Fundamentals of Software Architecture @ Chalmers University

# Scaling DevOps - GitHub's journey from 500+ to 1500+ people













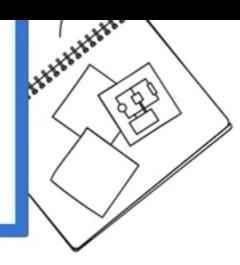


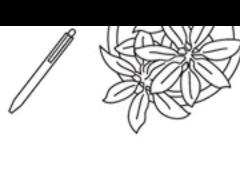
SCOUT 24



games like 2048—which inspired hundreds of clones—have invited even more people to play and create.

Apart from active game development, public source code for some of the most influential games of all time, including <a href="Doom">Doom</a> and <a href="Prince of Persia">Prince of Persia</a> are archived in public repositories —certainly a better fate than <a href="being buried in the New">being buried in the New</a> <a href="Mexican desert">Mexican desert</a>.





22 SI

#### GitHub Classroom starts school

GitHub Classroom makes it easier for teachers to distribute starter code and collect assignments on GitHub. Today, students at high schools, universities, and coding bootcamps are learning across 1.8 million Classroom repositories—but these are far from the only educational resources on GitHub.

From lists of resources to massive open online courses (MOOCs) like <a href="mailto:edX">edX</a> and <a href="ModCs">Udacity</a>, you've created thousands of ways to learn software development on GitHub. Courses at that top <a href="mailto:our list">our list</a> include <a href="mailto:Ada's Jumpstart program">Ada's Jumpstart program</a> and <a href="mailto:Stanford's TensorFlow Tutorials">Stanford's TensorFlow Tutorials</a>. STAT545 is teaching thousands of students to wrangle data, while <a href="mailto:CS50">CS50</a> is being adopted in classrooms across the United States.

In the last decade, community-driven programs like Django Girls have kicked off online tutorials, hosted in-person events, and broadened learning opportunities for students around the world. Our <u>Campus Experts</u> are also building tech communities on campuses in more than 15 countries.





