


 Hasso Plattner Institut Potsdam

Scaling DevOps - GitHub's journey from 500+ to 2000+ people

@jonico January 8th 2021



what's
in it for
ME?



GitHub is the largest developer platform on Earth

60M+

Developers

200M+

Private and public
repositories

1,000s

Top open source
communities

1B+

Contributions
per year

3M+

Organizations

50%

Fortune 500
companies

The most innovative companies

stripe

Spotify

airbnb

lyft

slack

Uber

twilio

shopify

NETFLIX

The top open source projects



docker

TensorFlow



APACHE

elasticsearch

npm

GraphQL



Leading global enterprises

SAP

DAIMLER



Continental



TESLA

QUALCOMM



Elektrobit

Largest DevOps ecosystem

RapidAPI

percy

GitKraken

SENTRY

ZenHub

Codecov

Jira

HOUND

HashiCorp



OTTO



zalando



DAIMLER




BOSCH



Delivery Hero

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 [Doom](#) and [Prince of Persia](#) are archived in public repositories—certainly a better fate than [being buried in the New Mexican desert](#).



22 SEP

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 [edX](#) and [Udacity](#), you've created thousands of ways to learn software development on GitHub. Courses at that top [our list](#) include [Ada's Jumpstart program](#) and [Stanford's TensorFlow Tutorials](#). STAT545 is teaching thousands of students to wrangle data, while [CS50](#) 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 [Campus Experts](#) are also building tech communities on campuses in more than 15 countries.



GitHub Now vs Then

Repositories

200 M+

vs. 36 M in 2016

Registered users

60 M+

vs. 14 M in 2016

Pull Requests

500 M+

vs. 130 M in 2017

Suggested fixes for security alerts

10 M+

vs. 800 K in 2018



A dark gray background featuring a complex network diagram. The diagram consists of numerous white circles of varying sizes, some of which are connected by thin white lines. Interspersed among the circles are several white speech bubble icons, each containing three dots. The overall pattern is dense and interconnected, suggesting a global or community network.






