failed CI &
merge conflicts




Temp branch

**That's
basically it!**

Green Master!





But wait
THERE'S MORE!












Our workflow requirements at Smarmets



Productivity Requirements

Code must flow:

- 70 devs, 11 teams, ~130 services
- Commits every few mins
- ~10 ships to prod/day
- want to preserve velocity

Your Projects		Starred Projects				
All	Push events	Merge events	Issue events	Comments	Team	
	Marge Bot pushed to branch mobile-add-card-for-everyone at smarkets / smarkets				less than a minute ago	
c7e381b4 · mobile: allow all users to access add new card page						
	Valerio pushed new branch mobile-add-card-for-everyone at smarkets / smarkets				4 minutes ago	
	Daniel Ghita pushed to branch dan.market-graphs at smarkets / smarkets				4 minutes ago	
ba47074a · web: add market charts						
	Sam Jones pushed to branch popular-categories-change at smarkets / smarkets				5 minutes ago	
65e4aca4 · marketing: Fix popular events logging						
	Sam Jones pushed to branch popular-categories-change at smarkets / smarkets				8 minutes ago	
cd89ce89 · marketing: Fix popular events logging ... and 16 more commits. Compare 7cca9a30...cd89ce89						
	Tornike Gogniashvili pushed to branch push-avp at smarkets / smarkets				10 minutes ago	
7ed90311 · notification: Few fixes and refactors						
	Marge Bot pushed to branch master at smarkets / smarkets				20 minutes ago	
67d917ad · Golf market names						
	Marge Bot pushed to branch golf_name at smarkets / smarkets				21 minutes ago	
67d917ad · Golf market names ... and 16 more commits. Compare 6588be23...67d917ad						
	Marge Bot pushed to branch master at smarkets / smarkets				21 minutes ago	
e5393f00 · mobile: bump version number 2.1.0 (4) ... and 11 more commits. Compare 91b982ea...e5393f00						
	Marge Bot pushed to branch release-adjustments at smarkets / smarkets				22 minutes ago	
e5393f00 · mobile: bump version number 2.1.0 (4) ... and 17 more commits. Compare fbbd1c29...e5393f00						
	Marge Bot pushed to branch marge_bot_batch_merge_job at smarkets / smarkets				28 minutes ago	
1b5cc6c4 · Golf market names						



Audits & Auditors

- We're in a regulated industry
- Requirements vary by country
- Goal #1: meet **all** at once
- Goal #2: don't cripple workflow



Regulatory Requirements

Auditors may want to know:

- Who wrote this code and when?
- Who signed it off?
- How was it validated?
- Why was it needed?
- When was it deployed and by whom?
- Who approved the deployment?

In fact fact **you** might well want to know these things, even if you're not audited!



Regulatory Requirements

Auditors may want to know:

- Who wrote this code and when?
- Who signed it off?
- How was it validated?
- Why was it needed?
- When was it deployed and by whom?
- Who approved the deployment?

Git-out of the box

Gitlab + Marge-bot (git trailers)

Our Ship tool + git notes

In fact fact **you** might well want to know these things, even if you're not audited!



Common thread: compliance centred around git

- Cryptographic
- Familiar
- Flexible
- Platform agnostic



Git workflows

Theory vs Practice



Theory

- Main repo as a collection of **subrepos** of independently developed (micro-)services/libs is a **Good Idea**
 - Macro-view of all that will run in prod in the main repo
 - Micro-view of individual services in subrepo
 - just your service's history/code, no need to check out GiBs etc.
 - can simplify and speed up CI
- **Merging** (feature) branches (into master) is a **Good Idea**
 - Macro-view of features on master
 - Micro-view of implementation steps in branch
 - history is inviolate!



Practice



Google Image Search for Git submodules



**SAY GIT SUBMODULE UPDATE --INIT --
RECURSIVE**



ONE MORE TIME

memegenerator.net

**EVERY TIME YOU USE GIT
SUBMODULE**



GOD KILLS A SMALL KITTEN

memegenerator.net

**AIN'T NOBODY GOT TIME FOR
DAT**

memegenerator.net


USERNAME



LLION OF APPS


DAYA

**PATA KARO GIT SUBMODULE
UPDATED HAI KI NAHI**




memegenerator.net

I CLONE REPO



WHY SUBMODULES ARE EMPTY

Just go Monorepo, don't resist assimilation





Reverting PRs the marge-bot way

```
git revert-mr 123
```




**Which one would
you rather do when
prod is broken?**



Now *which* PR actually broke this? (plain git way)

- I know, I'll use `git bisect` run to find out



- Only look at merge commits to master? **Needs 3rd party tool.**
- Feature branch commits will often not have passed CI, will get false positives.