### GitHub Now vs Then

Repositories

150 M+

vs. 36 M in 2016

Pull Requests

400 M+

vs. 130 M in 2017

Registered users

50 M+

vs. 14 M in 2016

Suggested fixes for security alerts

4 M+

vs. 800 K in 2018



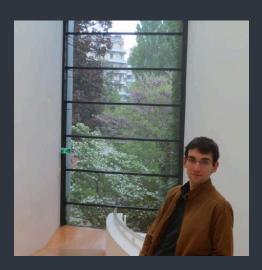
## Of software projects use open source

#### Code

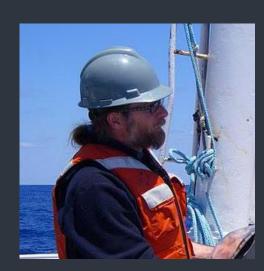
- 1 import kiwisolver
- 2 import numpy

3

#### Community contributors









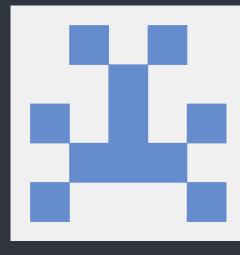








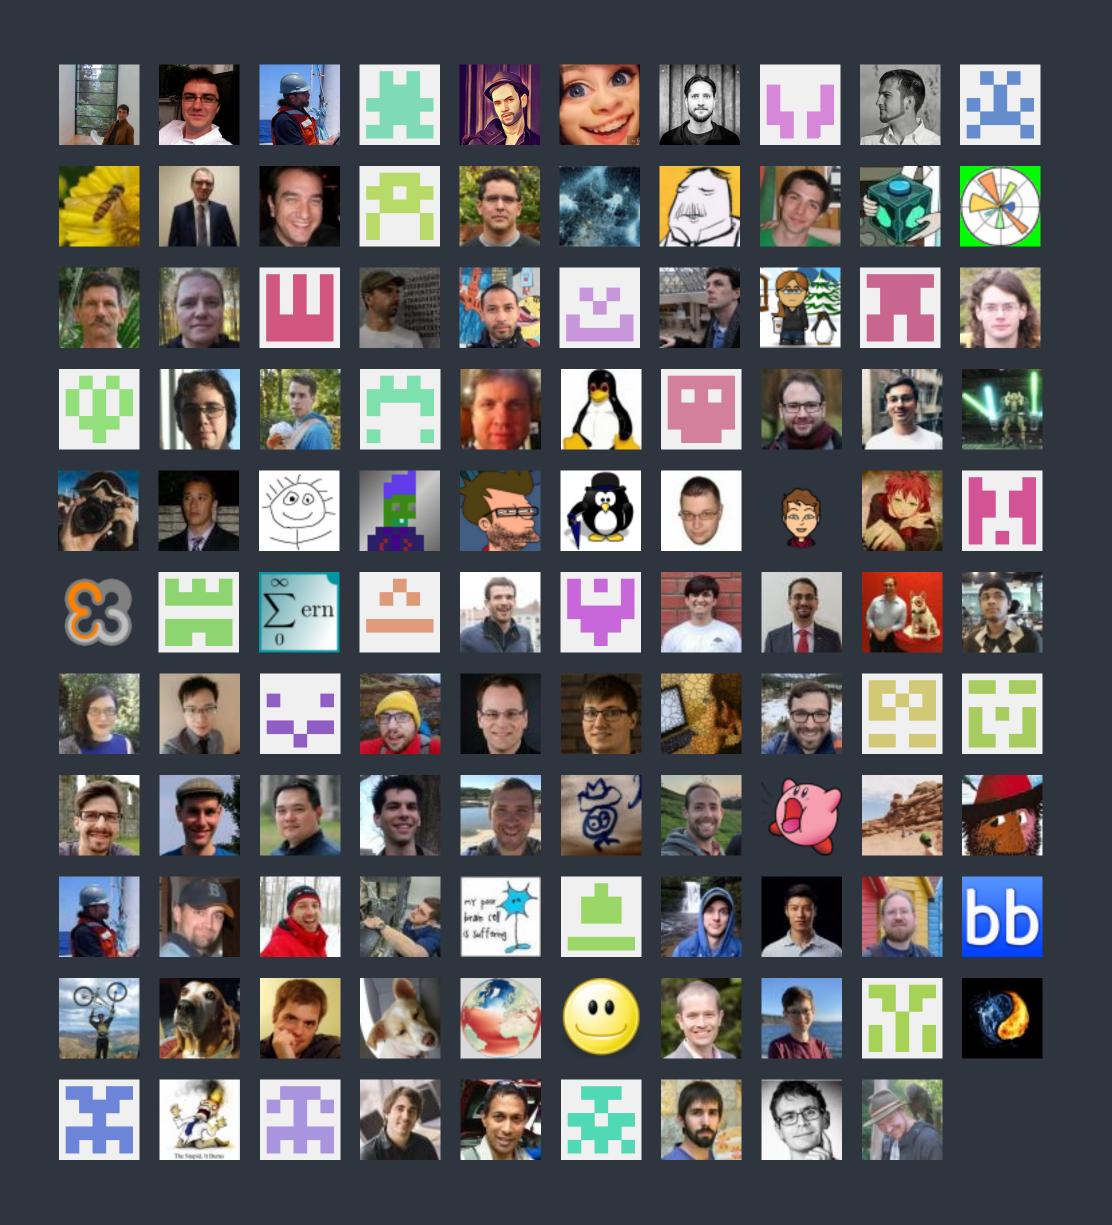




#### Code

- 1 import kiwisolver
- 2 import numpy
- 3 import matplotlib

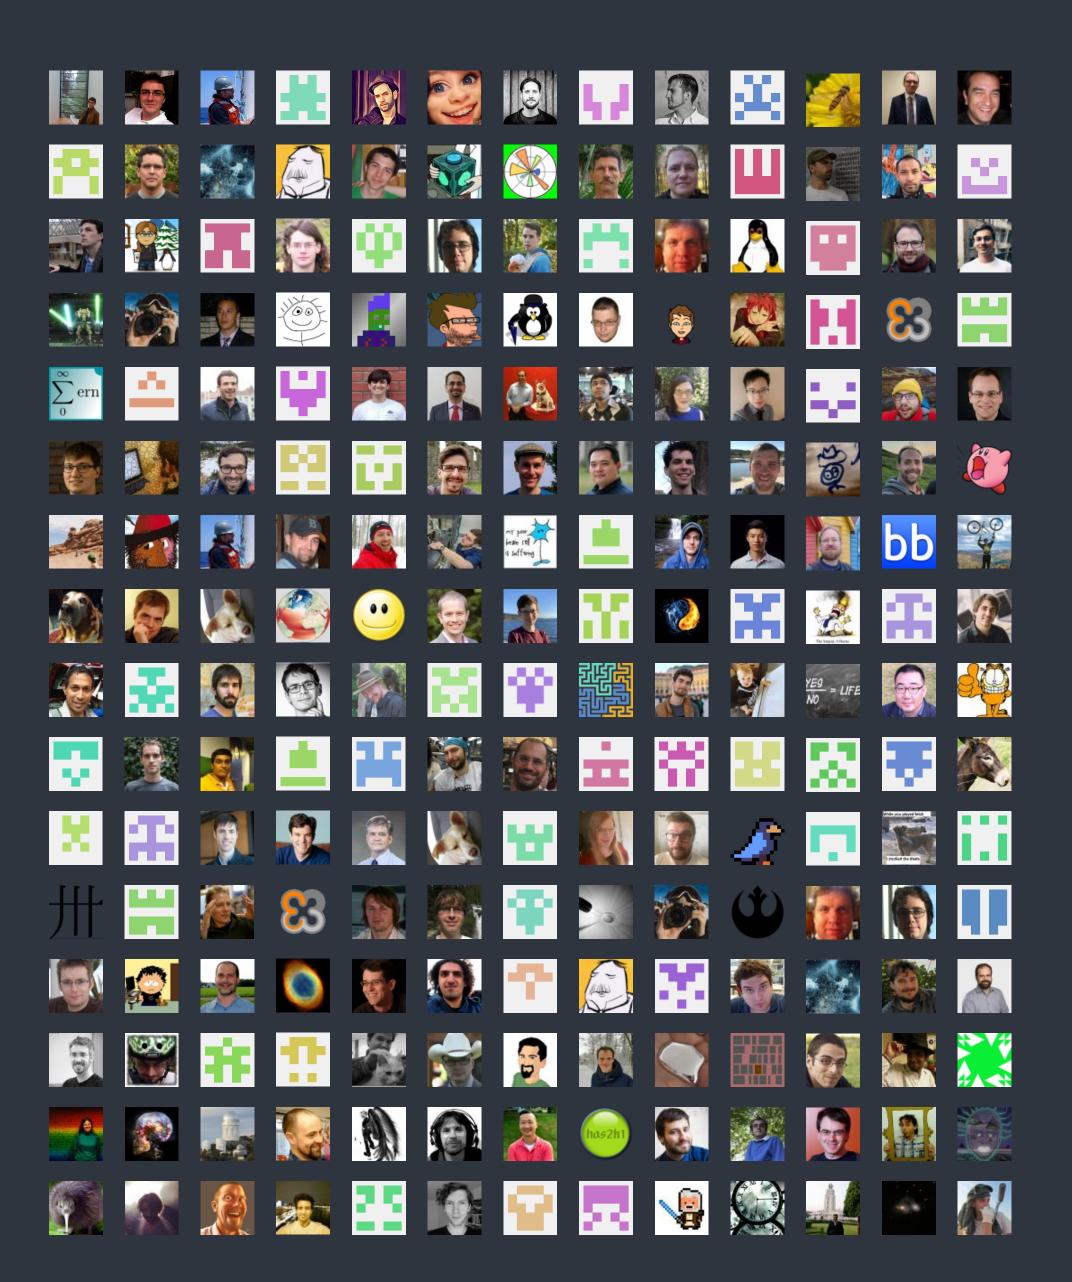
#### Community contributors



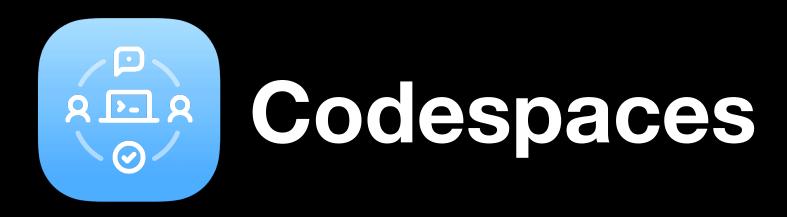
#### Code

- 1 import kiwisolver
- 2 import numpy
- 3 import matplotlib

#### Community contributors



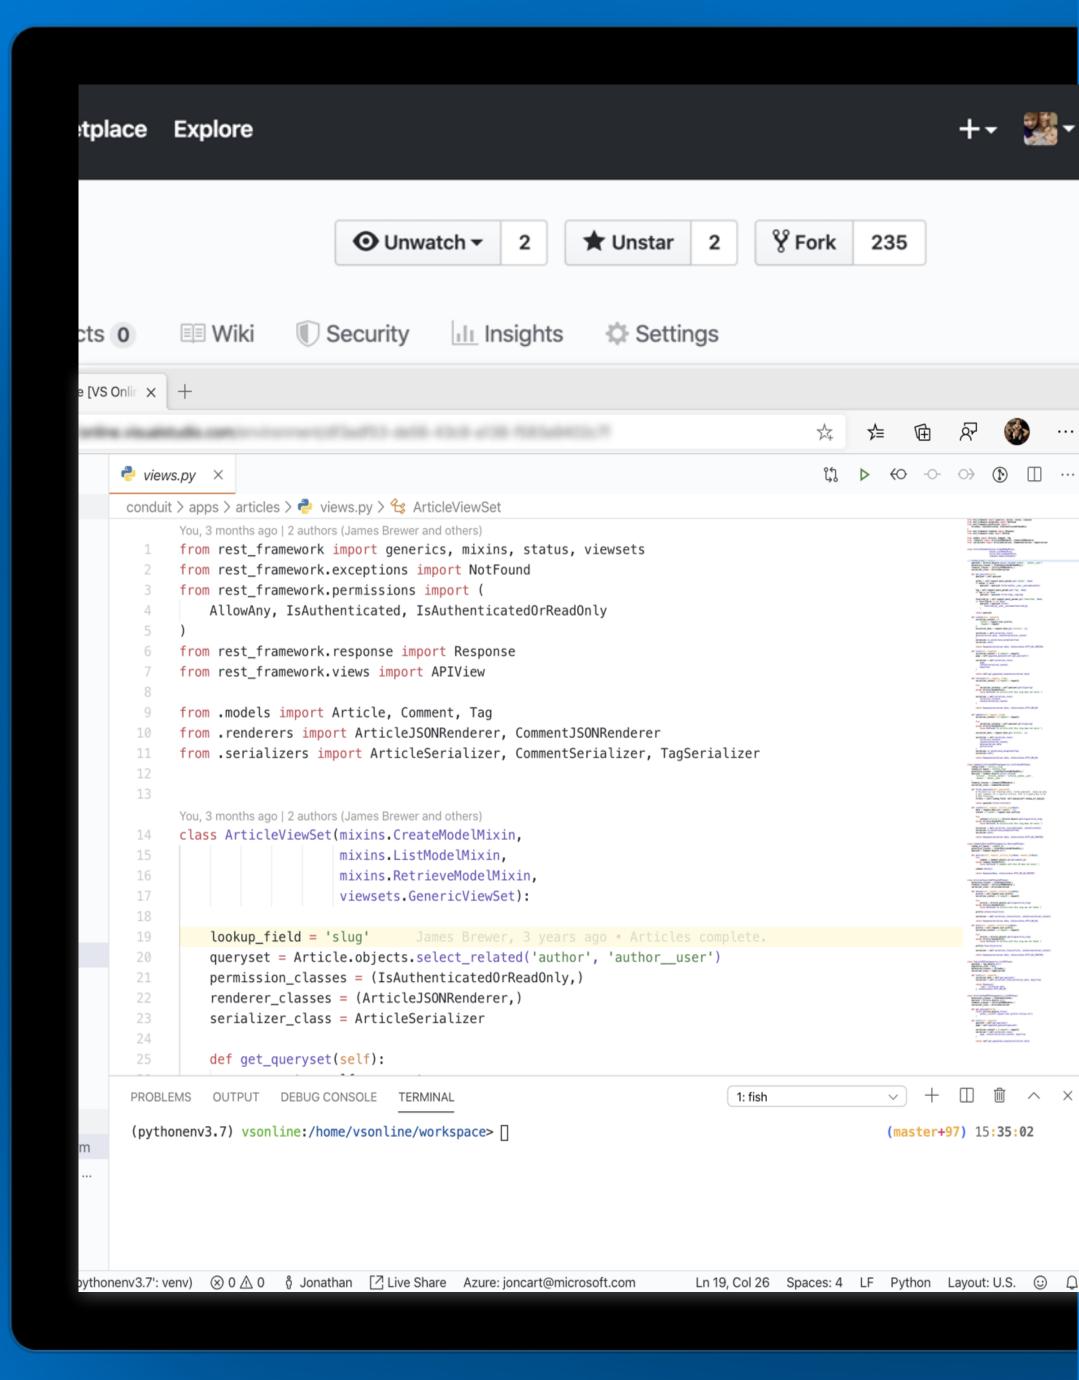




Reducing time to set up development environments from hours per week to seconds

Powered by web-based version of VS Code, or use with your desktop IDE

Fully integrated into the GitHub developer experience

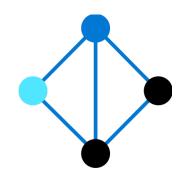




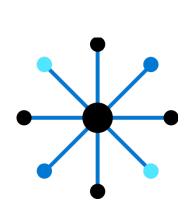




## Built with a shared engineering team



Azure DevOps and GitHub share the same leadership, reporting to the GitHub CEO



Bringing the requirements and insights of Azure DevOps customers to GitHub



Standardized tooling for 100k engineers at Microsoft



(i)

#### Test libzengithub package on main platforms

on: push

- **X** test (macOS-10.14)
- test (windows-2016)
- test (windows-2019)
- ✓ test (ubuntu-16.04)
- ✓ test (ubuntu-18.04)

- \$\langle \$\{\{\text{ format('Set up Python \{0\}', matrix.python-version) \}\}
- Install C and C++ header files

Complete job

- Install conan && conan package tools
- Generating conan user directory and building the solution

```
ZenGitHub/1.0@jonico/stable (test package): Running test()
1028
1029
                               . MMM
                  MMM.
1030
1031
                  MMMMMMMMMMMMMMM
1032
                 It's not fully shipped until it's fast. |
1033
                1034
1035
               MMMM::- -::::::- -::MMMM
1036
                MM~:~ 00~::::: ~ 00~:~MM
            .. MMMMM::.00:::+:::.00::MMMMM ..
1037
                 .MM::::: ._. :::::MM.
1038
                    MMMM;::::;MMMM
1039
1040
              -MM
              ^ M+
                      MMMMMMMM
1041
1042
                 MMMMMM MM MM MM
1043
                     MM MM MM
                     MM MM MM
1044
1045
                   .~~MM~MM~MM~MM~~.
1046
                 ~~~~MM:~MM~~~MM~:MM~~~~
1047
                ~~~~~==~==~~~
                ~~~~~==~==~==~~~
1049
                    :~==~==~=
1050
     [HOOK - attribute_checker.py] pre_export(): WARN: Conanfile doesn't have 'description'. It is recommended to add it as attribute
1052
```