99%

Of software projects use open source

Code

```
1 import kiwisolver
2 import numpy
3
```

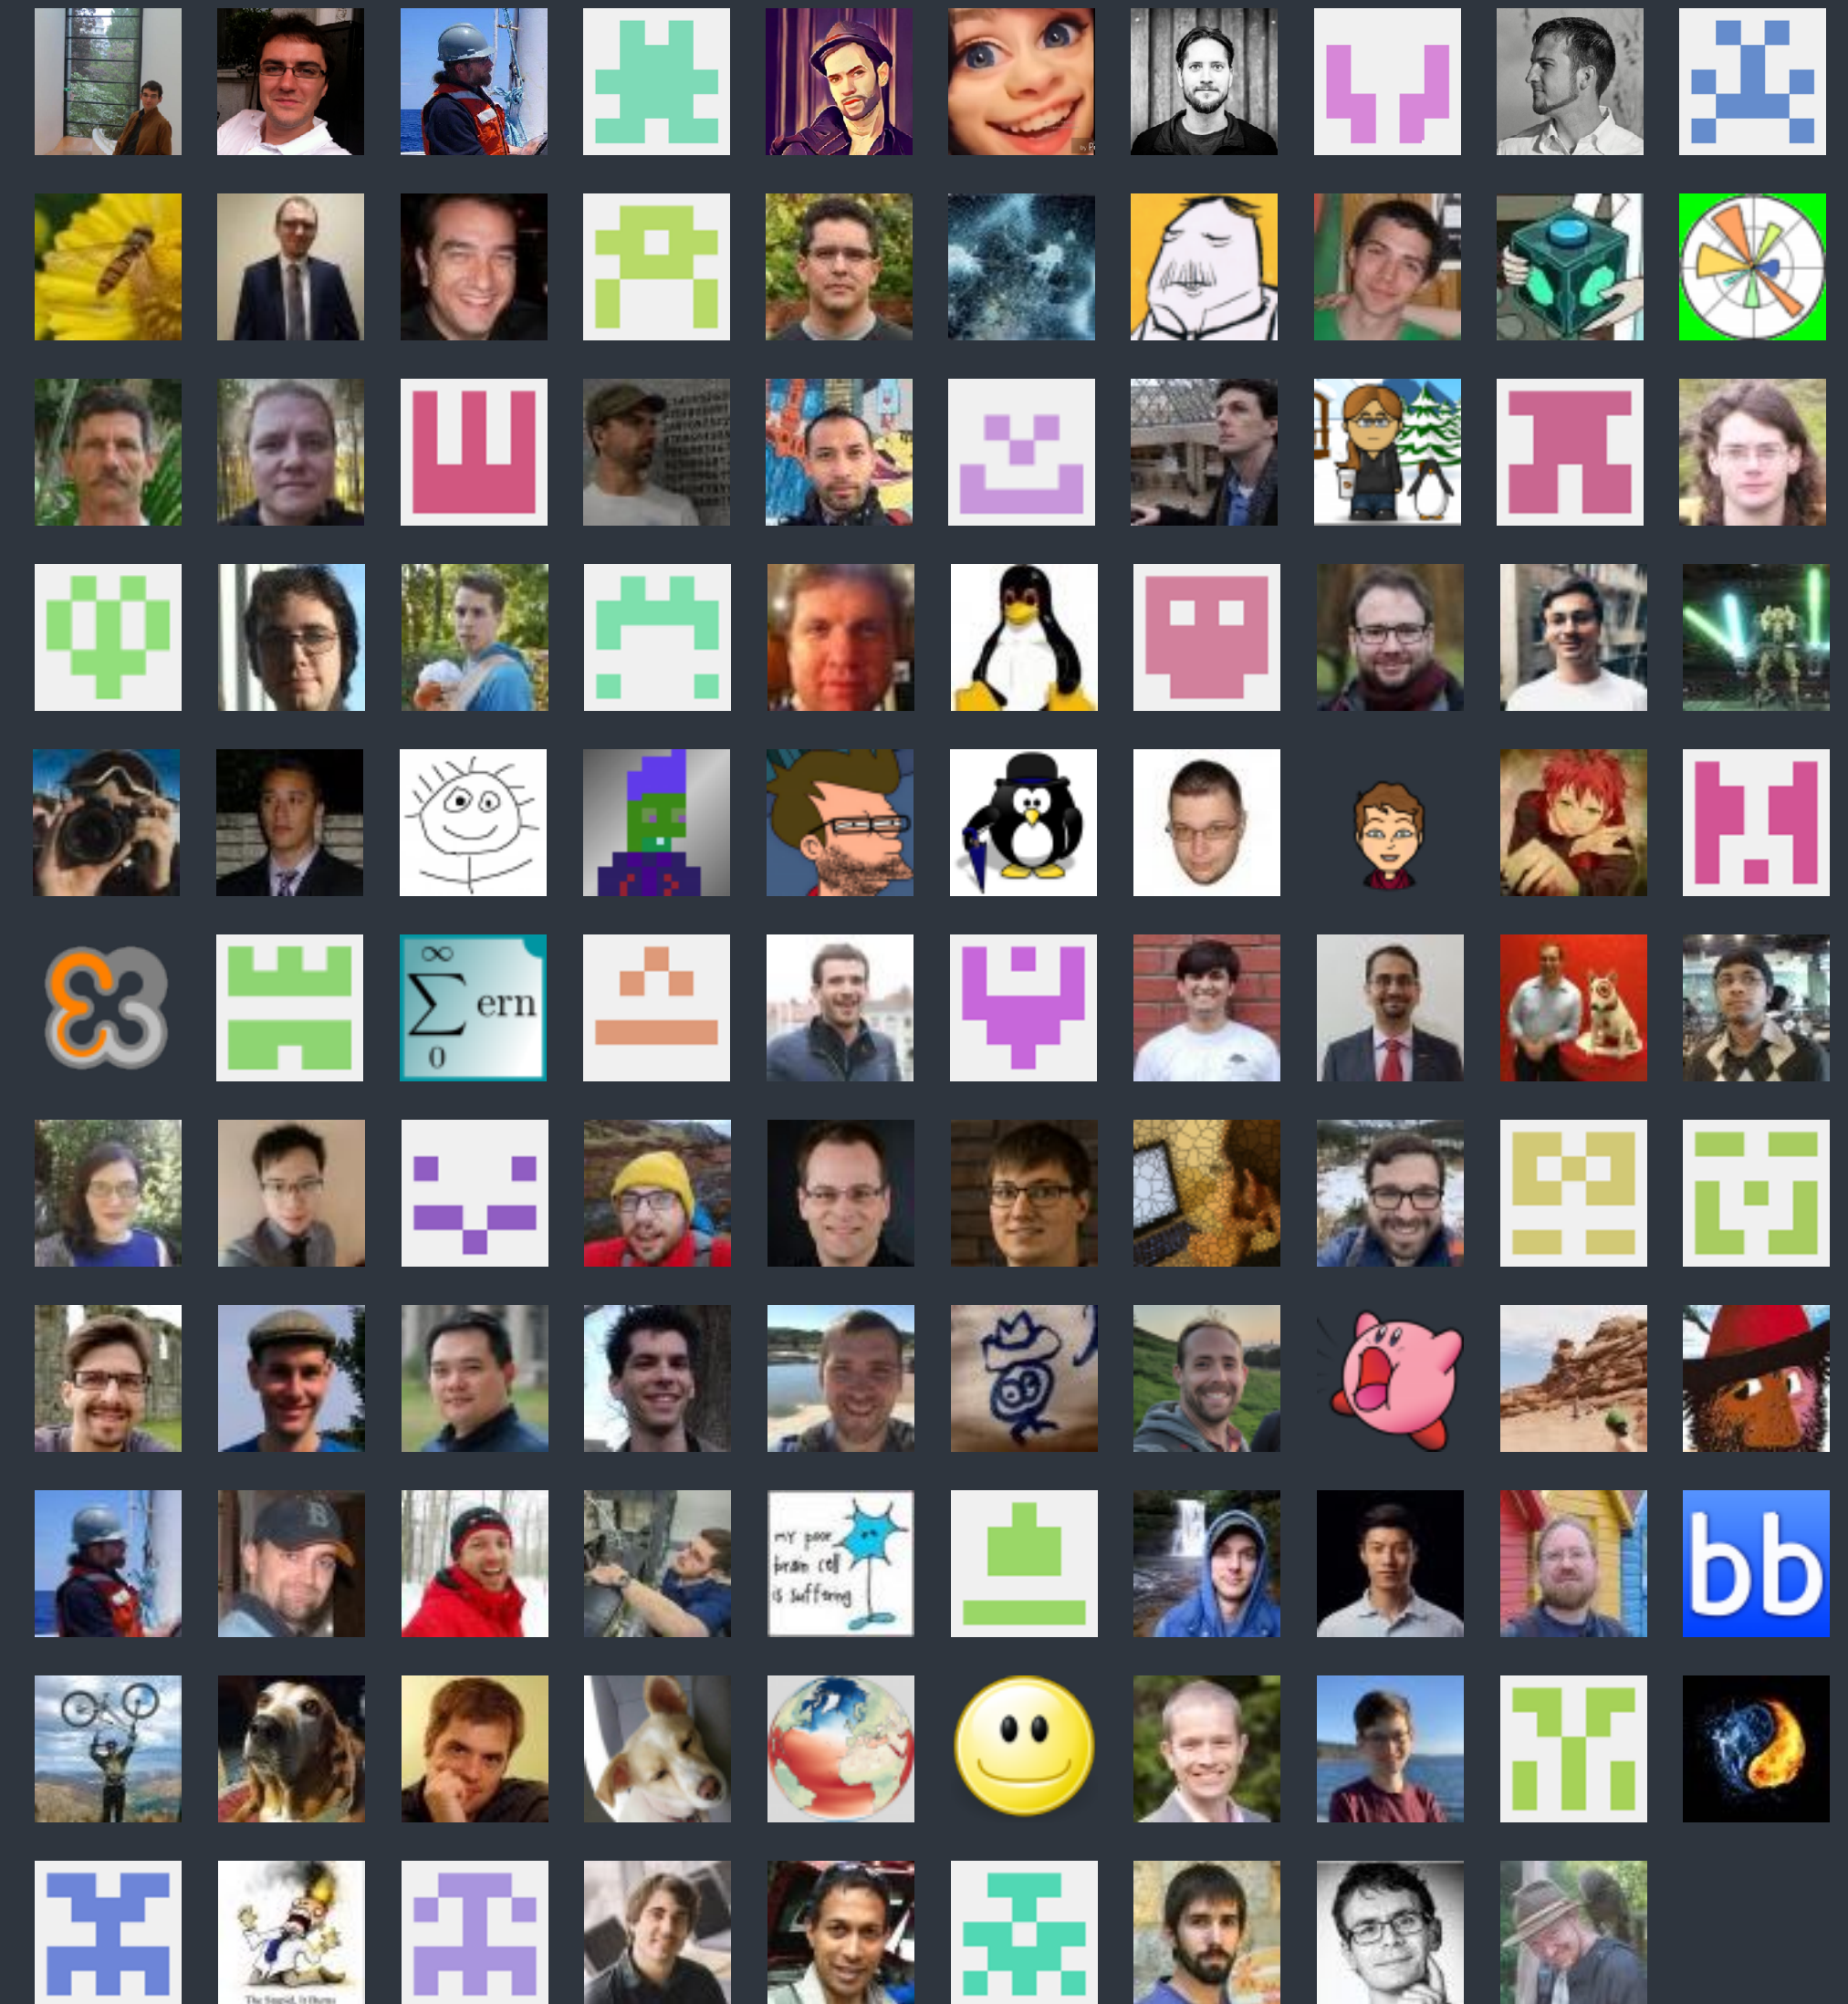
Community contributors



Code

```
1 import kiwisolver
2 import numpy
3 import matplotlib
```


Community contributors



Code

```
1 import kiwisolver
2 import numpy
3 import matplotlib
```

Community contributors





Microsoft



GitHub



One million pull requests
merged with Dependabot






+

PRM



+ Semmle

Open Source sei Dank: So fanden Experten eine Schwachstelle im Server der Corona-Warn-App



Die mobile App war nicht betroffen. (Foto: Firm / shutterstock)

23.11.2020, 14:13 Uhr • Lesezeit: 5 Min. Gerade keine Zeit? [Jetzt speichern und später lesen](#)

Ein Cybersecurity-Experte hat eine Sicherheitslücke in der Infrastruktur der Corona-Warn-App entdeckt. Die Hintergründe und der weitere Verlauf der Geschichte zeigen einmal mehr die Vorteile von Open-Source-Software.