Now *which* PR actually broke this? (marge-bot way)

```
git bisect-run-tested ./test.sh
```

Only runs on commits that passed CI

(marge-bot can add Tested-by: trailer to last commit of PR if you have mandatory CI switched on in Gitlab)



**How does it
work?**



Recap: bare bones marge setup

```
docker run --restart=on-failure \  
  --env-file=<(\  
    echo MARGE_AUTH_TOKEN="$(cat marge-bot.token)";\  
    echo MARGE_SSH_KEY="$(cat marge-bot-ssh-key)") \  
  smarkets/marge-bot \  
  --gitlab-url='http://your.gitlab.instance.com'
```



More --args, more features!

```
docker run --restart=on-failure \  
  --env-file=<(\  
    echo MARGE_AUTH_TOKEN="$(cat marge-bot.token)";\  
    echo MARGE_SSH_KEY="$(cat marge-bot-ssh-key)") \  
  smarkets/marge-bot \  
  --gitlab-url='http://your.gitlab.instance.com' \  
  --batch \           Try to optimistically batch several PRs for faster CI\  
  --add-part-of \     Adds a "Part-of: <PR-url>" Trailer to commits\  
  --add-tested \      Adds a "Tested-by: <PR-url>" Trailer to final commit in PR\  
  --add-reviewers \   Adds a "Reviewed-by: <guy who approved PR>" Trailer to commits\  
  --impersonate-approvers    Re-approve after rewriting commits to add Trailers
```



Bare bones commit

```
commit 5ddc293ec55408ecc101eacf6495421a16182633
Author: Jaime Lennox <jaime.lennox@smarkets.com>
Date:   Mon Jul 23 11:46:41 2018 +0100
```

```
marge-bot: bump to version 0.7.0
```

There's a new version of Marge available, so let's update our in-house version to match.



Optional Features: Audit with --add-reviewers

```
commit 146891a956fd35cf8ab6445d7ec76fddf4230925
Author: Jaime Lennox <jaime.lennox@smarkets.com>
Date:   Mon Jul 23 11:46:41 2018 +0100
```

```
marge-bot: bump to version 0.7.0
```

There's a new version of Marge available, so let's update our in-house version to match.

```
Reviewed-by: Tornike Gogniashvili <tornike.gogniashvili@smarkets.com>
```



Optional Features: Bisect with --add-tested

```
commit ca582b7ab9f03f496509291b1fa2e8f768a76f05
Author: Jaime Lennox <jaime.lennox@smarkets.com>
Date:   Mon Jul 23 11:46:41 2018 +0100
```

```
marge-bot: bump to version 0.7.0
```

There's a new version of Marge available, so let's update our in-house version to match.

```
Reviewed-by: Tornike Gogniashvili <tornike.gogniashvili@smarkets.com>
```

```
Tested-by: <https://git.corp.smarkets.com/smarkets/smarkets/merge\_requests/9727>
```



Optional Features: Revert with --add-part-of

```
commit 088bf8546b73b559322a8744e867cf8949fe6225
Author: Jaime Lennox <jaime.lennox@smarkets.com>
Date:   Mon Jul 23 11:46:41 2018 +0100
```

```
marge-bot: bump to version 0.7.0
```

There's a new version of Marge available, so let's update our in-house version to match.

```
Reviewed-by: Tornike Gogniashvili <tornike.gogniashvili@smarkets.com>
Tested-by: <https://git.corp.smarkets.com/smarkets/smarkets/merge\_requests/9727>
Part-of: <https://git.corp.smarkets.com/smarkets/smarkets/merge\_requests/9727>
```



Then it's just a bunch of git aliases!

```
git config --global alias.bisect-run-tested \  
  'f() { git bisect run /bin/sh -c  
    "if !(git log -1 --format %B  
      | fgrep -q \"Tested-by: Marge Bot\");  
    then exit 125;  
    else \"$@\"; fi\"; }; f'
```

```
git config --global alias.mr-revs \  
  '!f() { git log --grep "^Part-of.*/" "$1" --pretty="%H"; }; f'  
git config --global alias.mr-url \  
  '!f() { git log -1 --grep "^Part-of.*/" "$1" --pretty="%b" |  
    grep "^Part-of.*/" "$1" | sed "s/.*<\\(.*\\)>/\\1/"; }; f'  
git config --global alias.revert-mr \  
  '!f() { REVS=$(git mr-revs "$1"); URL=$(git mr-url "$1");  
    git revert --no-commit $REVS;  
    git commit -m "Revert <$URL>$(echo;echo; echo "$REVS"  
      | xargs -I% echo "This reverts commit %.")"; }; f'
```



What if you can't use Marge-bot (not on Gitlab)?



What if you can't use Marge-bot (not on Gitlab)?

- At least now you know what you're missing ;)
- If you don't need something general, roll-your-own is often pretty easy
 - we more or less did that at my last 3 employers
- Also, we welcome PRs to <https://github.com/smarkets/marge> (github backend would be cool, and probably straightforward!)
- Similar tools might exist; e.g. for github there's also Rust's [homu](#)



Summary

- A good PR workflow runs tests against “future” master not just the feature branch
- good \neq common (but now you know you want it and how to get it!)
- <https://github.com/smarkets/marge-bot> will do it for gitlab, the way you want (merge or rebase-based)
- Marge-bot can also add Trailers to show who Reviewed commit and what PR
 - Combines best of Merge and Rebase based workflows (e.g. you can still see what PR commits belonged to)
 - Extra perks:
 - git bisect that actually works (at PR level)
 - git revert that actually works (at PR level)
- In practice: Monorepo and rebasing PRs > subrepos and merging PRs (usually)



What's wrong with this picture?