**REMOTE** 

BASED IN SF

**NEW HIRES** 

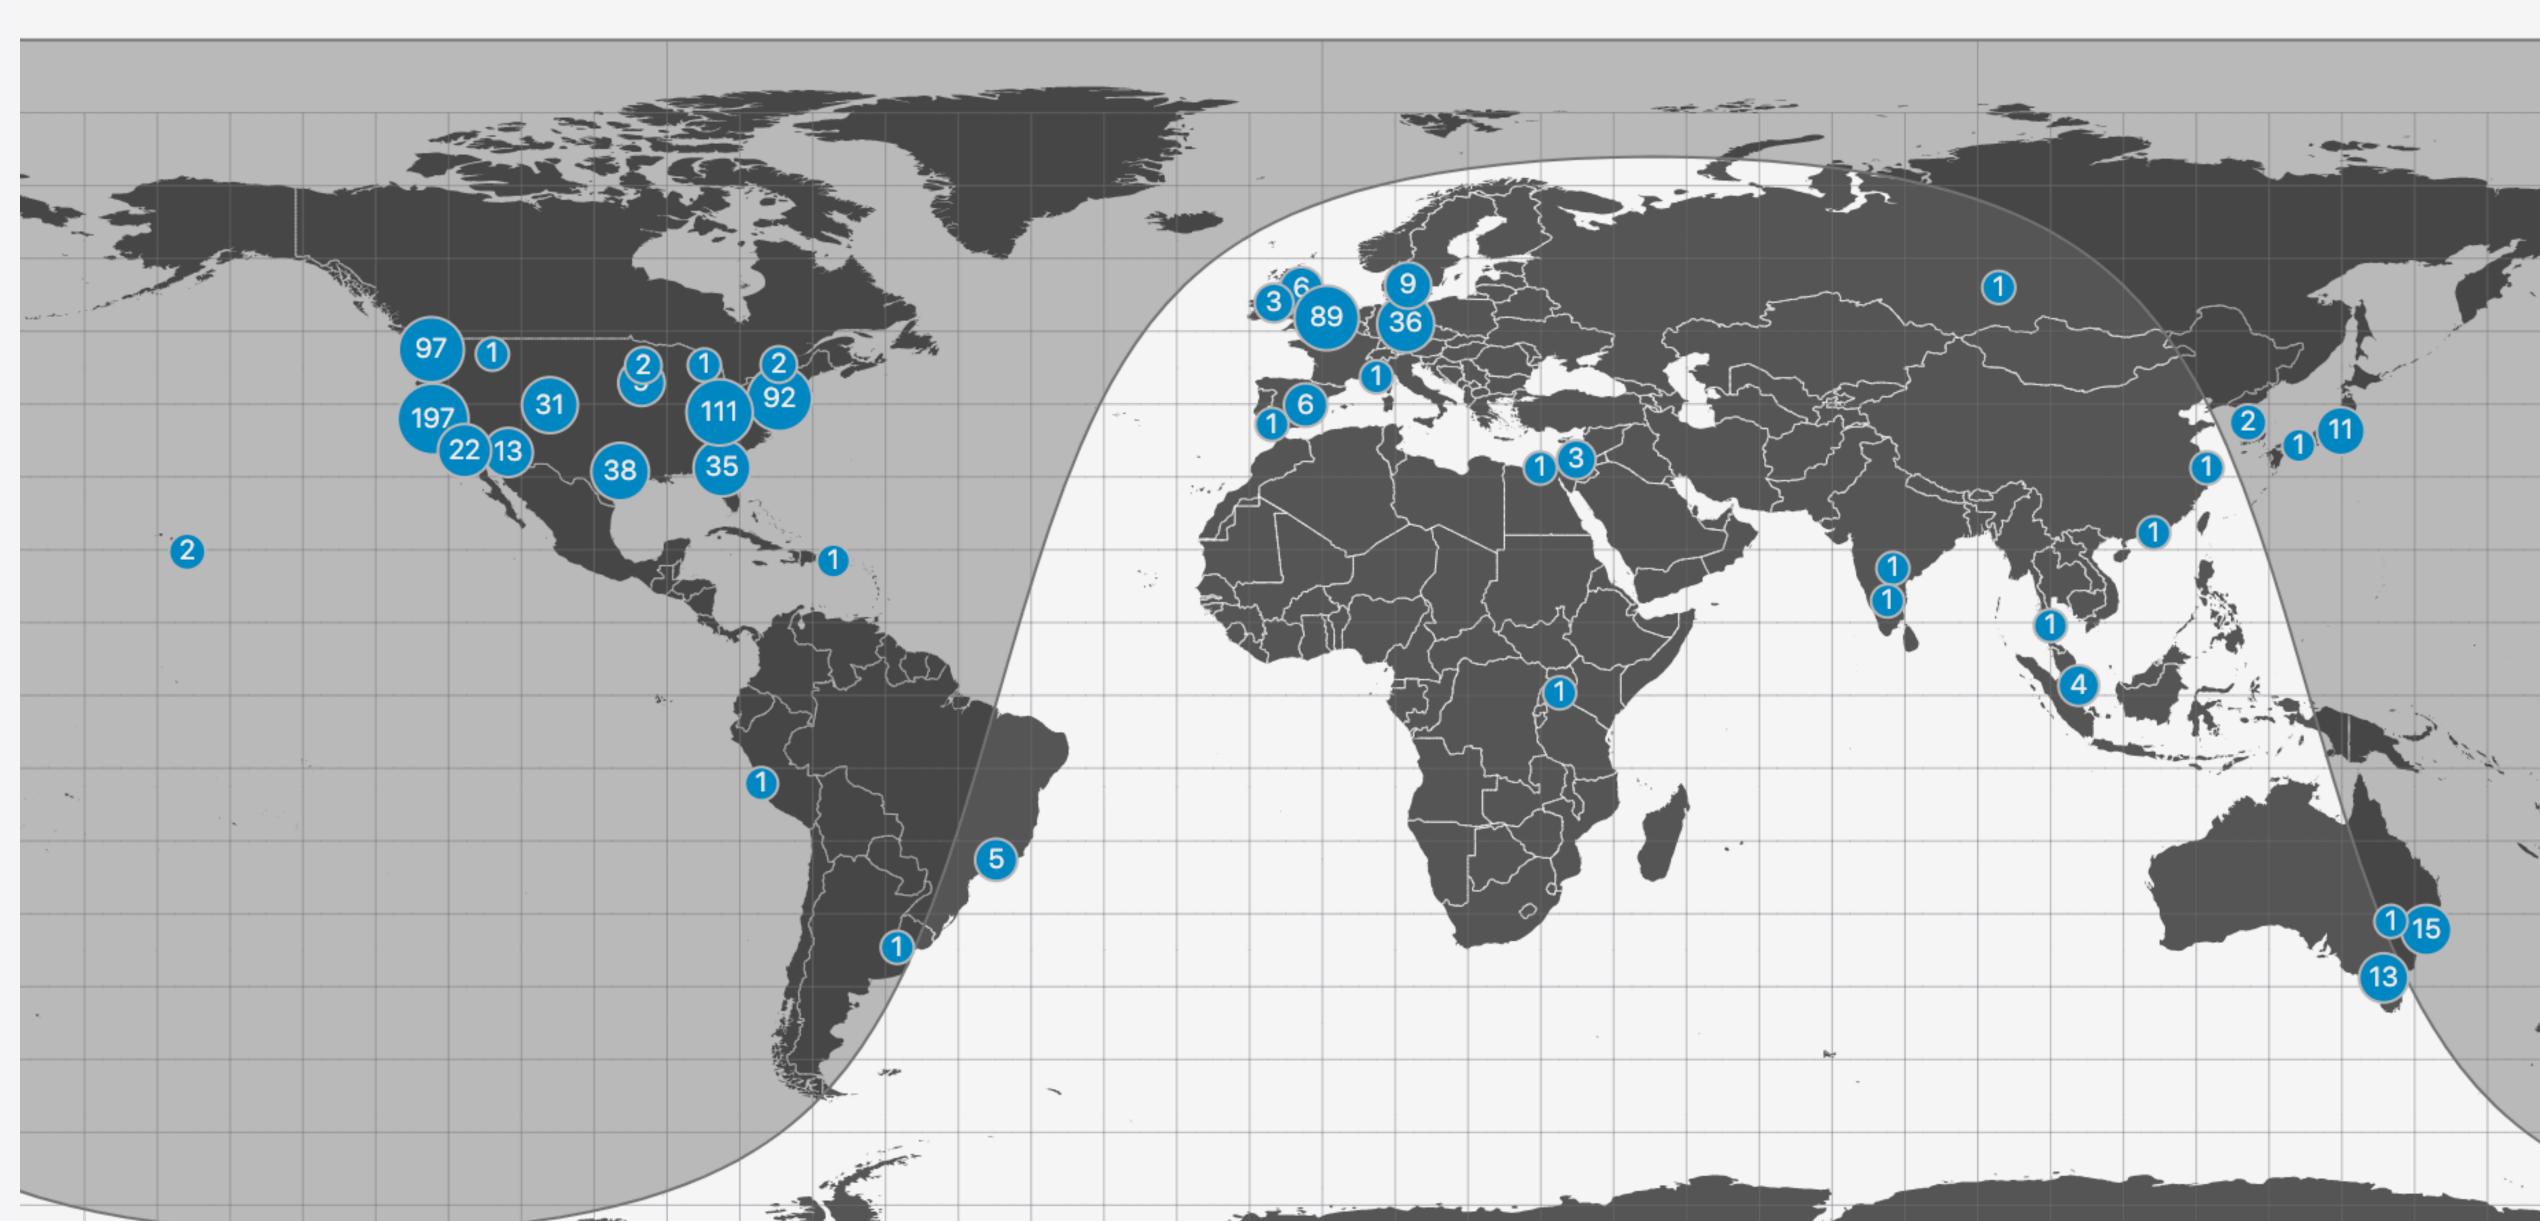






1613 3 1229/76% 383/23% 195/12%



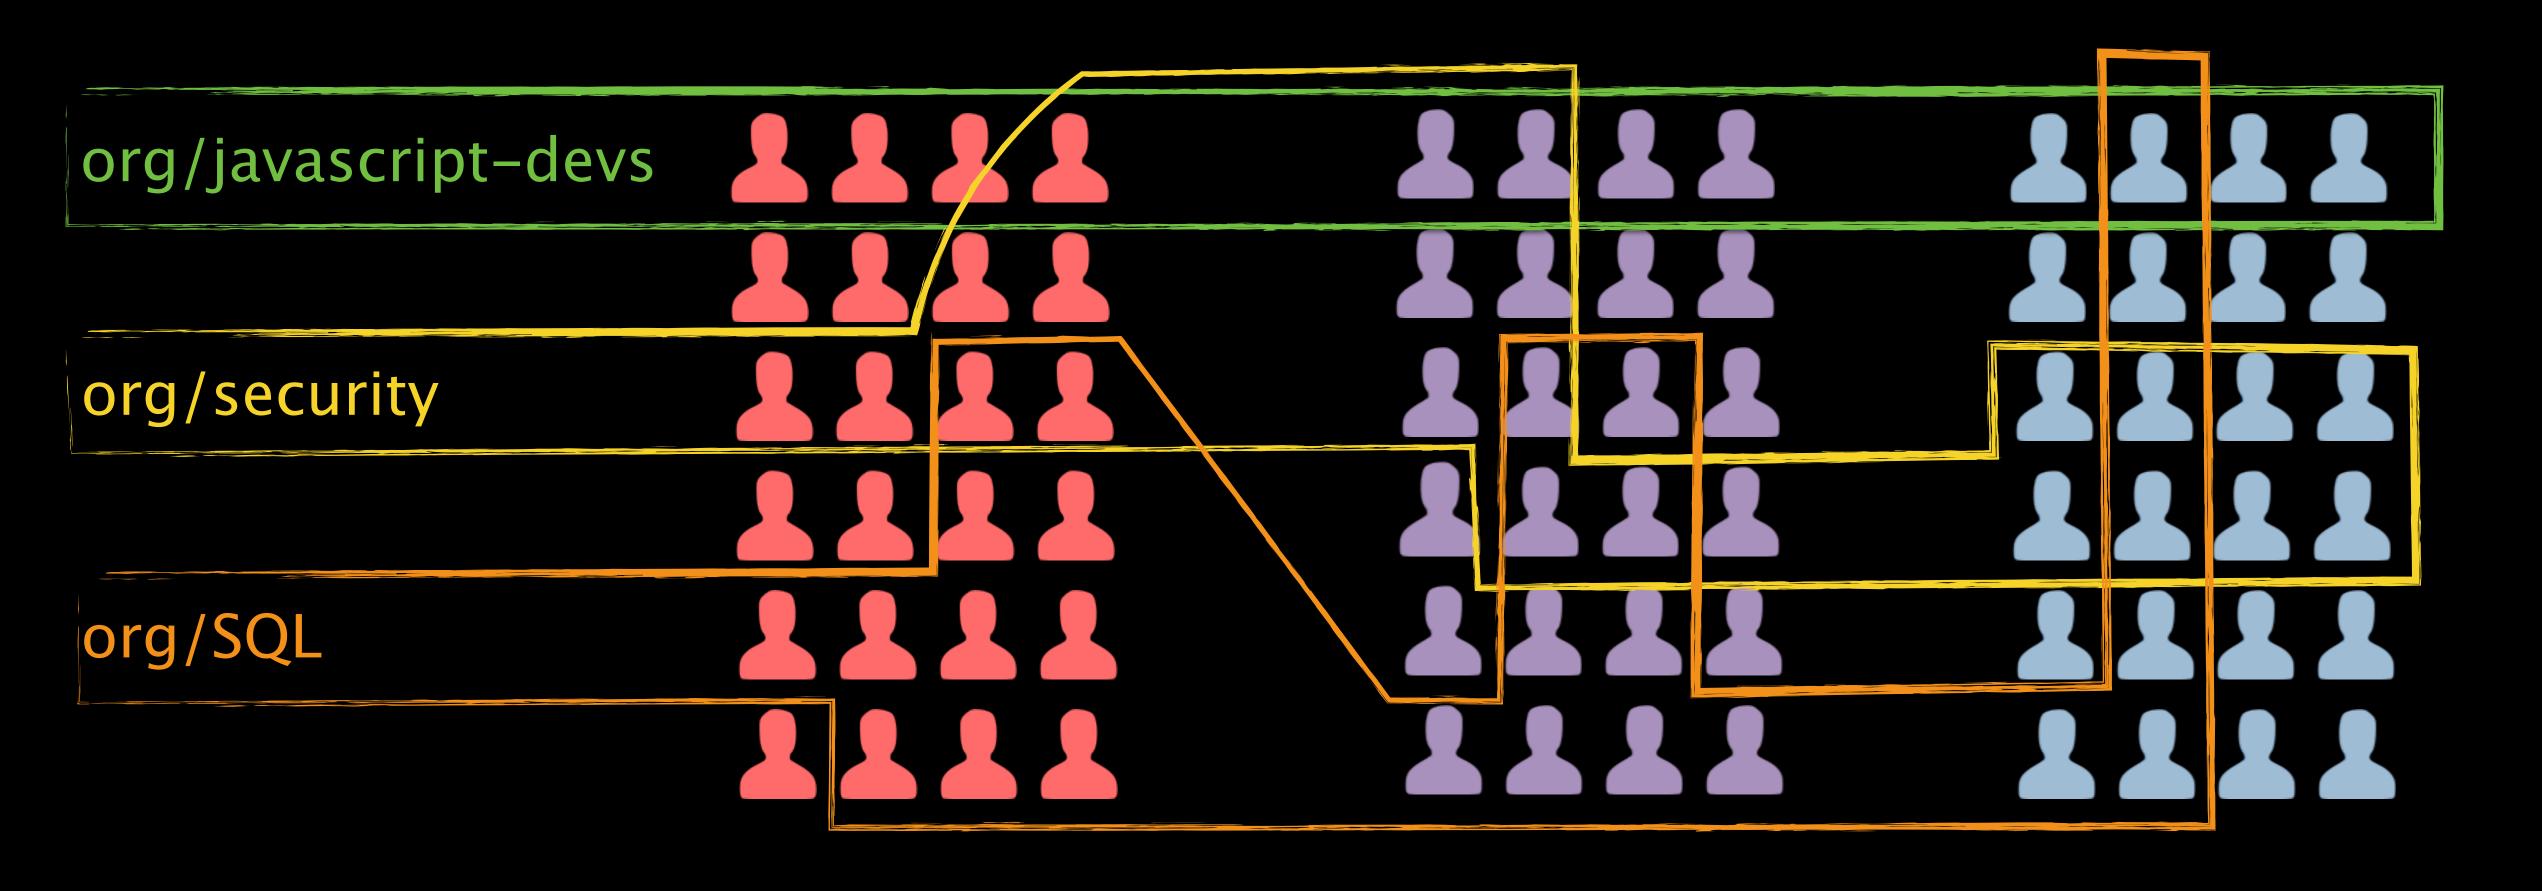




## Avoid use of emails

If there is no link it does not exist





org/devGroupA

org/devGroupB

org/devGroupC



## More teams than people

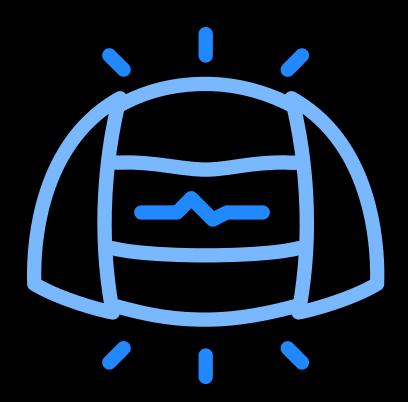
Use team @-mentions to communicate across departments