Alvaro Muñoz' Job beim GitHub Security Lab ist es, Schwachstellen in Open-



Kathrin Stoll

Redakteurin Entwicklung & Design bei t3n

Verwandte Themen

Apps

Corona-Warn-App

Corona

Leak


Open Source

Server

**NIX MEHR VERPASSEN:
UNSERE NEWSLETTER**

E-Mail-Adresse

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



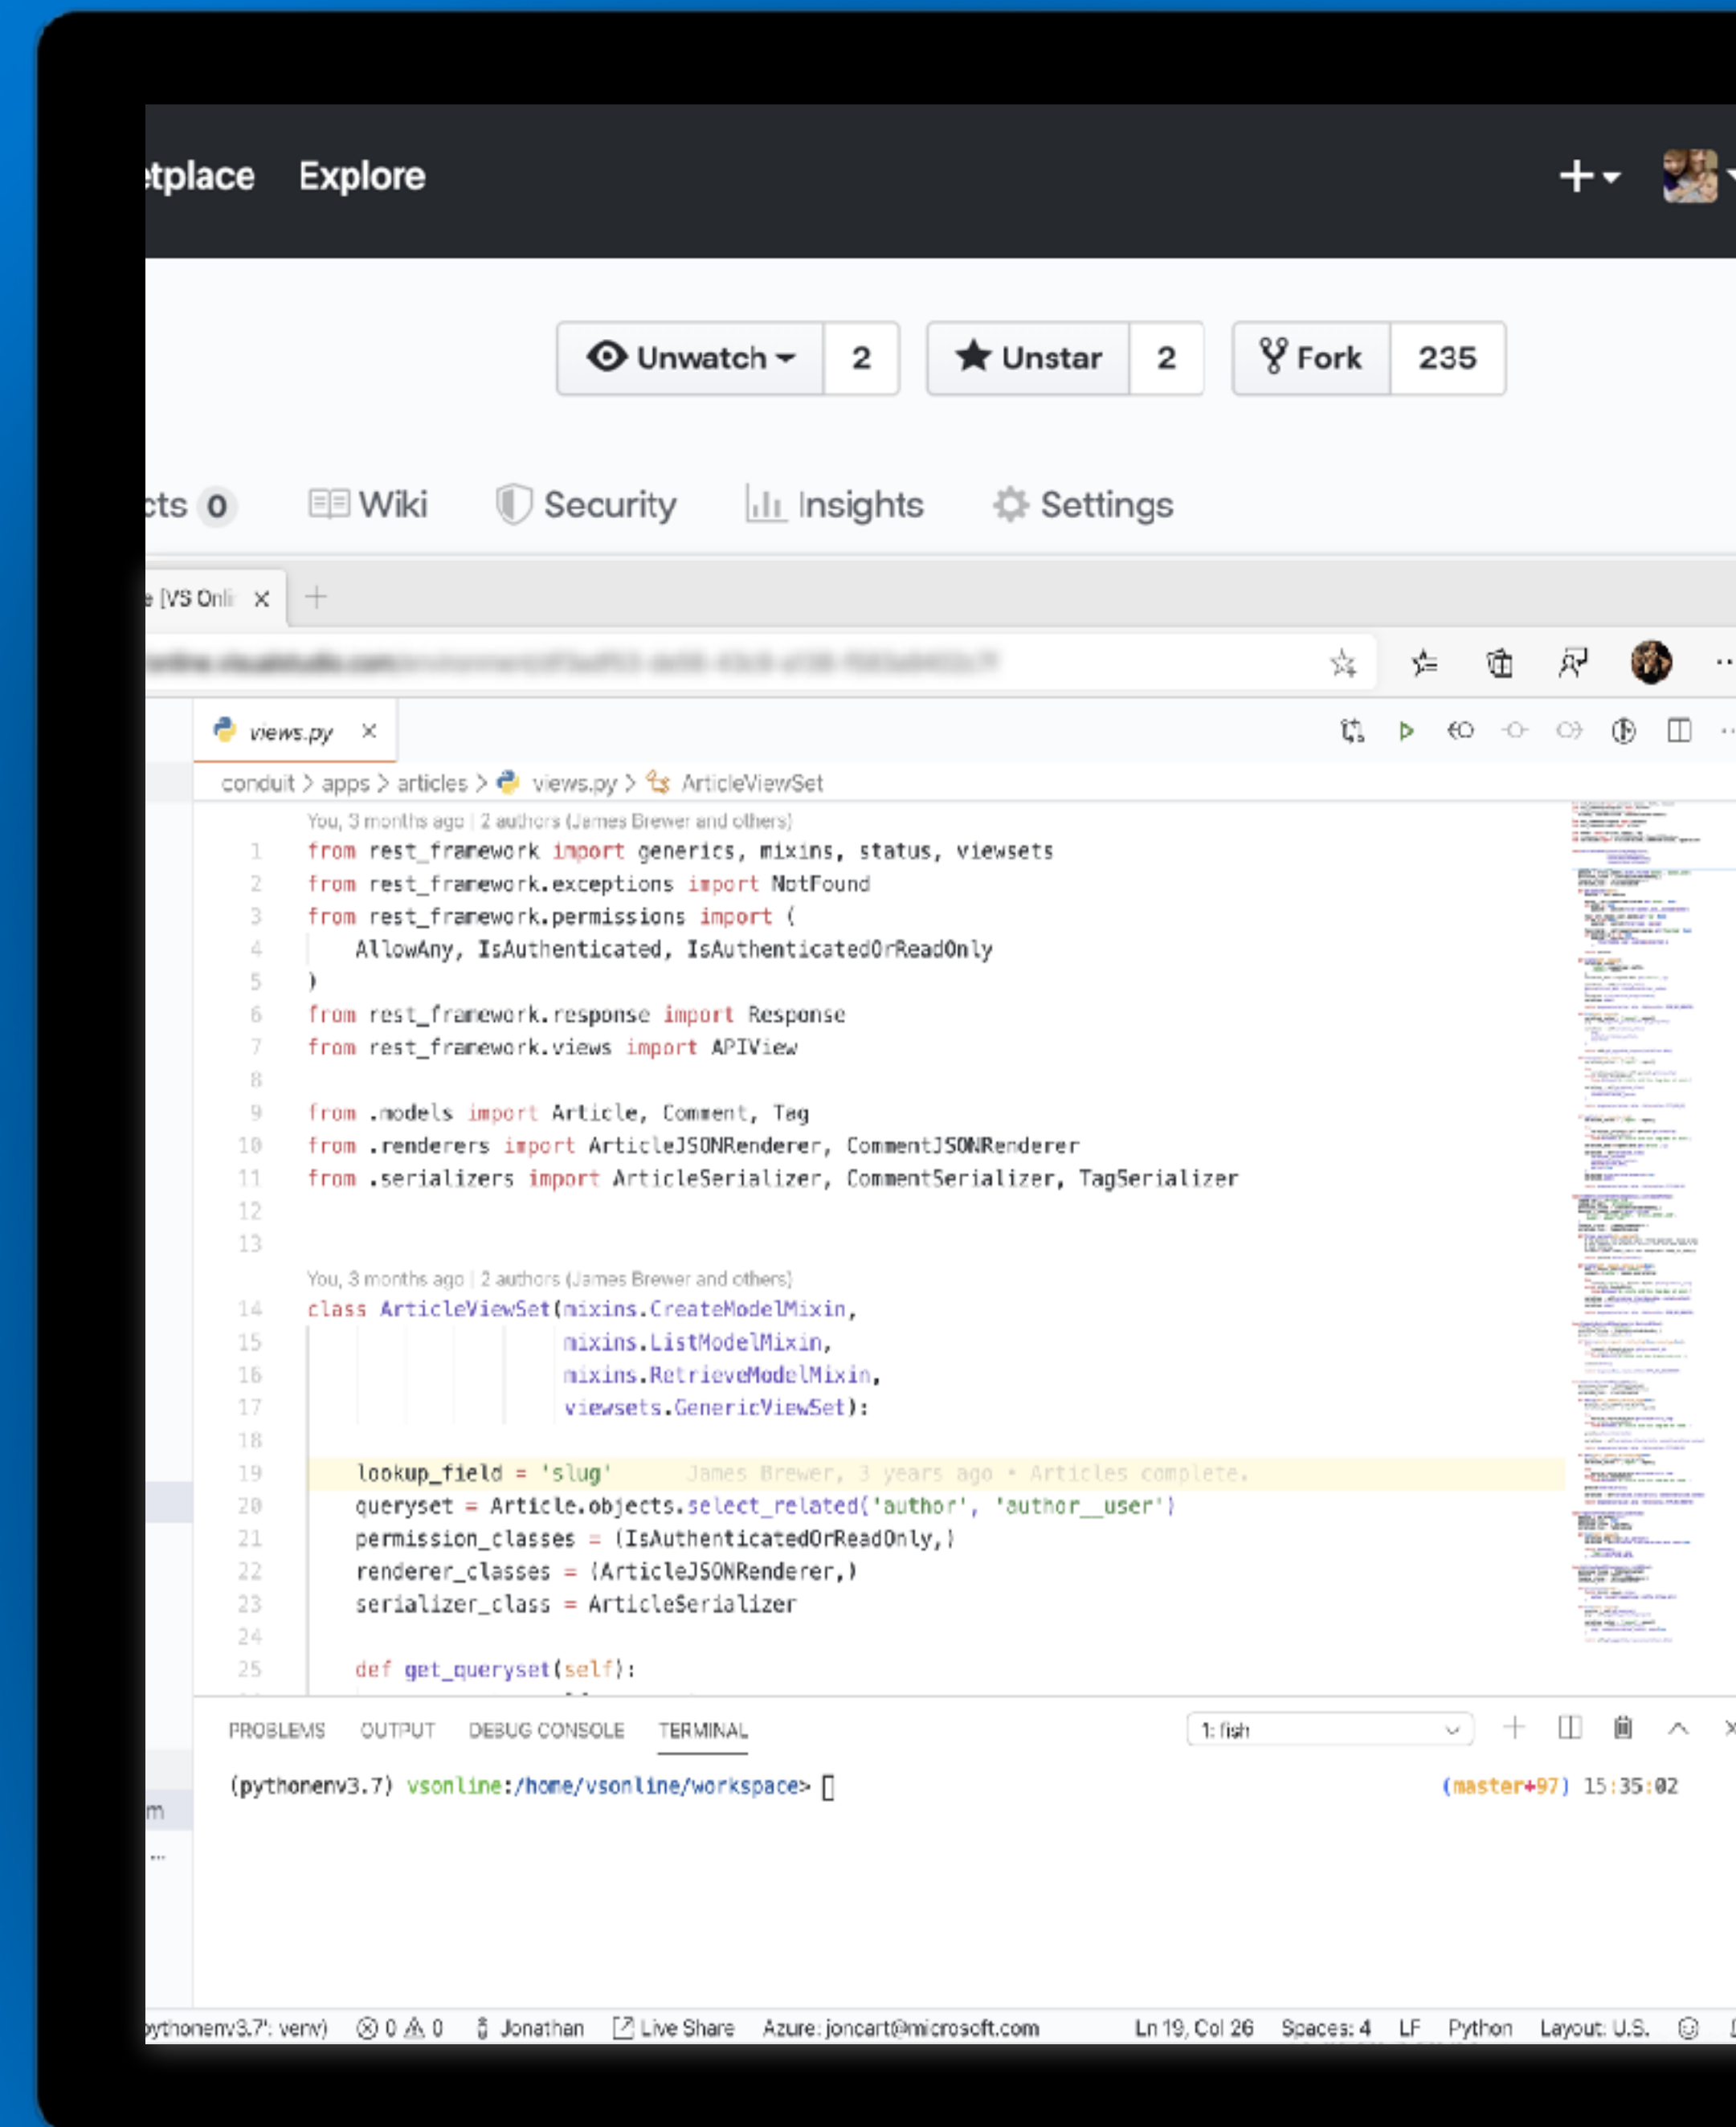


Codespaces

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





Test libzengithub package on main platforms

on: push

✗ test (macOS-10.14)



● test (windows-2016)

● test (windows-2019)

✓ test (ubuntu-16.04)

✓ test (ubuntu-18.04)

✓ `${{ format('Set up Python {0}', matrix.python-version) }}`

✓ Install C and C++ header files

✓ Install conan && conan package tools

✓ Generating conan user directory and building the solution

1027 ZenGitHub/1.0@jonico/stable (test package): Running test()

1028

1029 MMM. .MMM

1030 MMMMMMMMMMMMMMMMMMM

1031 MMMMMMMMMMMMMMMMMMM

1032 MMMMMMMMMMMMMMMMMMM

1033 MMMMMMMMMMMMMMMMMMM

1034 MMMMMMMMMMMMMMMMMMM

1035 MMMM::- -:::-- --:MMM

1036 MM~::~ 00~::~:~ 00~::~MM

1037 .. MMMMM:::00::+:::00::MMMM ..

1038 .MM::::: ._. ::::MM.

1039 MMMM;::::;MMMM

1040 -MM MMMMMMM

1041 ^ M+ MMMMMMM

1042 MMMMMMM MM MM MM

1043 MM MM MM MM

1044 MM MM MM MM

1045 .~MM~MM~MM~MM~.

1046 ~~~MM:~MM~MM~MM~

1047 ~~~~~

1048 ~~~~~

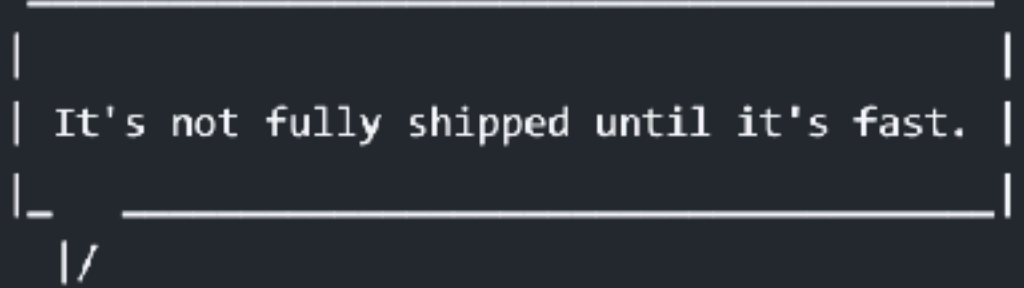
1049 :~~~~

1050

1051 [HOOK - attribute_checker.py] pre_export(): WARN: Conanfile doesn't have 'description'. It is recommended to add it as attribute