README.md

build passing

Marge-bot

Marge-bot is a merge-bot for GitLab that, beside other goodies, implements [the Not Rocket Science Rule Of Software Engineering](#):

automatically maintain a repository of code that always passes all the tests.

— Graydon Hoare, main author of Rust

This simple rule of thumb is still nowadays surprisingly difficult to implement with the state-of-the-art tools, and more a way that scales with team size (also see our [blog post](#)).



What's wrong with this picture?

README.md

build passing

Marge-bot

Marge-bot is a merge-bot for GitLab that, beside other goodies, implements [the Not Rocket Science Rule Of Software Engineering](#):

automatically maintain a repository of code that always passes all the tests.

— Graydon Hoare, main author of Rust

This simple rule of thumb is still nowadays surprisingly difficult to implement with the state-of-the-art tools, and more a way that scales with team size (also see our [blog post](#)).



Credits

- [Daniel Gorin](#) (initial design & implementation)



- [Jaime Lennox](#) (batching, slack integration, general maintenance)



- [Alexander Schmolck](#) (trailers, nix-based build, general maintenance)

- [Marian Rusu](#) (batching, build & CI improvements)

- all our users who submitted PRs, suggestions and bug reports



Extra Material



Gitlab Wish list

- Commit message changes should not trigger CI
- Rebasing the target branch into the source branch should not reset approvals
- Getting approver email should not require excessive perms
- Merge API should
 - take optional expected hash of master (don't merge if it doesn't match)
 - allow overriding CI trigger and force merge (e.g. “trust me, it's tested” for batch mode)
- PRs from forks are kinda broken
 - CI is on fork (which probably hasn't the right setup)



Implementation details...

... that I think worked well



Architecture

- keeping it simple:
 - stateless
 - no concurrency
 - crash on network or Gitlab failures (HTTP 50x)
 - rely on
 - ``docker run --restart=on-failure`` or `systemd`
- Advantages:
 - simple! (no messy python-style concurrency, linear log of actions easy to grok)
- Disadvantages:
 - extra requests/slower (e.g. cloning big repo again after crash); not an issue for us so far
- yield-on-sleep and retrying requests might be good complexity/benefit trade-off



ConfigArgParse: mix and match --args, ENV and config files

- Marge accepts most options as env-var --arg or yaml config option
- Various benefits
 - quickly override option in config file for test (e.g. new auth token)
 - config file best for complex setup, command line args sufficient for vanilla setup
- Can still customize:
 - we disallow passing secrets as commandline args for security reasons



Derive from namedtuples for lightweight biz logic classes

```
class Version(namedtuple('Version', 'release edition')):
    @classmethod
    def parse(cls, string):
        release_string, edition = string.split('-', 1)
        release = tuple(int(number) for number in release_string.split('.'))
        return cls(release=release, edition=edition)

    @property
    def is_ee(self):
        return self.edition == 'ee'
```

(constructor, repr, updated-copy and enforced immutability for free)



Better mocking with state machines

```
class MockLab(object):
    def __init__(self, gitlab_url=None):
        self.api = api = ApiMock('...', initial_state='initial')

        api.add_transition(GET('/version'), Ok({'version': '9.2.3-ee'}))
        # [...]

        api.add_transition(
            GET('/projects/1234/repository/commits/%s' % rewritten_sha),
            Ok(commit_after_pushing),
            from_state='pushed', to_state='passed',
        )
        api.add_transition(
            GET('/projects/1234/repository/commits/%s' % rewritten_sha),
            Ok(_commit(id=rewritten_sha, status='success')),
            from_state=['passed', 'merged'],
        )
```

(Nicer than typical mocking, IMO, but learning curve has kept it out of recent code)



nix: Completely reproducible (docker image) builds (w/o docker)

- fast, correct, small (pick 3)
- Travis supports it
- drives CI and images
- nix-shell --pure = virtualenv on steroids
 - also handles non-py deps, e.g. git, ssh
 - actually works

```
{ pkgs ? import ./pinnedNixpkgs.nix }:  
let callPackage = pkgs.lib.callPackageWith (pkgs);  
    marge = callPackage ./marge.nix {};  
    version = marge.version;  
in  
pkgs.dockerTools.buildImage {  
  name = "smarkets/marge-bot";  
  tag = "${version}";  
  # minimal user setup, so ssh won't whine 'No user exists for uid 0'  
  runAsRoot = ''  
    #!${pkgs.stdenv.shell}  
    ${pkgs.dockerTools.shadowSetup}  
    mkdir -p /root/.ssh  
    '';  
  contents = [marge pkgs.bash pkgs.coreutils pkgs.openssh pkgs.glibcLocales];  
  config = {  
    Entrypoint = [ "/bin/marge.app" ];  
    Env = [ "LANG=en_US.UTF-8" "'LOCALE_ARCHIVE=/lib/locale/locale-archive'" ];  
  };  
}
```