## What stayed (and evolved)



Remote First No Emails

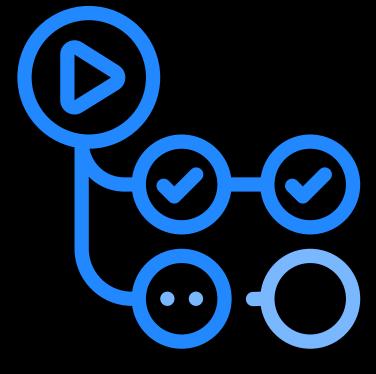


ChatOps Hubot

## What changed heavily



Architecture Languages



Flaky Test Detection



Deploy Trains

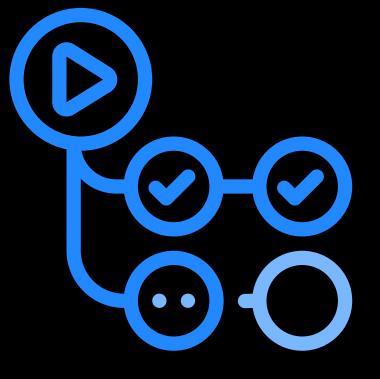




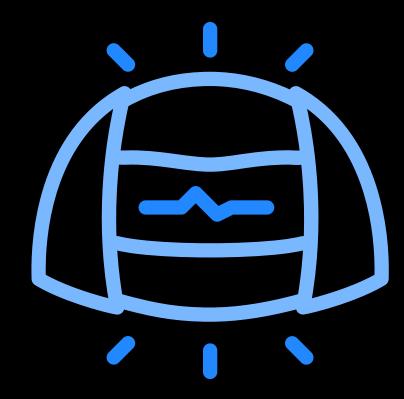
## Focus For Today



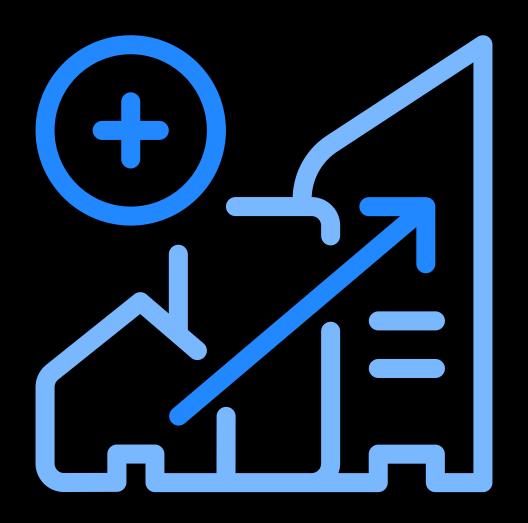
Architecture Languages



Flaky Test Detection



ChatOps Hubot



## Architecture / Languages







### From monolith to micro services

From one Ruby/Rails monolith to many micro services running on Kubernetes, managed by Moda - GitHub's internal service platform







**North America 2017** 

## Keynote: Kubernetes at GitHub

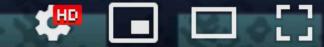
Jesse Newland, Principal Site Reliability Engineer, GitHub

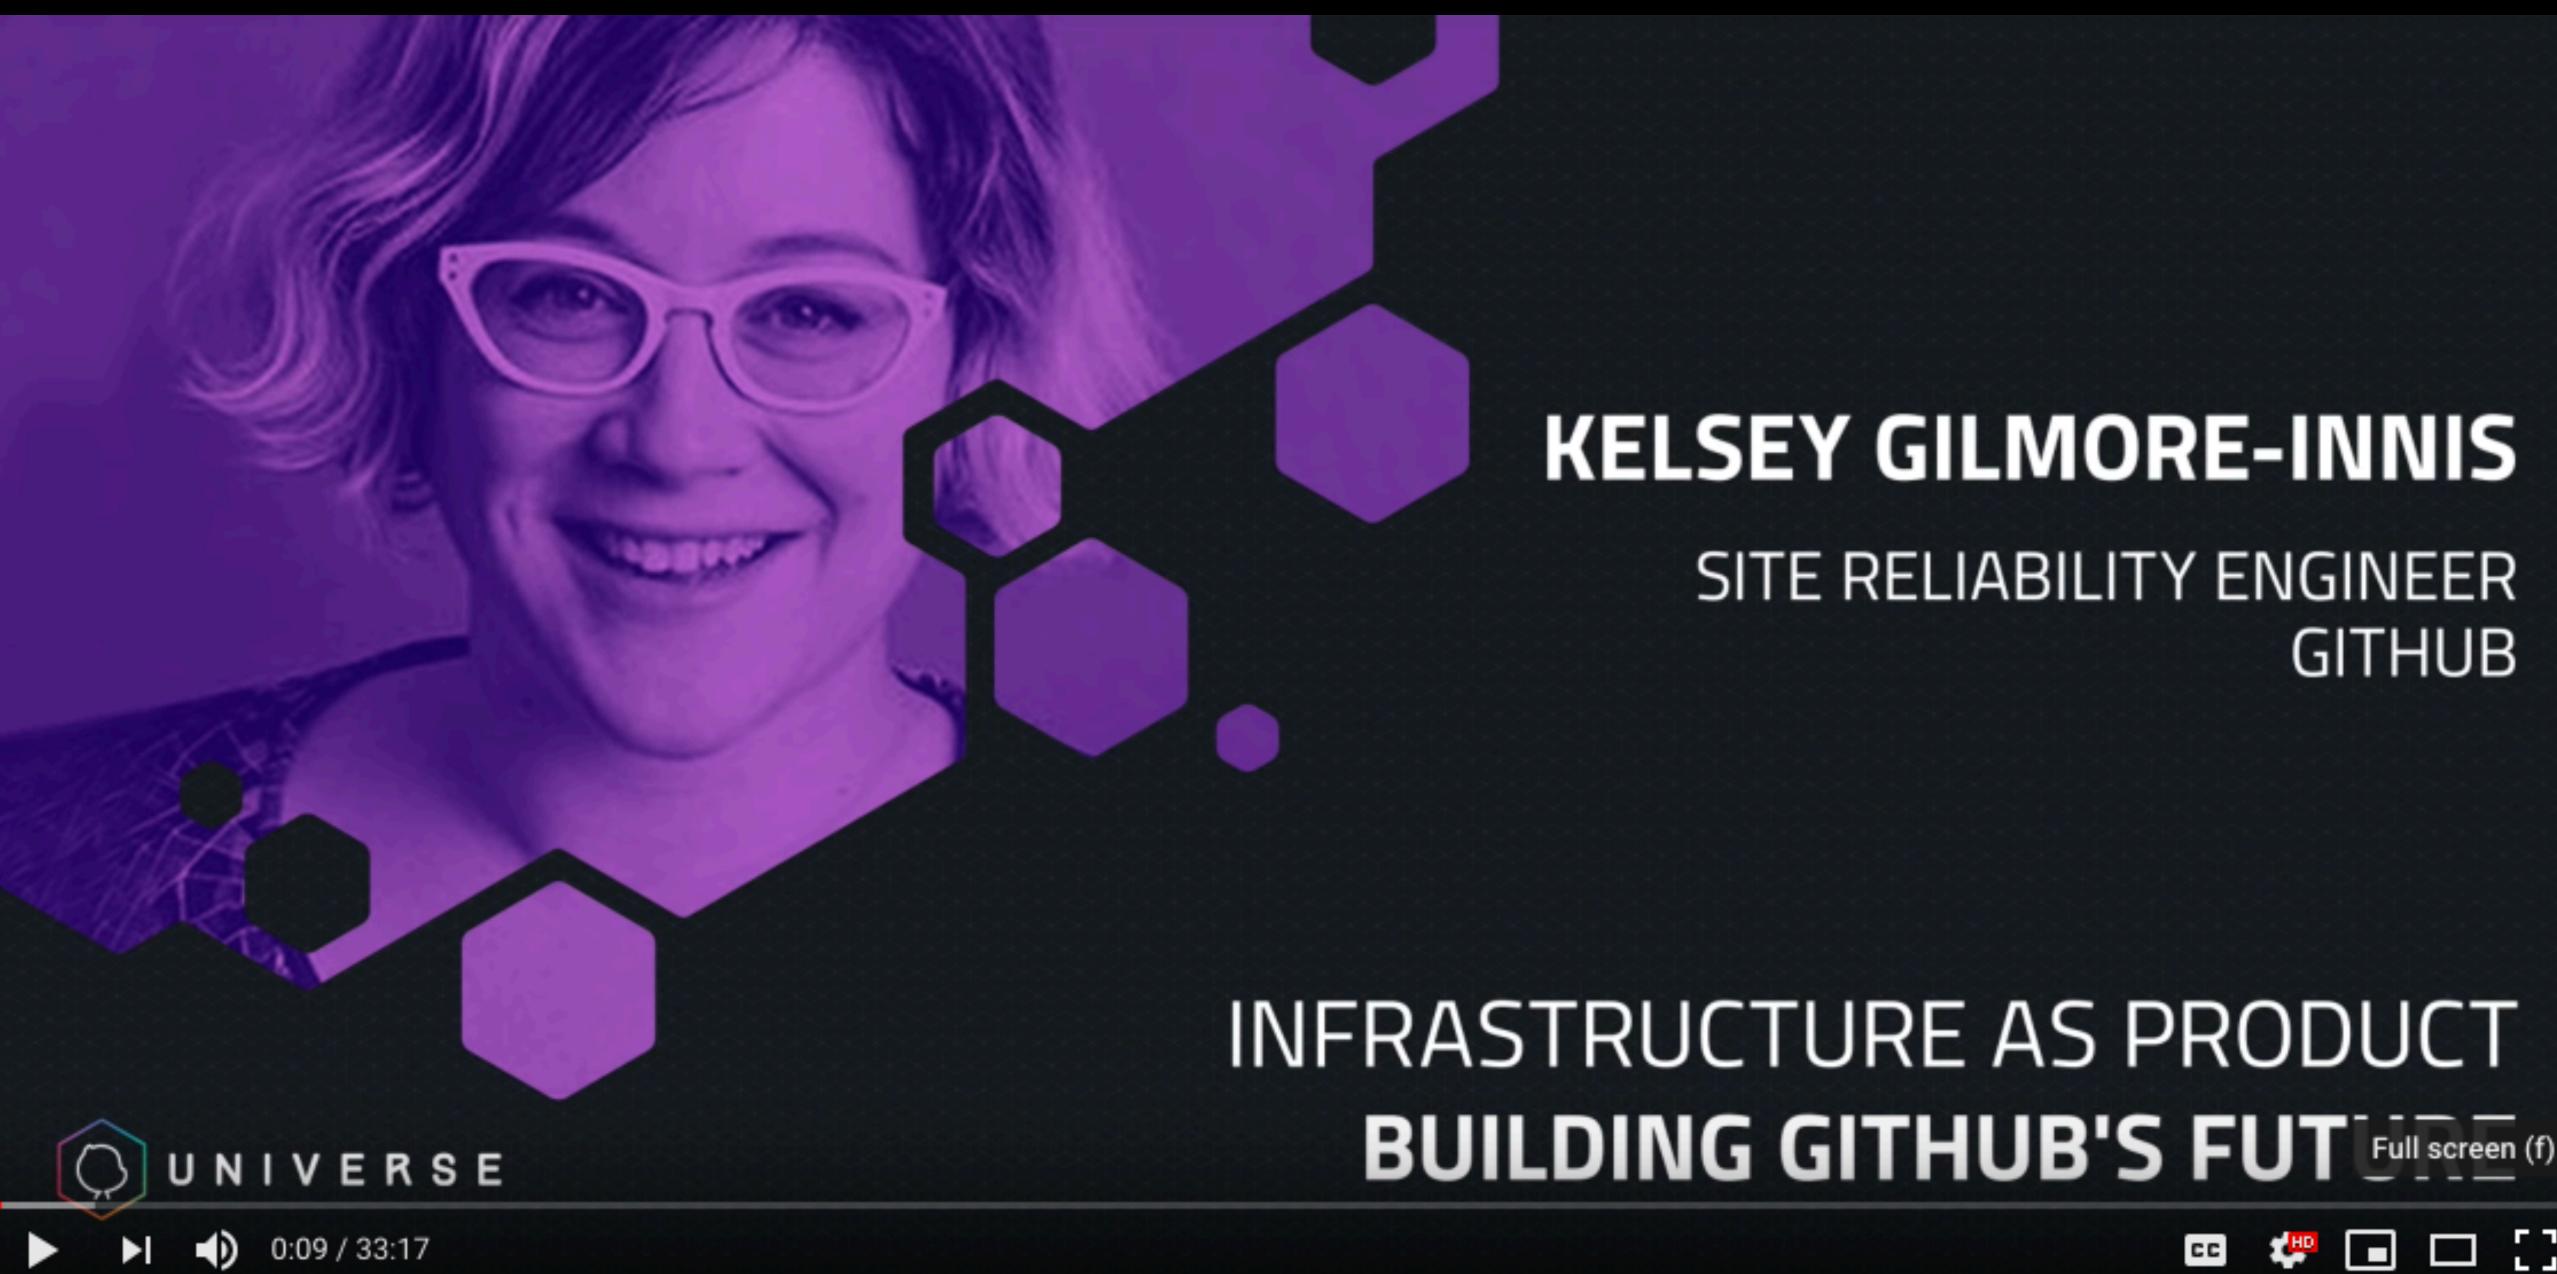






















## Paved roads for Ruby, Go, Java, C#

sponsoring team, loosely moderated environment, staffed and moderated environment, containerized releases available to production, test, and development clients, issue tracking for defects and feature requests, release and deprecation policy, roadmap with a plan for future improvements, style guide, supported by automated validation and formatting tooling, set of guidelines / recommendations for when using this language is appropriate, recommended and supported client libraries and workflows, recommendations for IDEs and plugins, automated and documented development workflow approved by relevant stakeholders, automated and documented CI / CD approach approved by relevant stakeholders, vulnerability detection and patch management workflow approved by relevant stakeholders, sample application designed to serve as a consolidated example of best practices, ...