1052

✓ Complete job



GITHubBERS



574

REMOTE



315 / 54%

CURRENTLY IN SF

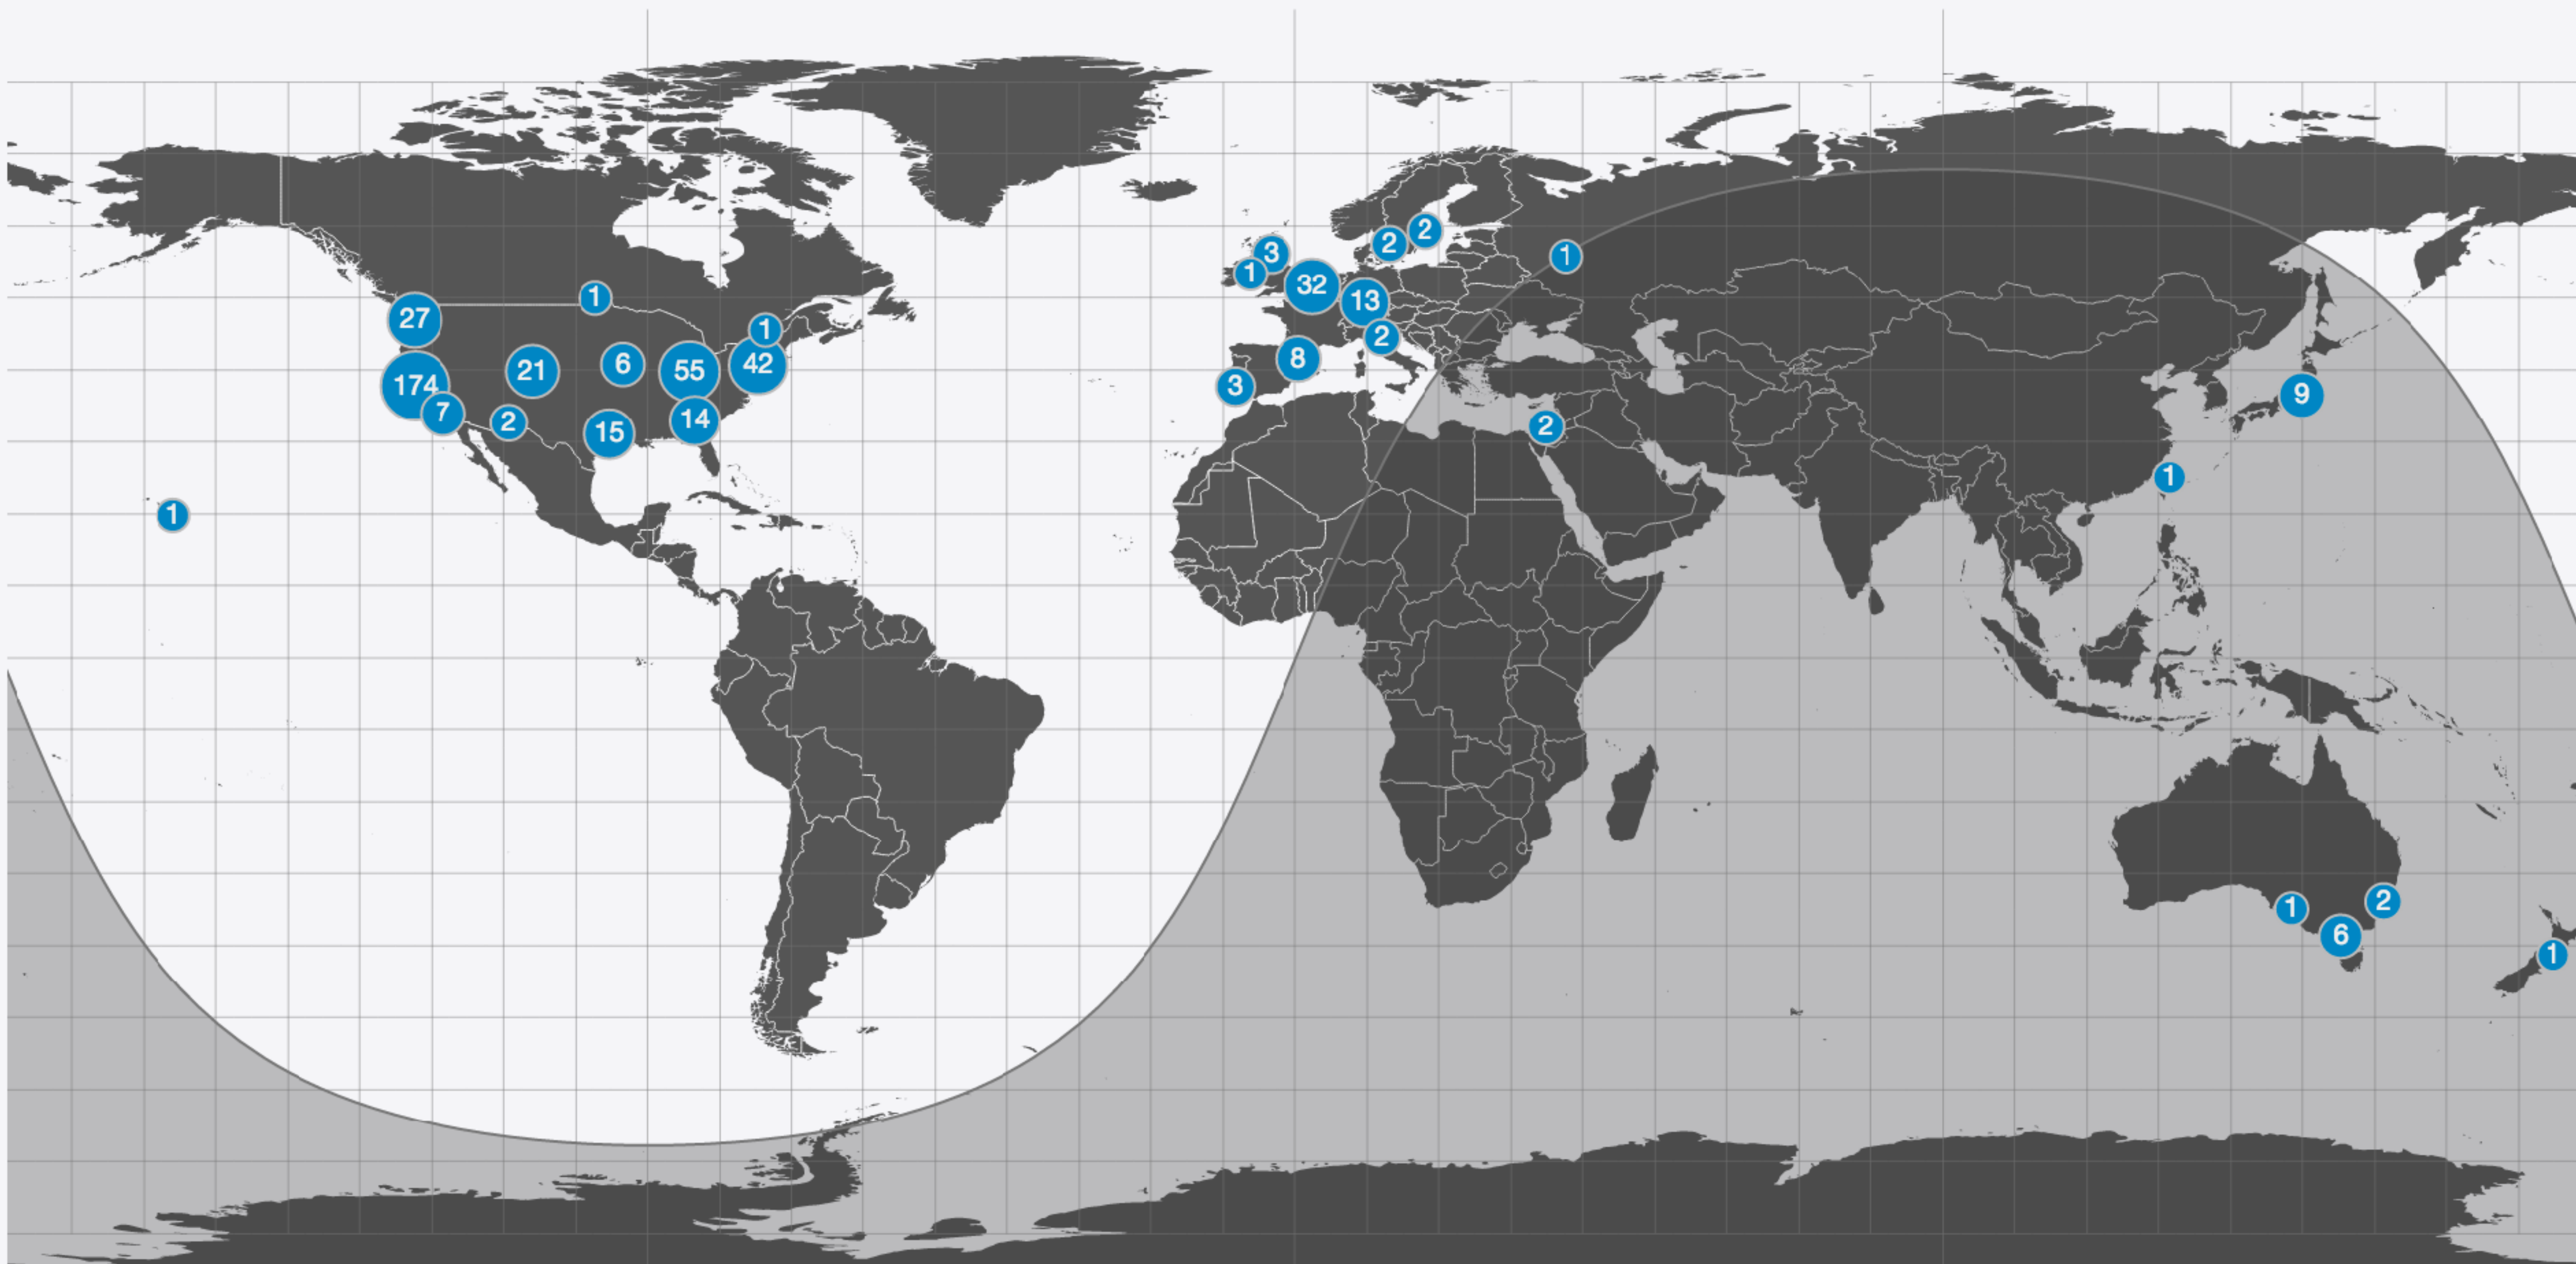


256 / 44%

NEW HIRES



18 / 3%



GITHUBBERS



1613

REMOTE



1229 / 76%

BASED IN SF

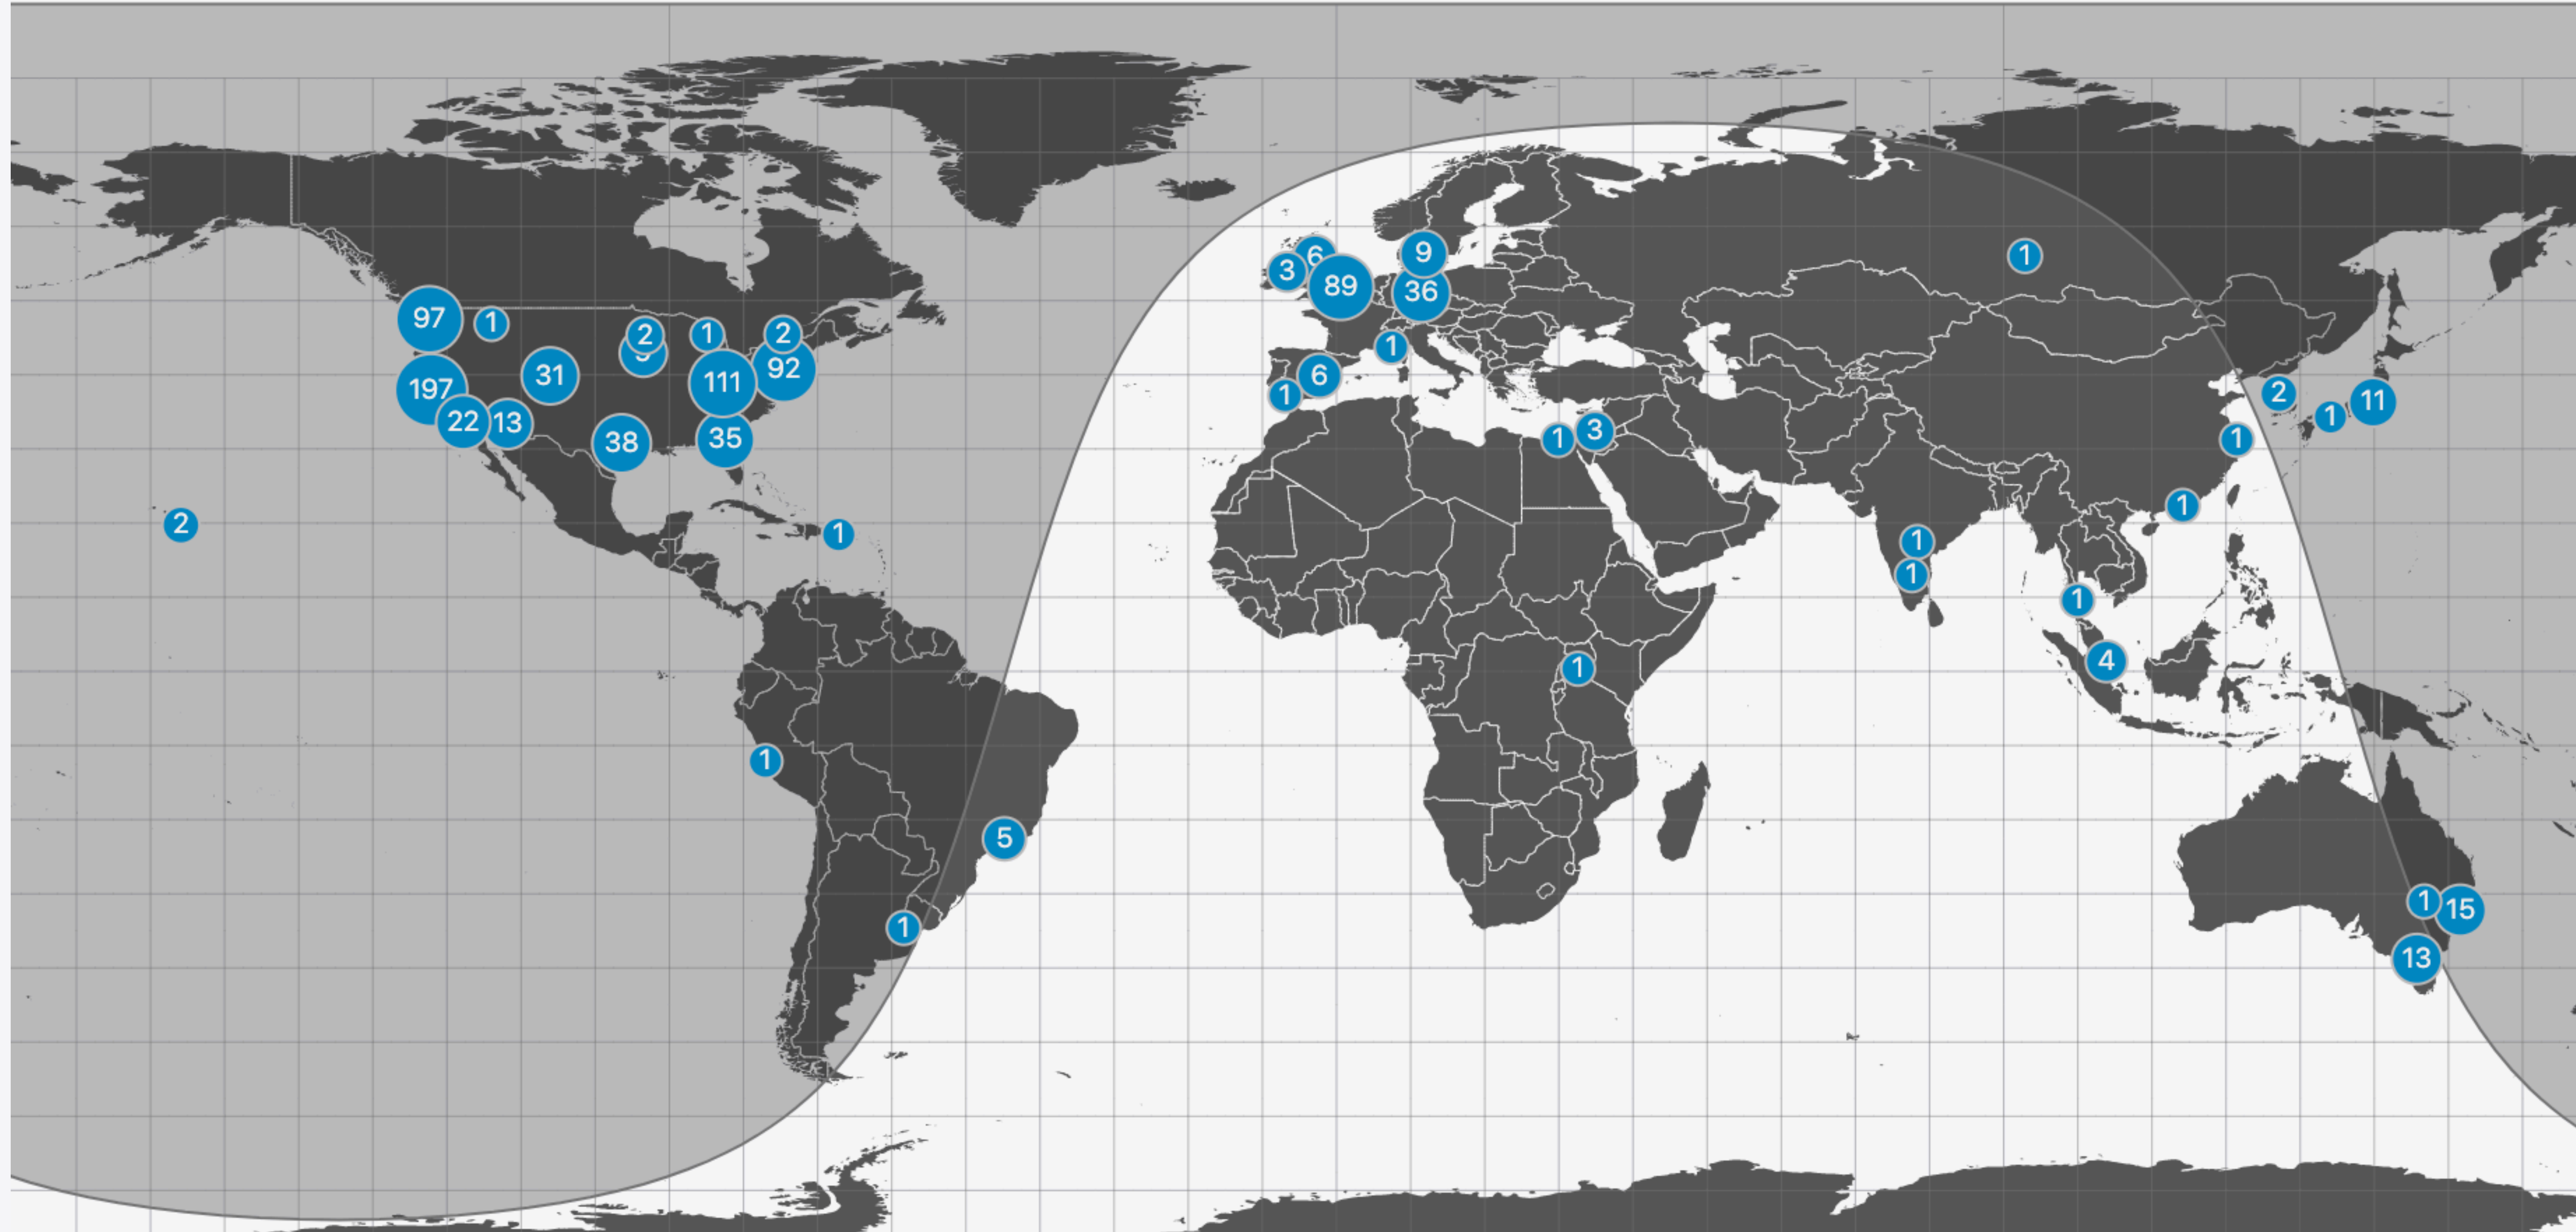


383 / 23%

NEW HIRES



195 / 12%



GITHUBBERS



2087

REMOTE



1687 / 80%

BASED IN SF

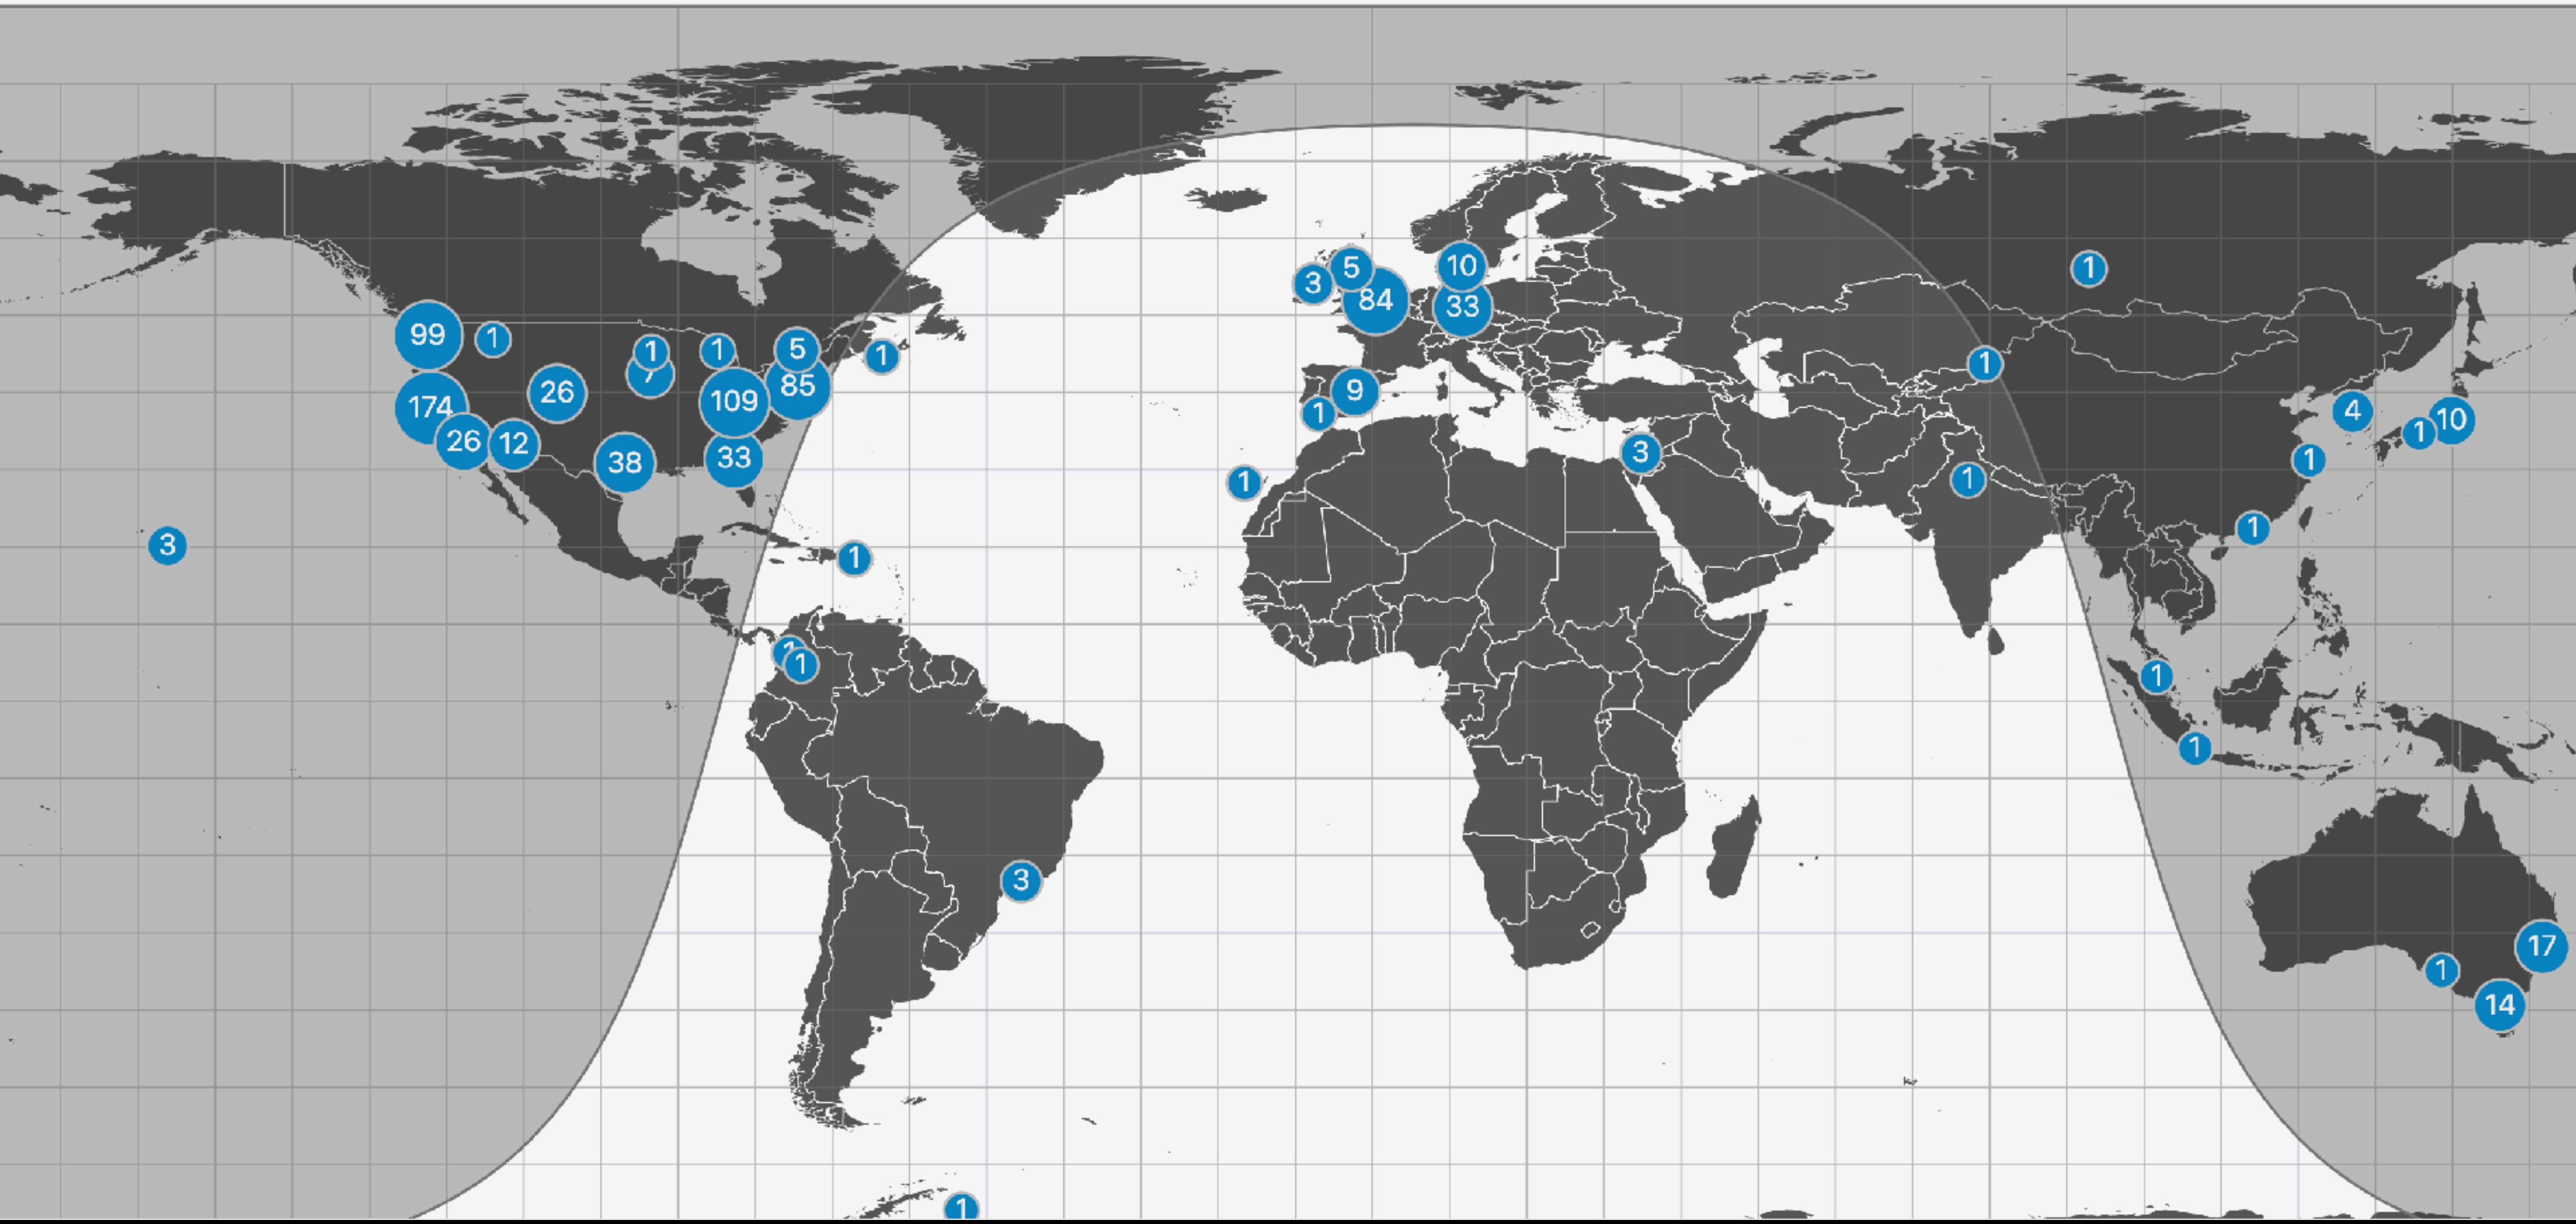



399 / 19%

NEW HIRES



39 / 1%

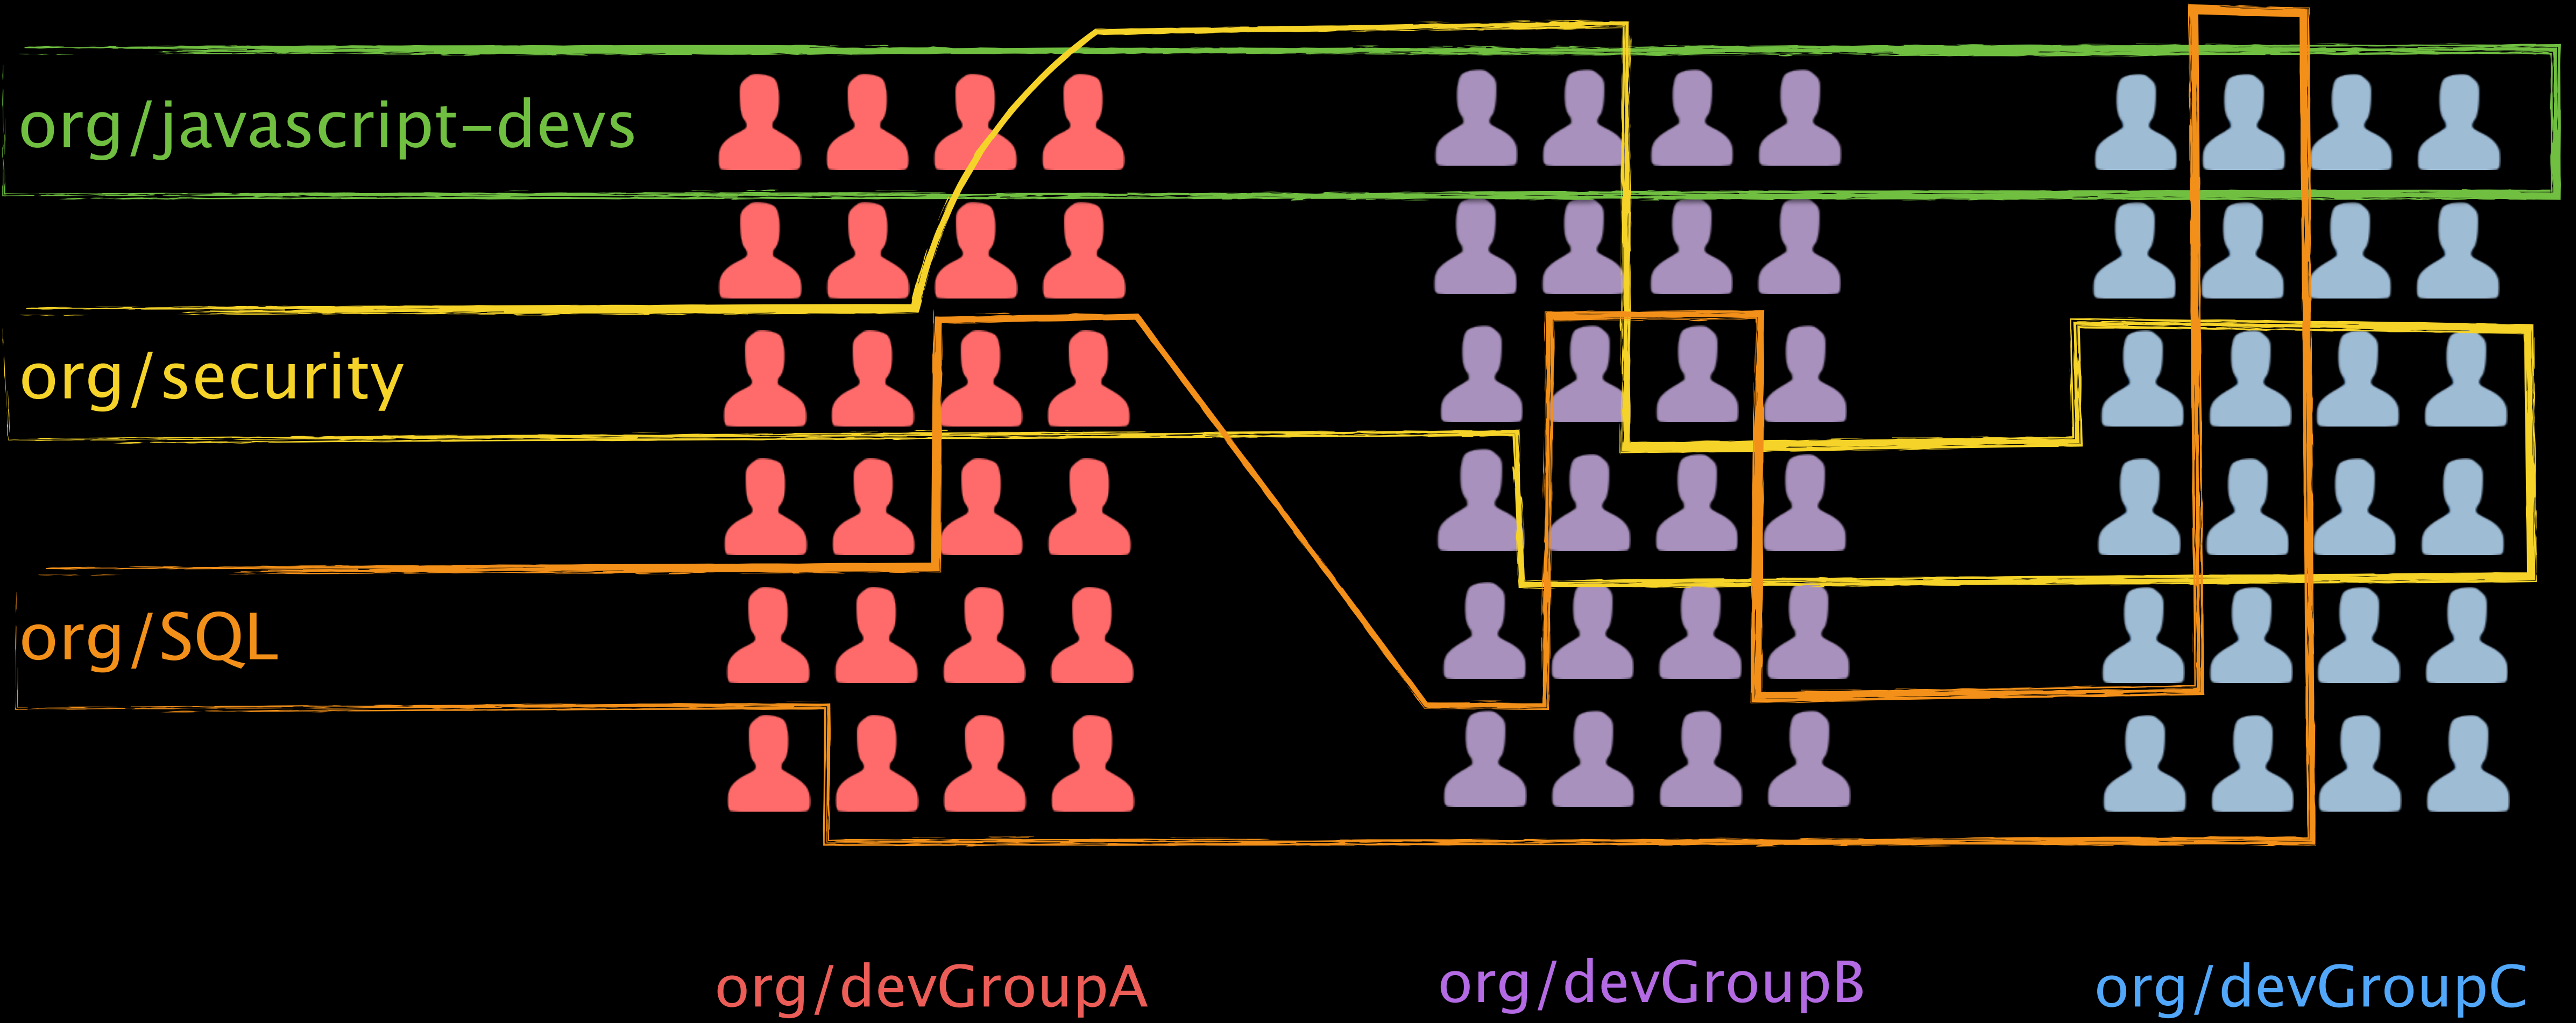






Avoid use of emails

If there is no link it does not exist



What stayed (and evolved)




Remote First
No Emails




ChatOps
Hubot

What changed heavily



Architecture
Languages



Flaky Test
Detection




Deploy
Trains


what's
in it for
ME?




Focus For Today



Architecture
Languages




Flaky Test
Detection



ChatOps
Hubot





Architecture / Languages



github
MODDA

APPS IN STYLE



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





KubeCon



CloudNativeCon

North America 2017

Keynote: Kubernetes at GitHub

Jesse Newland, Principal Site Reliability Engineer, *GitHub*



KELSEY GILMORE-INNIS


SITE RELIABILITY ENGINEER
GITHUB

INFRASTRUCTURE AS PRODUCT
BUILDING GITHUB'S FUTURE

Full screen (f)



▶ ⏪ 🔊 0:09 / 33:17






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](#)