#### Unicorns???



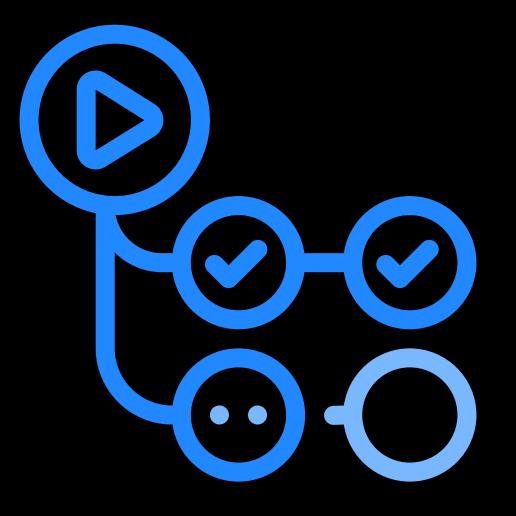


#### This page is taking way too long to load.

Sorry about that. Please try refreshing and contact us if the problem persists.

Contact Support GitHub Status — @githubstatus





## Flaky Test Detection



#### **Confession time**

Please raise your hand if you had to deal with flaky tests, i.e. tests that sometimes fail because of

- side effects of test execution order
- Reliance on external services, e.g. a remote HTTP API or time servers
- Current phase of the lunar cycle, cosmic rays, bad luck/karma, etc.









#### Flaky test math class

#### Let's assume a single flaky test fails in 1 out of thousand runs

How big is the like likelihood that an entire test suite run with **n** tests of that sort creates a least one intermittent test failure

 $P_{fail}(n)=1-0.999$ 





$$P_{fail}(n)=1-0.999^n$$

1 test

 $P_{fail}(1)=0.001$ 

100 tests

 $P_{fail}(100)=0.0952$ 

10 tests

 $P_{fail}(10)=0.009955$ 

1000 tests

P<sub>fail</sub>(100)=0.6323





# 30 percent of builds affected

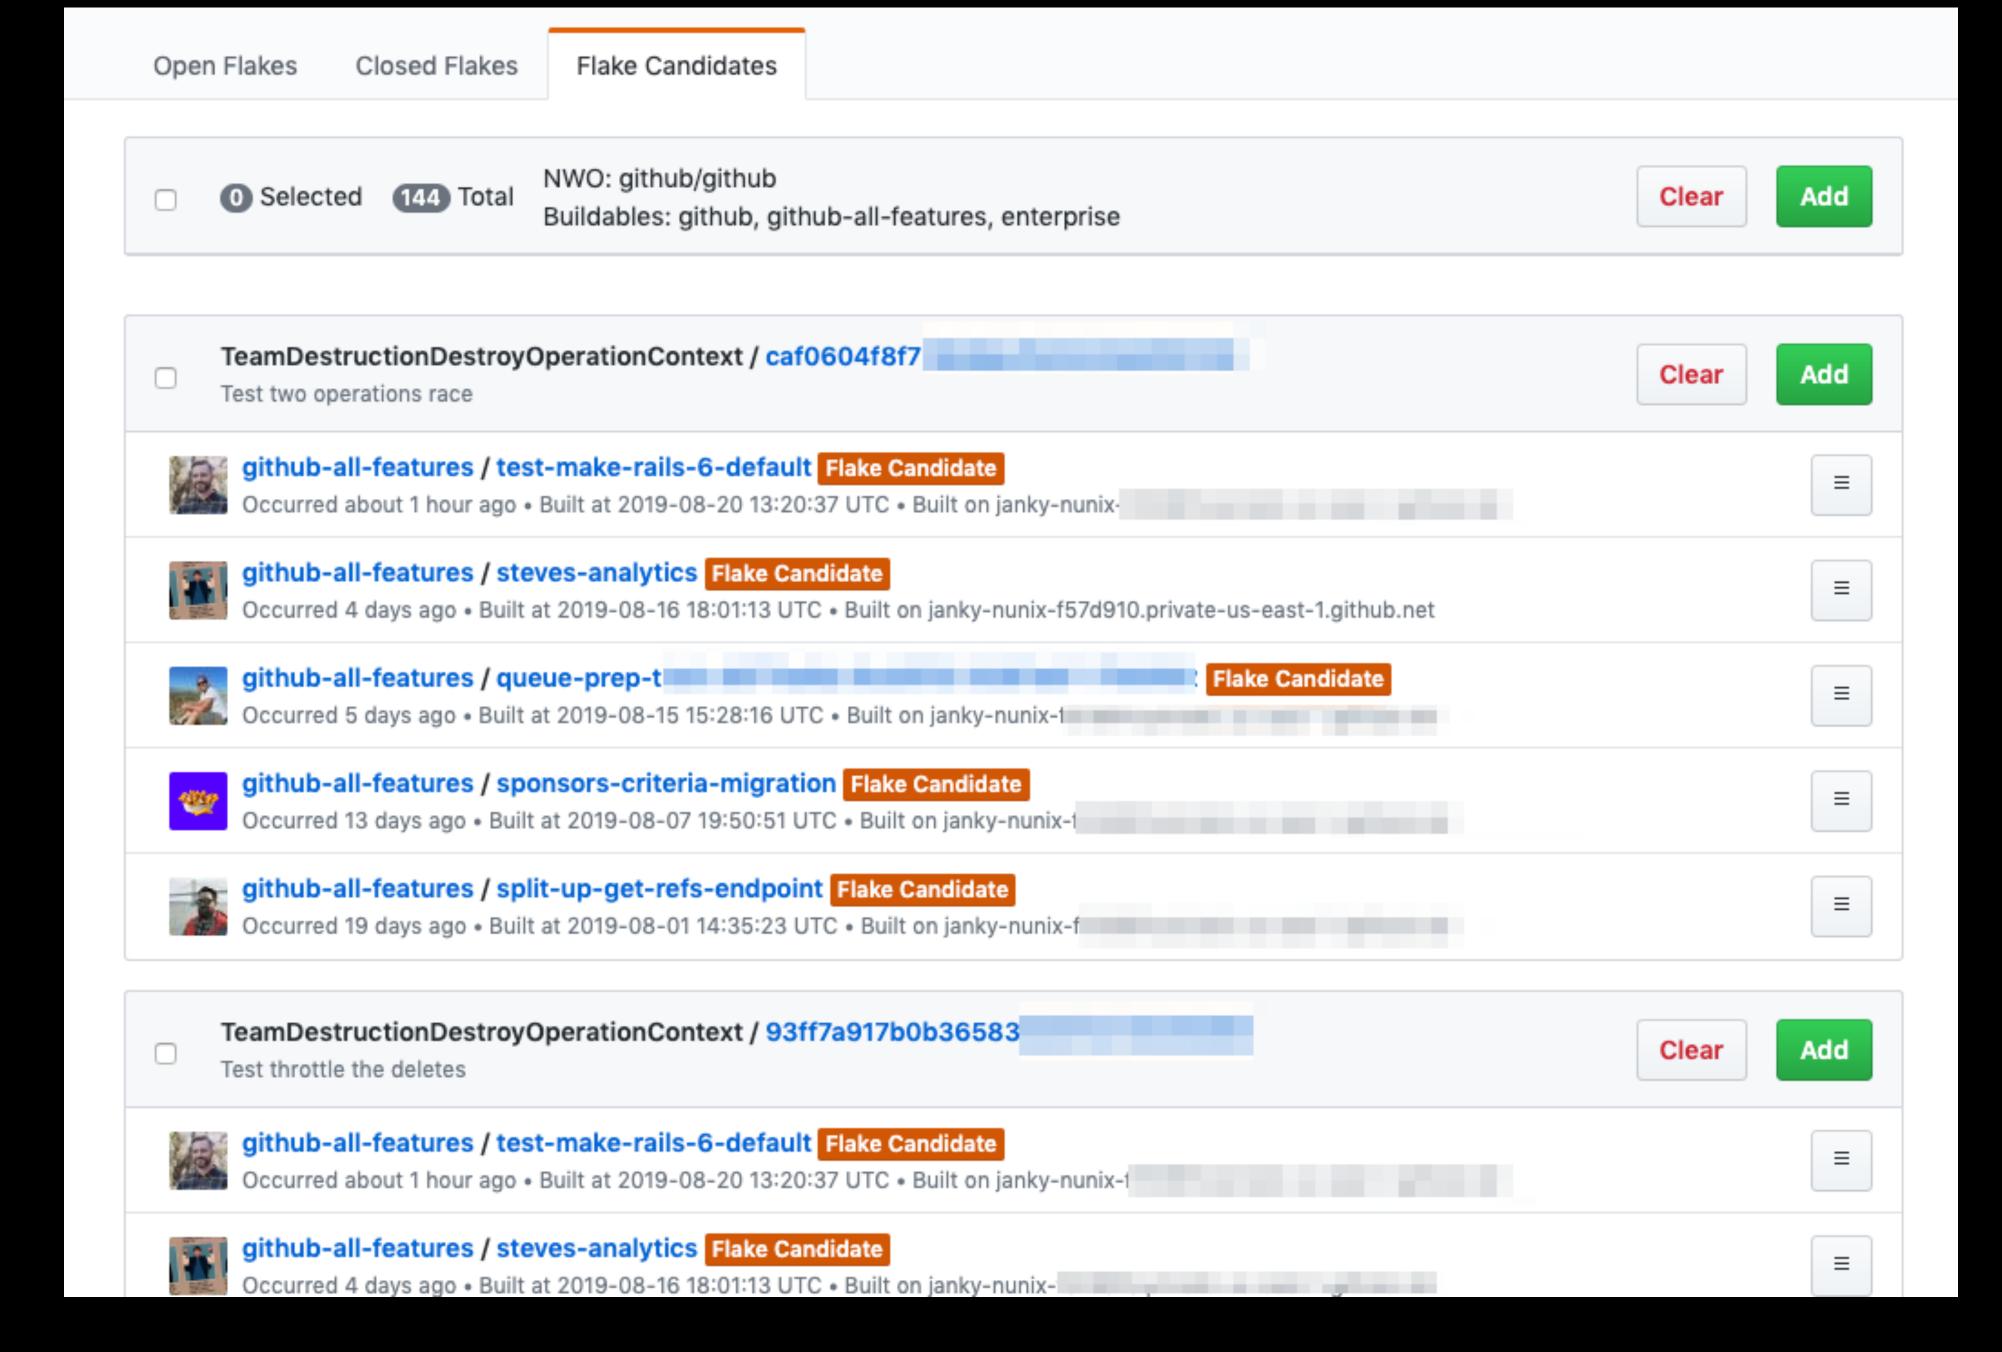
On 2018-12-06, 167 of 545 times our github-all-features test suite would have failed if it were not for marked flakes



## What's a flake candidate

- A flake candidate is a test that
  - both passes and fails against the same code (git tree) where less than 50 tests were failing
  - have been experienced by more than 1 user, on more than 1 branch
  - will be automatically captured and logged by the CI system by re-running failed tests
- Each flake candidate, has a SHA as identifier. The SHA is determined by:
  - Test name
  - Test suite
  - The first significant line of the exception stack trace





## What's a flake

- Flake candidates that are manually inspected by a developer and identified as intermittent test failures can be turned into real flakes (via ChatOps and Janky UI).
- CI builds will still succeed if < 30 flakes failed (and all other tests pass)</li>
- Flakes are identified as issues in github/github with the label of failing-test, multiple flake candidates can belong to the same issue
- All open flake have assigned owners and a deadline (1 week) to fix, else the test should rather be deleted
- Closing a flake is as simple as closing the issue. If an intermittent test failure reappears within 5 days, it will reopen the flake issue.



## Flaky test detection allows 30 percent more deployments



#### Flaky tests detection

Build #15378054 (481472e) of github-build-deploy-tarball-bp/master was successful (437s, queued 3s)

Build #15378045 (481472e) of enterprise/master failed with 8 known flakes - 9 failures (433s, queued 4s)

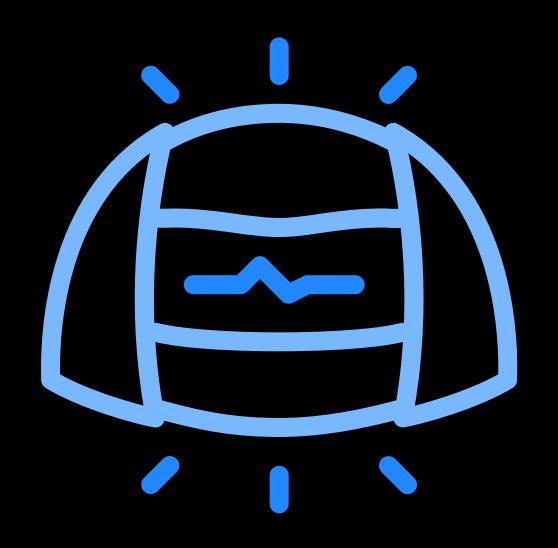
Build #15378071 (481472e) of github-rails-next/master was successful with 12 known flakes (440s, queued 3s)



latentflip 8:59 AM

.ci flakes add github

186b00f5f68c436c94e0dc4912bc3e80 https://github.com/github/github/issues/109058



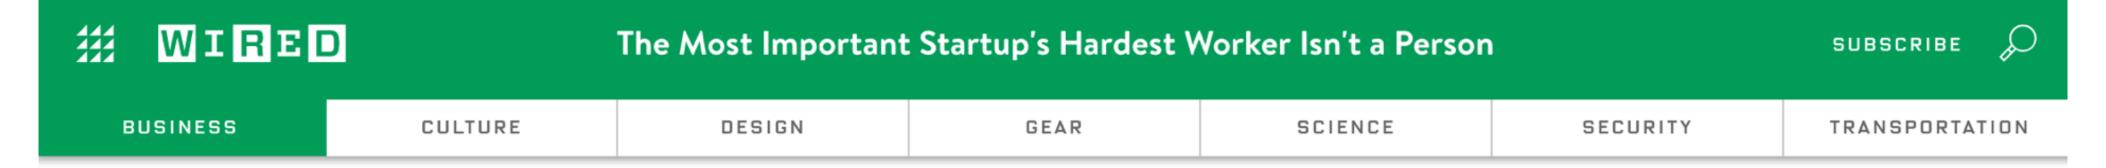
# ChatOps

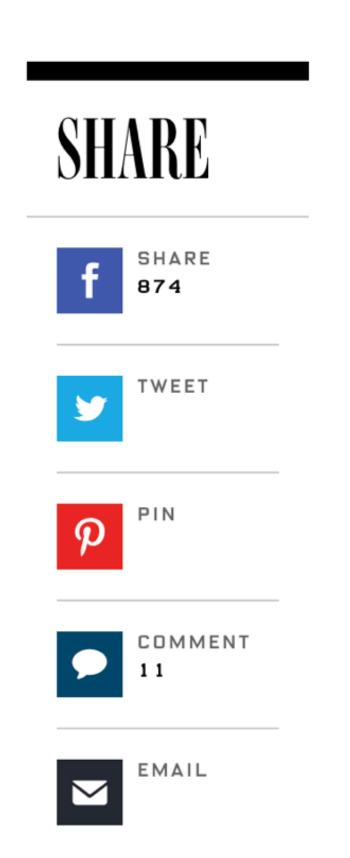






#### Hubot





CADE METZ BUSINESS 10.23.15 7:00 AM

'It's a new way of working.'

-SAM LAMBERT, GITHUB

Sam Lambert, the director of systems at GitHub, calls Hubot "the hardest working GitHubber." That's a company-wide in-joke. Hubot isn't really a GitHubber. He's a bit of software that plugs into the GitHub chat system. About five years ago, a guy named Ryan Tomayko built Hubot as an easier way for the company's engineers to manage and

modify all the hardware and software underpinning GitHub.com. Simply by sending a message to Hubot—much as they'd send a message to anyone else from inside the GitHub chat client—engineers could update the operating





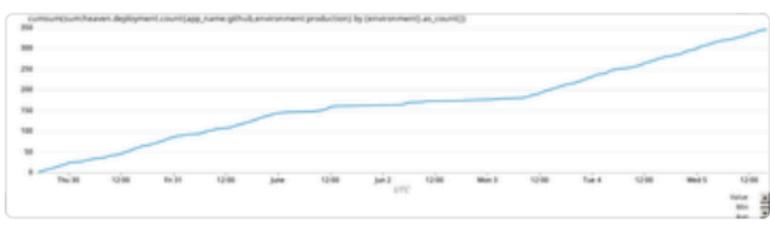
#### brntbeer 6:08 PM

.dgraph -1w cumsum(sum:heaven.deployment.count{app\_name:githu



hubot APP 6:09 PM

cumsum(sum:heaven.deployment.count{app\_name:github,environby {environment}.as\_count()) (23 kB) •





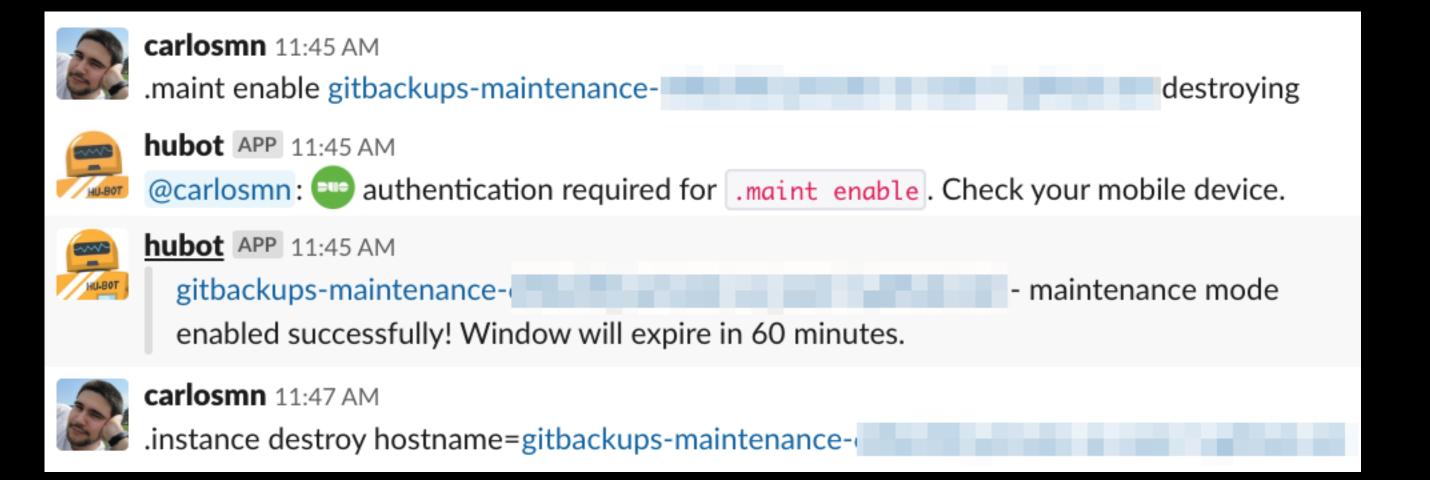
#### brntbeer 1:41 PM

.dash me -1h glb.services --glb-service dotcom



#### hubot APP 1:41 PM

glb.services (117 kB) ▼





#### **brntbeer** 1:51 PM

.catalog owner notifications



#### hubot APP 1:51 PM

Service: notifications

Maintainer: shayfrendt

Team: github/pe-notifications

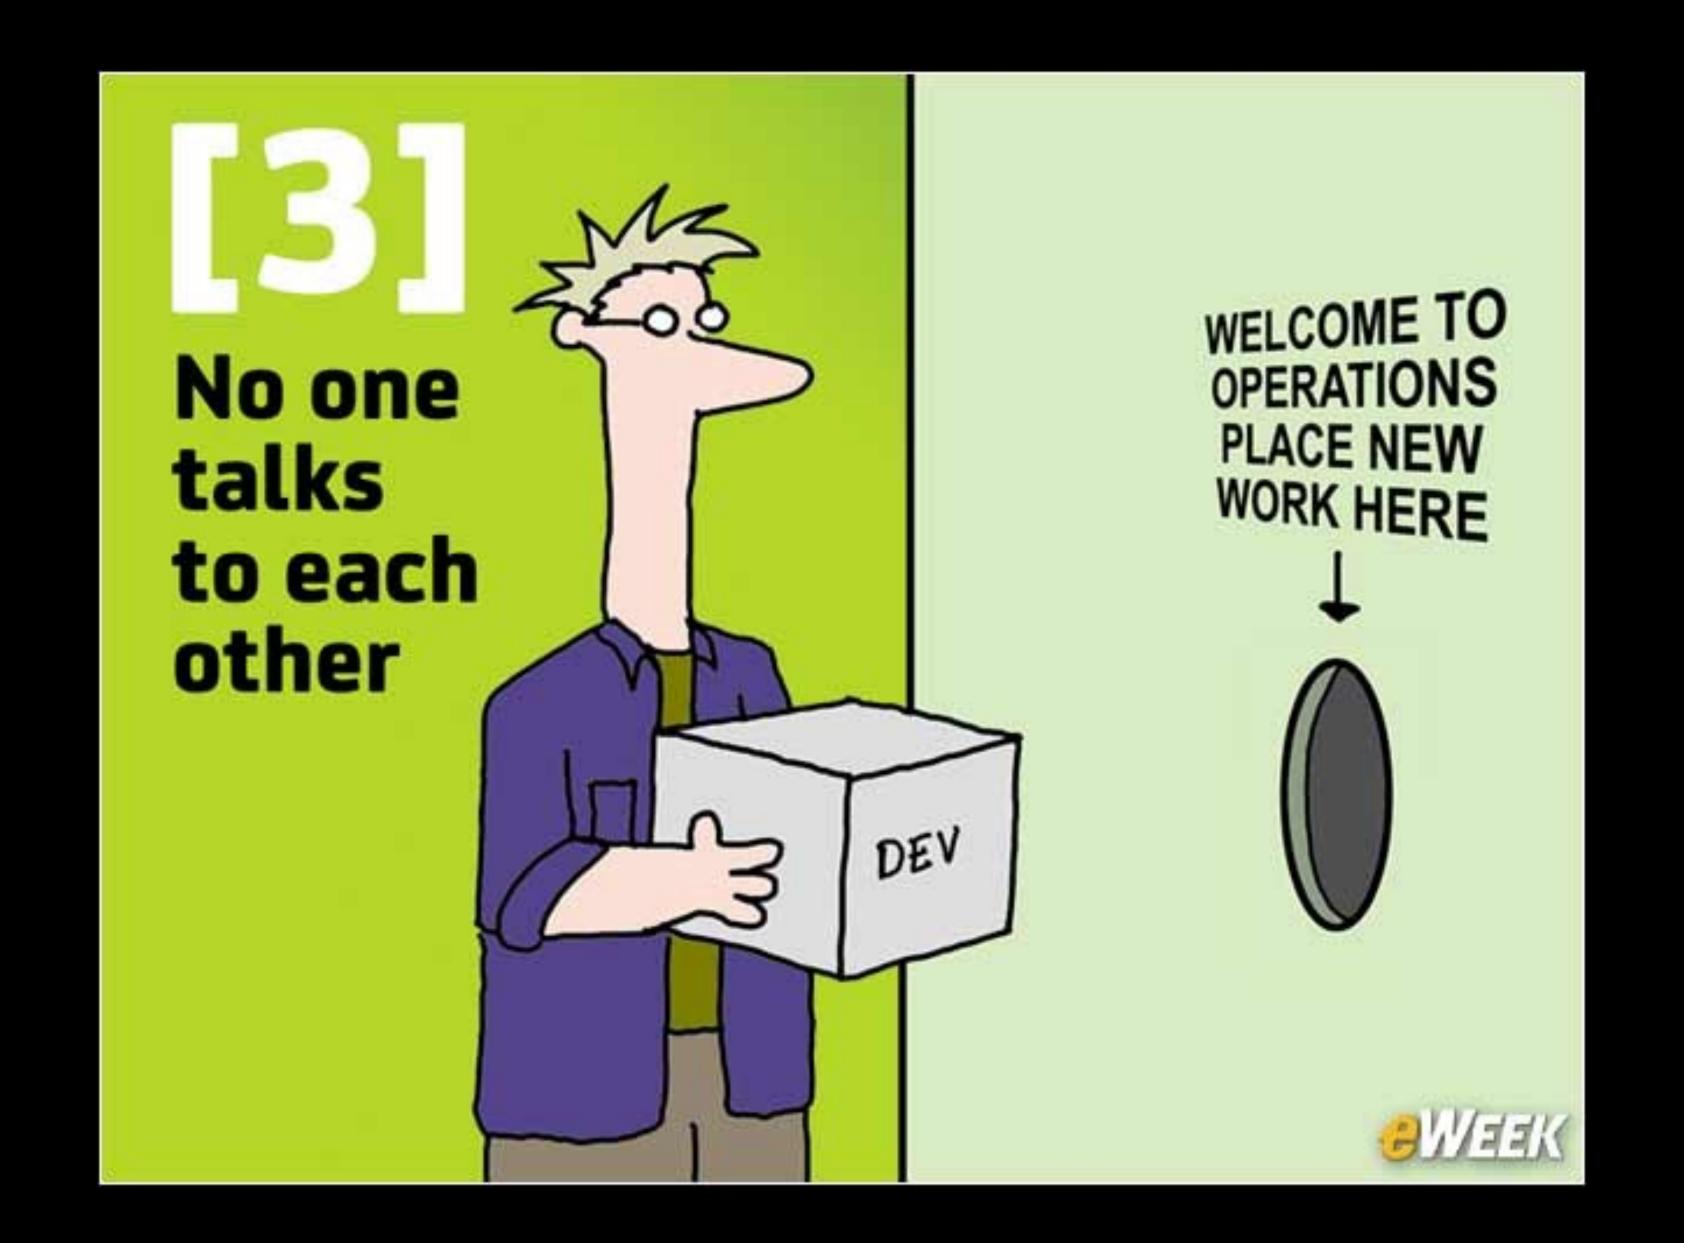
#### **Escalation Paths:**

sev1: \_\_\_\_ .pager trigger github-dotcom-oncall critical <reason>

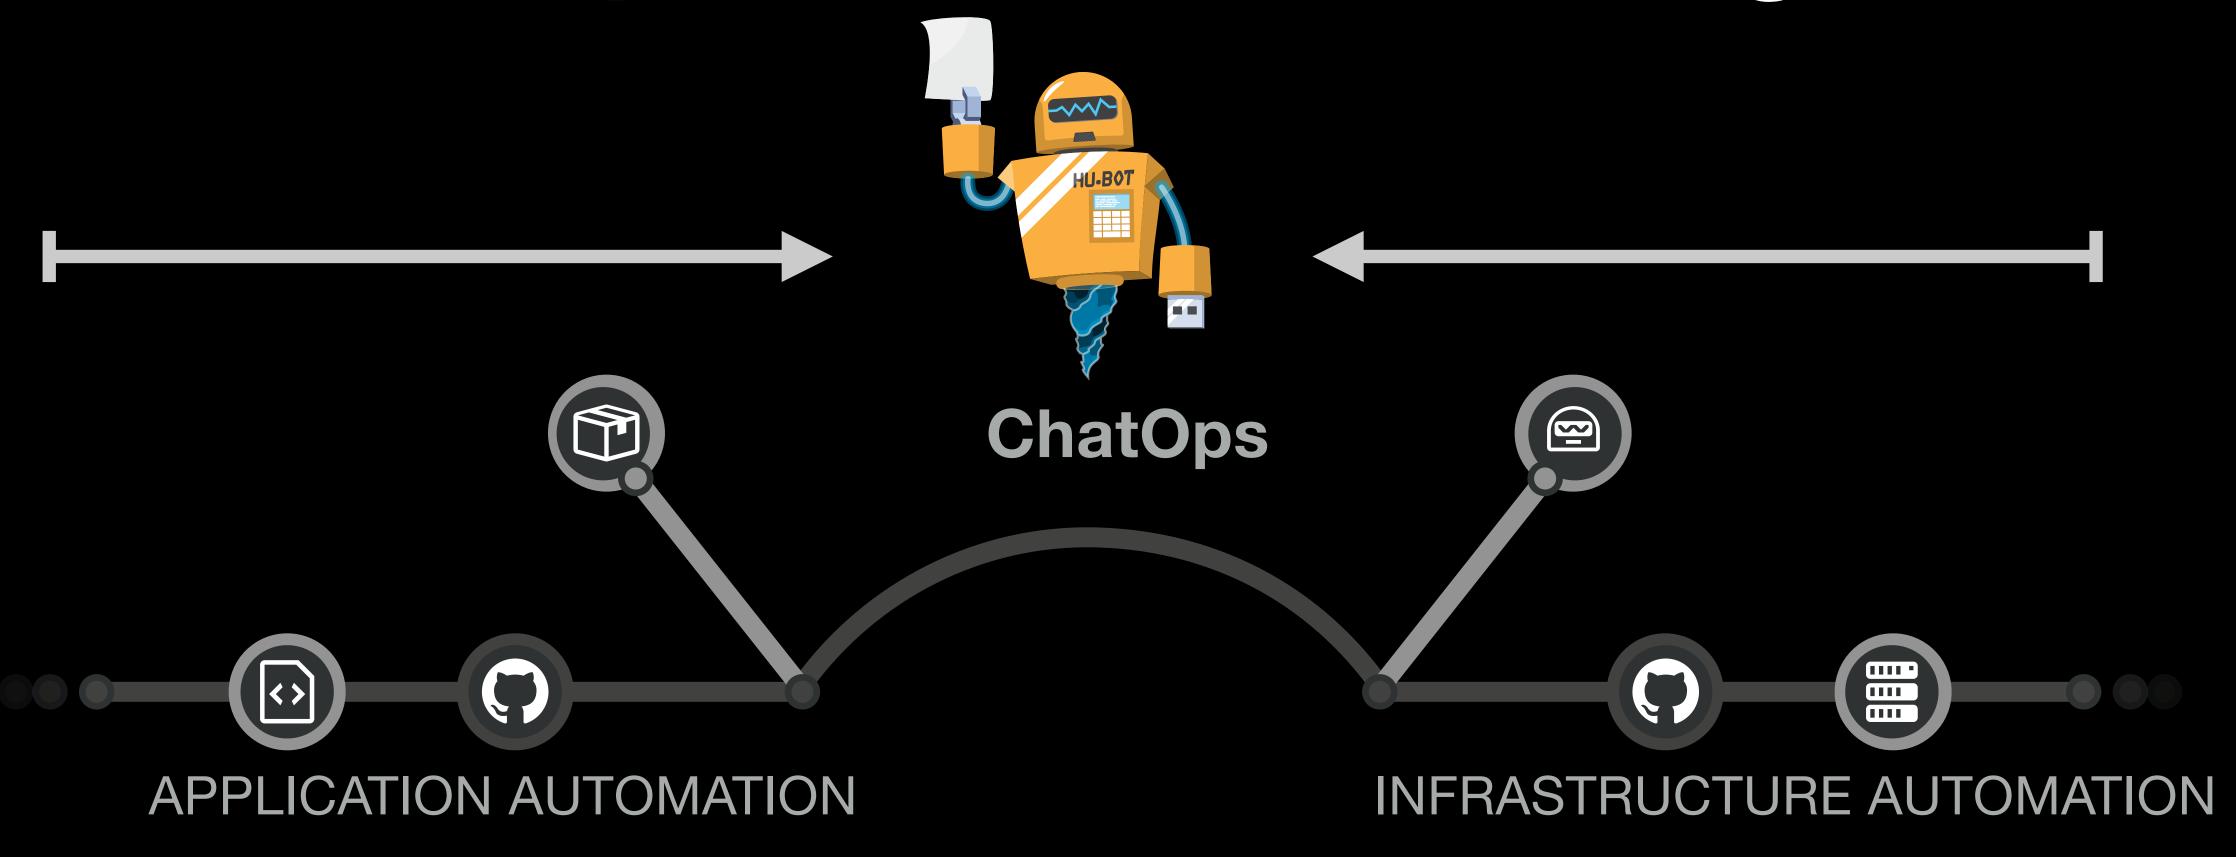
sev2: Thttps://github.com/github/github/issues

sev3: #pe-notifications

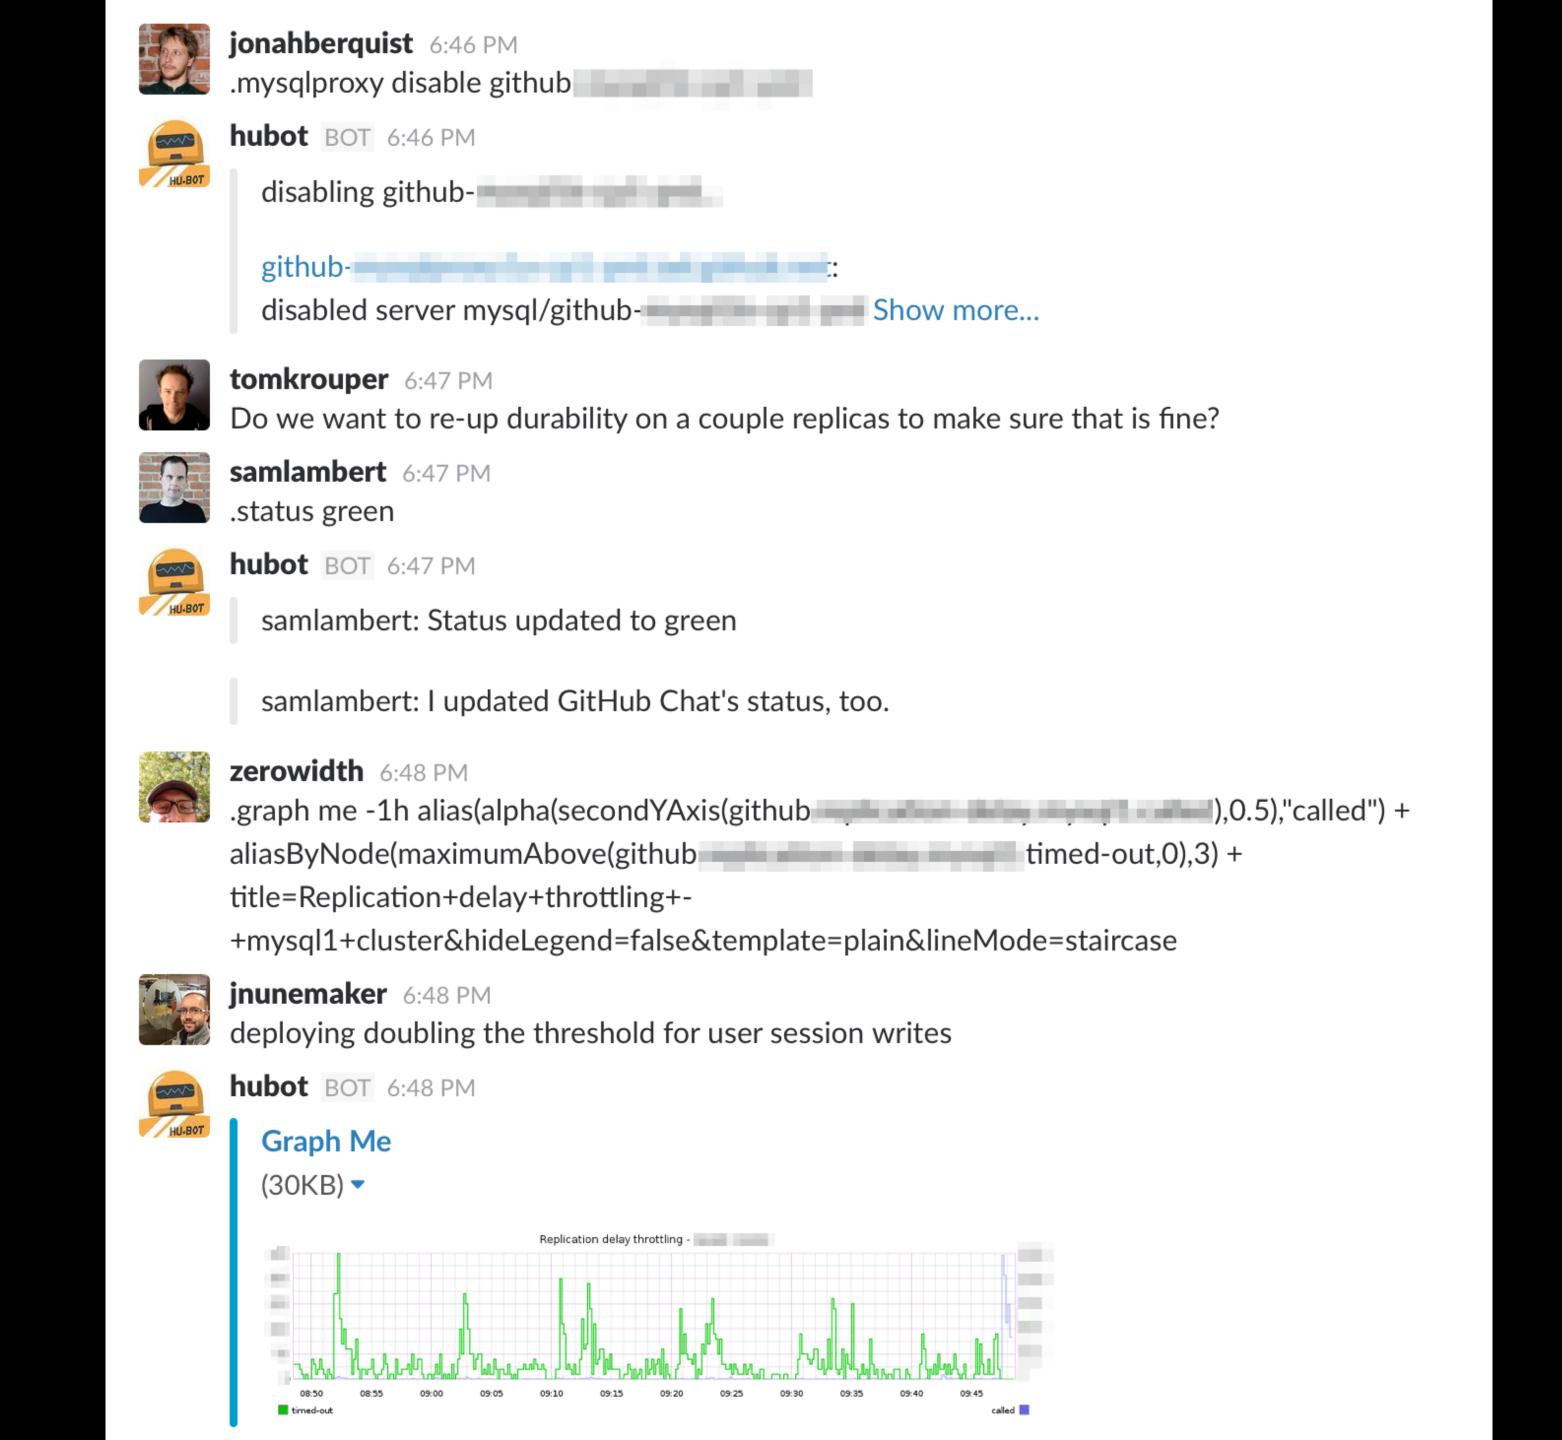




# ChatOps as the cultural glue

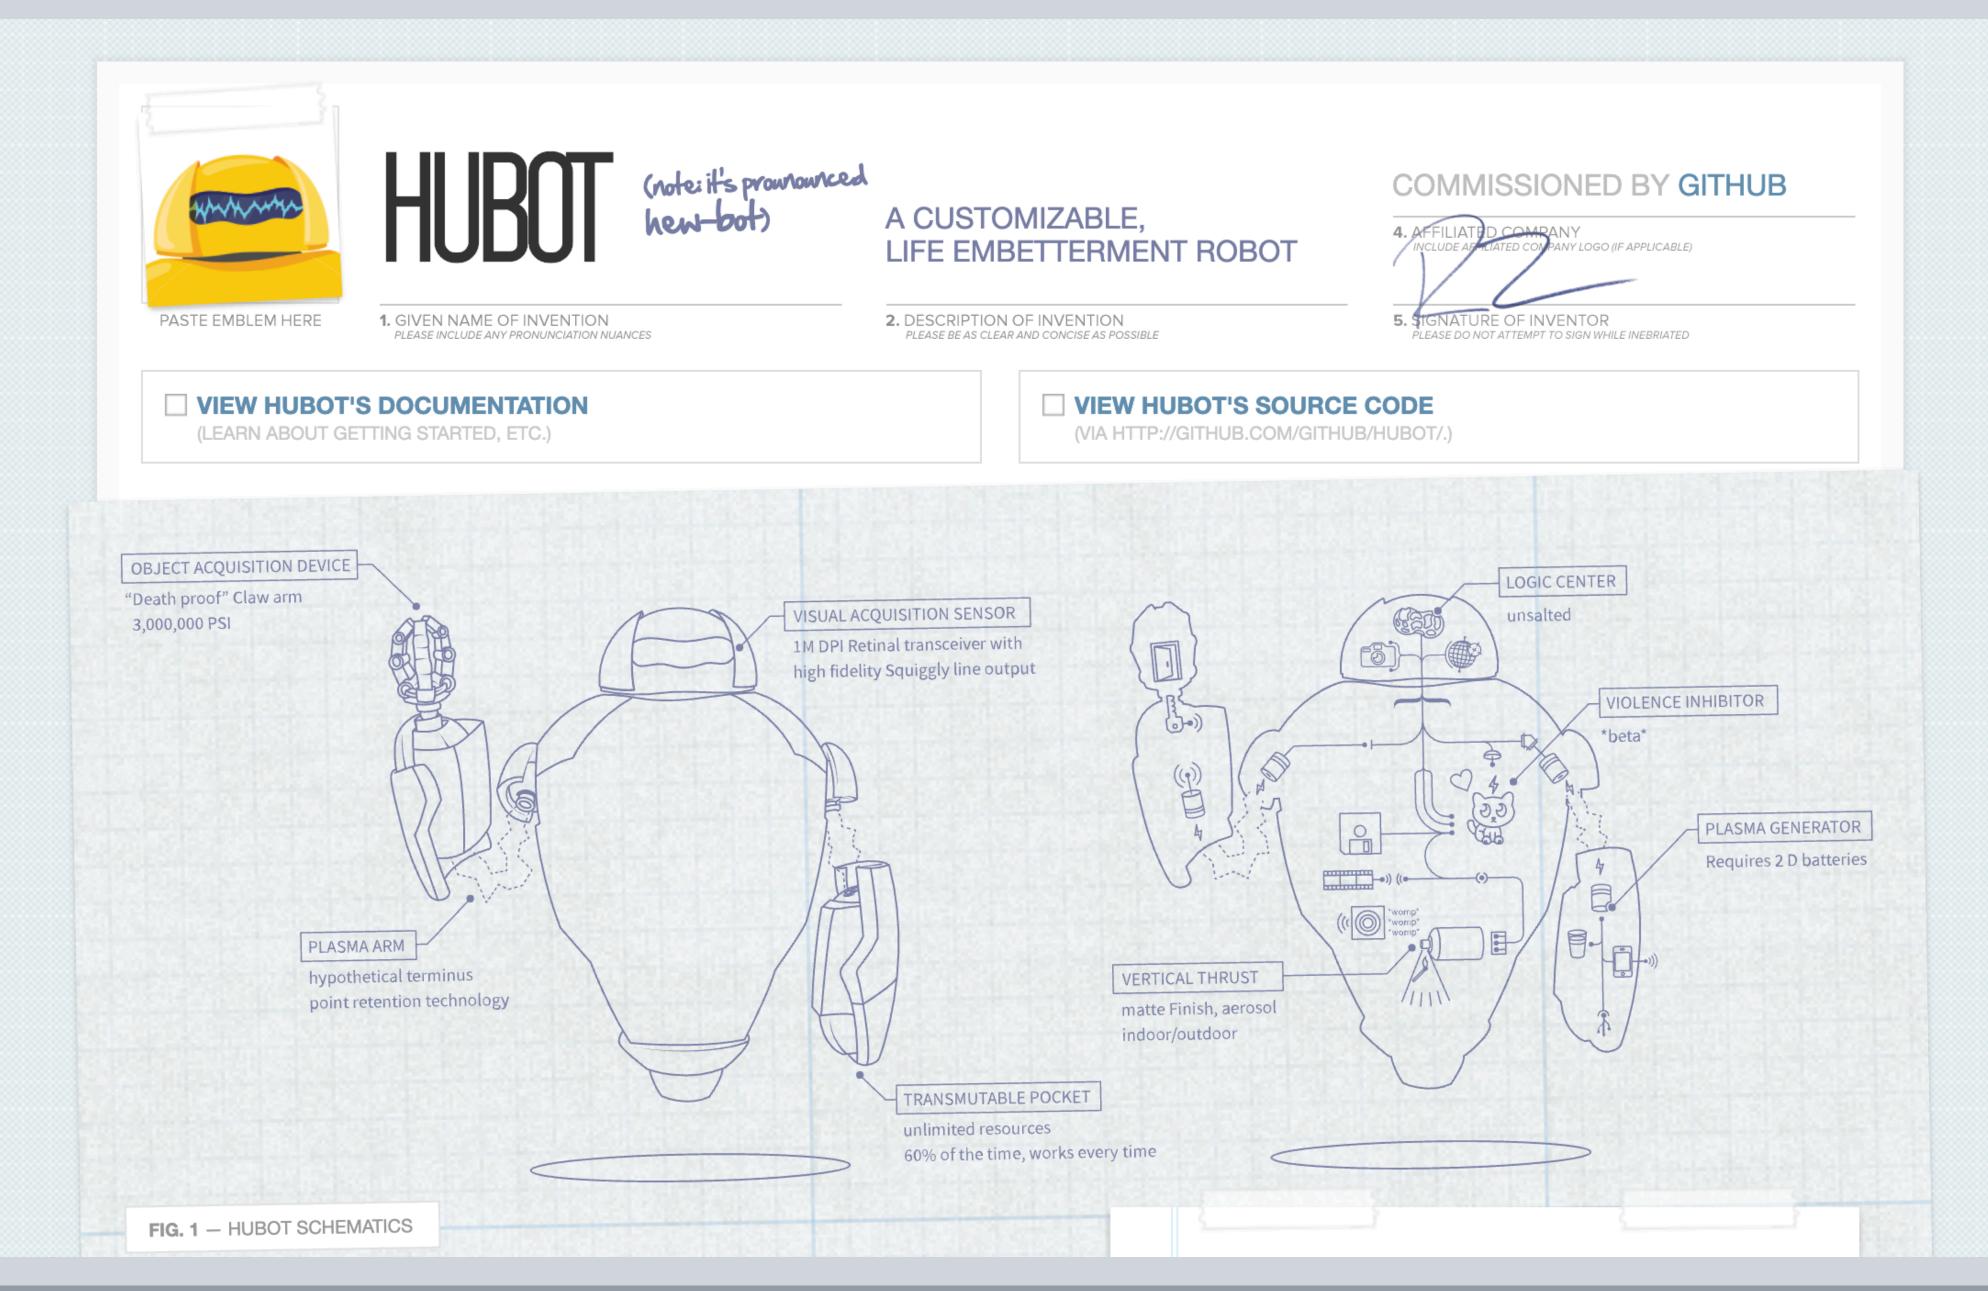






### https://hubot.github.com







# Summary

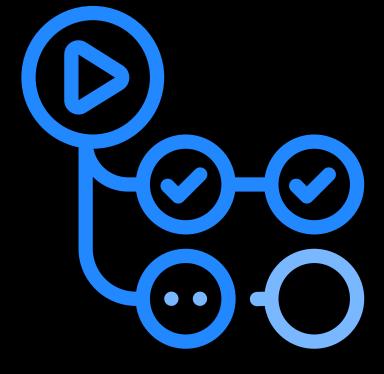




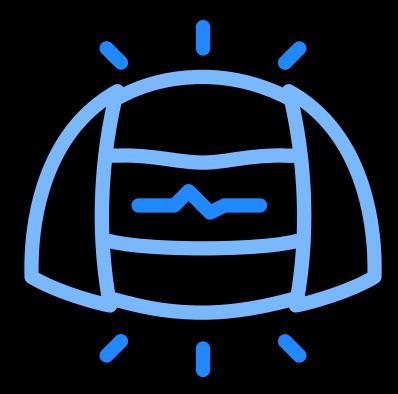
# Focus For Today



Architecture Languages



Flaky Test Detection



ChatOps Hubot

# Realität T

Pony hof

#### Was bei einem neuen Job wichtig ist

In Deutschland stehen bei der Beurteilung eines Arbeitsplatzes die spezifischen Technologien im Vordergrund, gefolgt von der Büroumgebung und der Unternehmenskultur. Die Technikbranche kämpft insgesamt mit Problemen rund um die Vielfalt und im Allgemeinen machen Entwickler sie nicht zu einer Priorität bei der Suche nach einem Job. Deutsche Entwickler, die einer geschlechtsspezifischen Minderheit angehören (Frauen, nicht binär), schätzen die Vielfalt in einer Organisation eher.



## Wie Entwickler potenzielle Arbeitsplätze bewerten: Höchste Priorität

Die Sprachen, Frameworks und andere Technologien, mit denen ich arbeiten würde

19,4 %

Das Arbeitsumfeld oder die Unternehmenskultur

17,4 %

Die Vergütung und die angebotenen Zusatzleistungen

15,1 %

Die Möglichkeiten der beruflichen Weiterentwicklung

14,2 %

Die spezifische Abteilung oder das Team, in der/in dem ich arbeiten würde

9,8 %

Die Möglichkeit, von zu Hause aus /remote zu arbeiten

8,2 %

Die Branche, in der ich arbeiten würde

Wie weit verbreitet oder einflussreich das Produkt oder die Dienstleistung, an der ich arbeiten würde, ist 5,2 %

Die wirtschaftliche Leistungsfähigkeit / der Finanzierungsstatus des Unternehmens oder der Organisation

# All Respondents Men → Women Non-binary Distracting work environment 41.8% Meetings 36.6% Being tasked with non-development work 36.5% Not enough people for the workload 33.6% Lack of support from management 26.6% Inadequate access to necessary tools 21.1% Toxic work environment 20.8% Non-work commitments (parenting, school work, hobbies, etc.)







## HOW THE FUTURE OF SOFTWARE HAS TO BE DIFFERENT

"Es ist nicht deine Schuld, dass der Code ist wie er ist, es wär' nur deine Schuld, wenn er so bleibt."





## Thank you