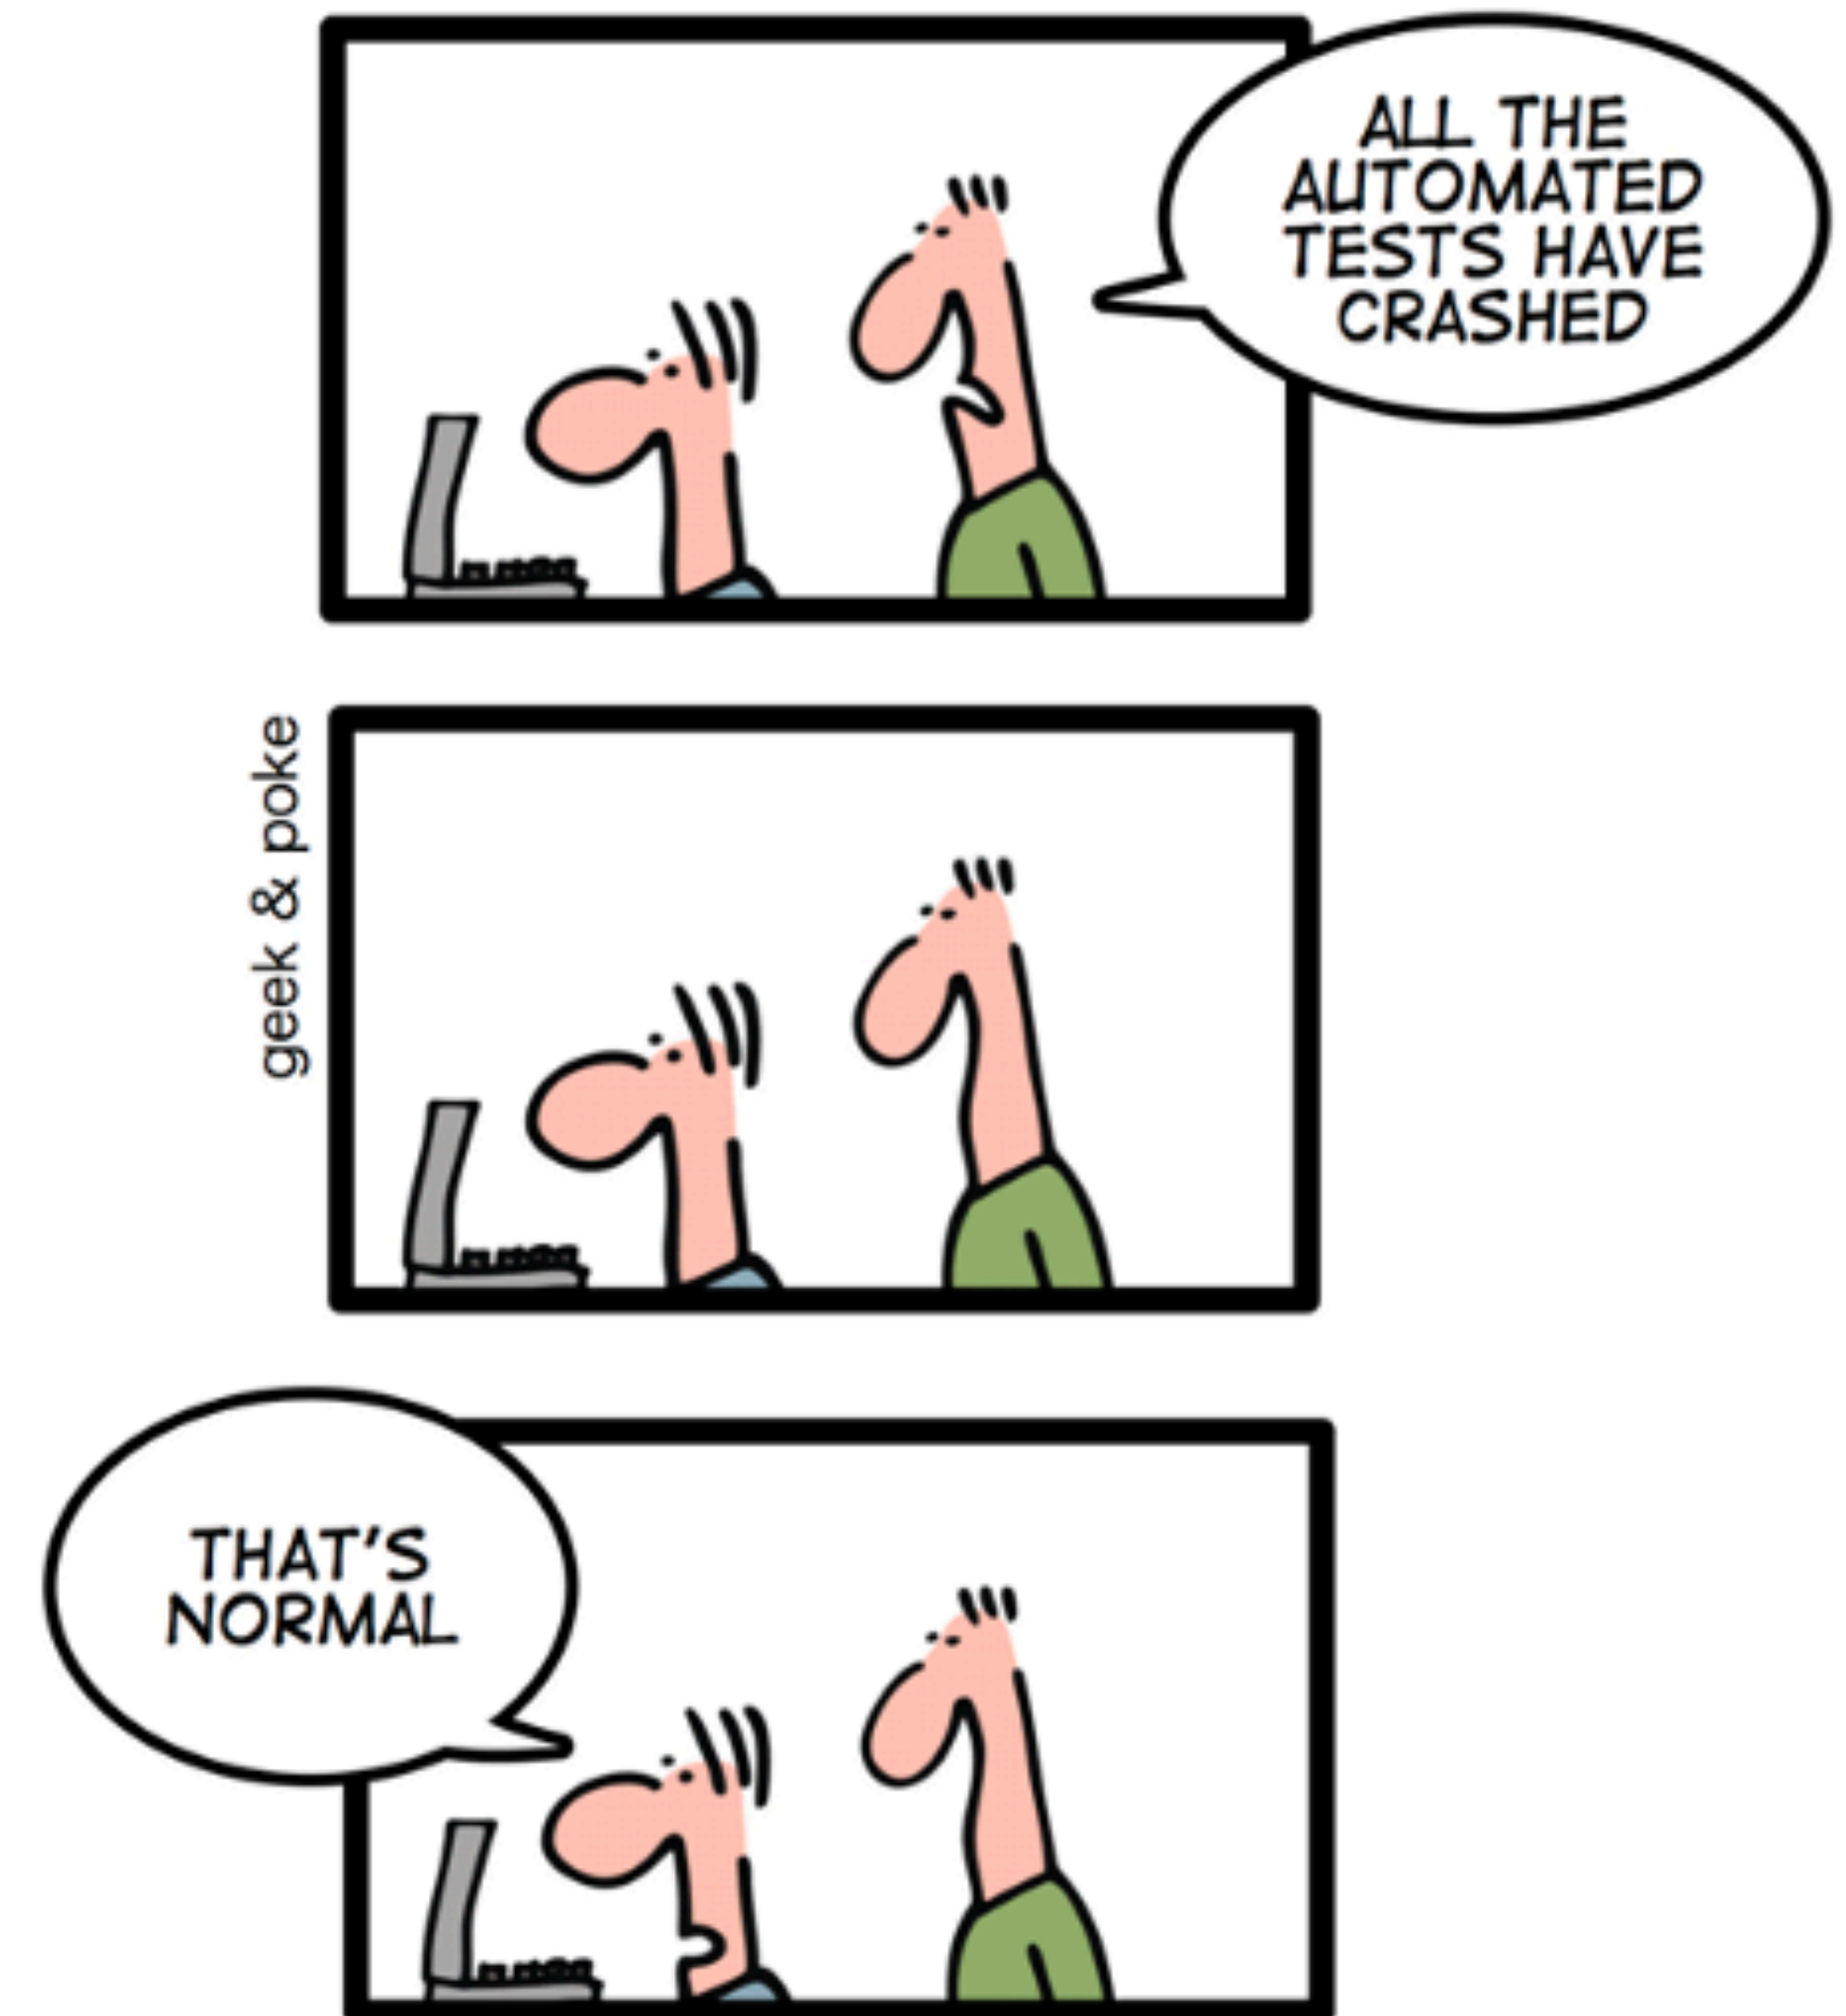


Flaky Test Detection

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

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



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

1 test

$$P_{\text{fail}}(1) = 0.001$$

10 tests

$$P_{\text{fail}}(10) = 0.009955$$

100 tests

$$P_{\text{fail}}(100) = 0.0952$$

1000 tests

$$P_{\text{fail}}(1000) = 0.6323$$





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



Open Flakes

Closed Flakes

Flake Candidates



0 Selected

144 Total

NWO: github/github

Buildables: github, github-all-features, enterprise

Clear

Add



TeamDestructionDestroyOperationContext / caf0604f8f7

Test two operations race

Clear

Add



github-all-features / test-make-rails-6-default **Flake Candidate**

Occurred about 1 hour ago • Built at 2019-08-20 13:20:37 UTC • Built on janky-nunix-



github-all-features / steves-analytics **Flake Candidate**

Occurred 4 days ago • Built at 2019-08-16 18:01:13 UTC • Built on janky-nunix-f57d910.private-us-east-1.github.net



github-all-features / queue-prep-t **Flake Candidate**

Occurred 5 days ago • Built at 2019-08-15 15:28:16 UTC • Built on janky-nunix-1



github-all-features / sponsors-criteria-migration **Flake Candidate**

Occurred 13 days ago • Built at 2019-08-07 19:50:51 UTC • Built on janky-nunix-1



github-all-features / split-up-get-refs-endpoint **Flake Candidate**

Occurred 19 days ago • Built at 2019-08-01 14:35:23 UTC • Built on janky-nunix-f



TeamDestructionDestroyOperationContext / 93ff7a917b0b36583

Test throttle the deletes

Clear

Add



github-all-features / test-make-rails-6-default **Flake Candidate**

Occurred about 1 hour ago • Built at 2019-08-20 13:20:37 UTC • Built on janky-nunix-1



github-all-features / steves-analytics **Flake Candidate**

Occurred 4 days ago • Built at 2019-08-16 18:01:13 UTC • Built on janky-nunix-



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)
- 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





SHARE

 SHARE
874

 TWEET

 PIN

 COMMENT
11

 EMAIL

CADE METZ BUSINESS 10.23.15 7:00 AM

THE MOST IMPORTANT STARTUP'S HARDEST WORKER ISN'T A PERSON

'It's a new way of working.'

—SAM LAMBERT, GITHUB

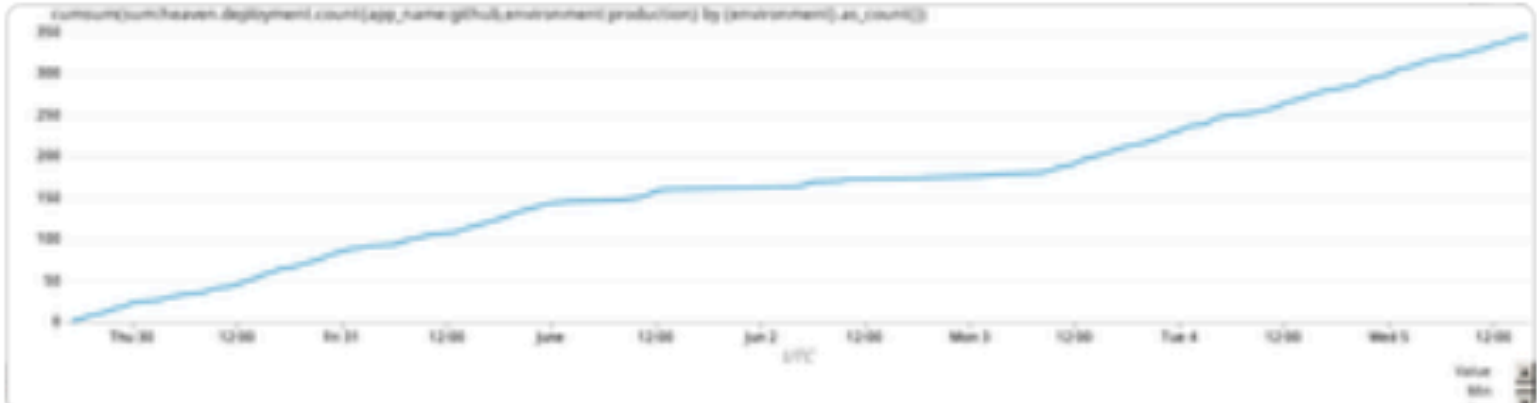
Sam Lambert, the director of systems at GitHub, calls Hubot “the hardest working GitHubber.” That’s a company-wide in-joke. Hubot isn’t really a GitHubber. He’s a bit of software that plugs into the GitHub chat system. About five years ago, a guy named Ryan Tomayko built Hubot as an easier way for the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

hubot APP 6:09 PM
`cumsum(sum:heaven.deployment.count{app_name:github,environment:production} by {environment}.as_count()) (23 kB)`



The chart displays a steady upward trend in the cumulative deployment count for the 'github' application in the 'production' environment from late 2016 to early 2017. The y-axis represents the count, ranging from 0 to 300, and the x-axis shows dates from late 2016 to early 2017.

carlosmn 11:45 AM
`.maint enable gitbackups-maintenance-[redacted] destroying`

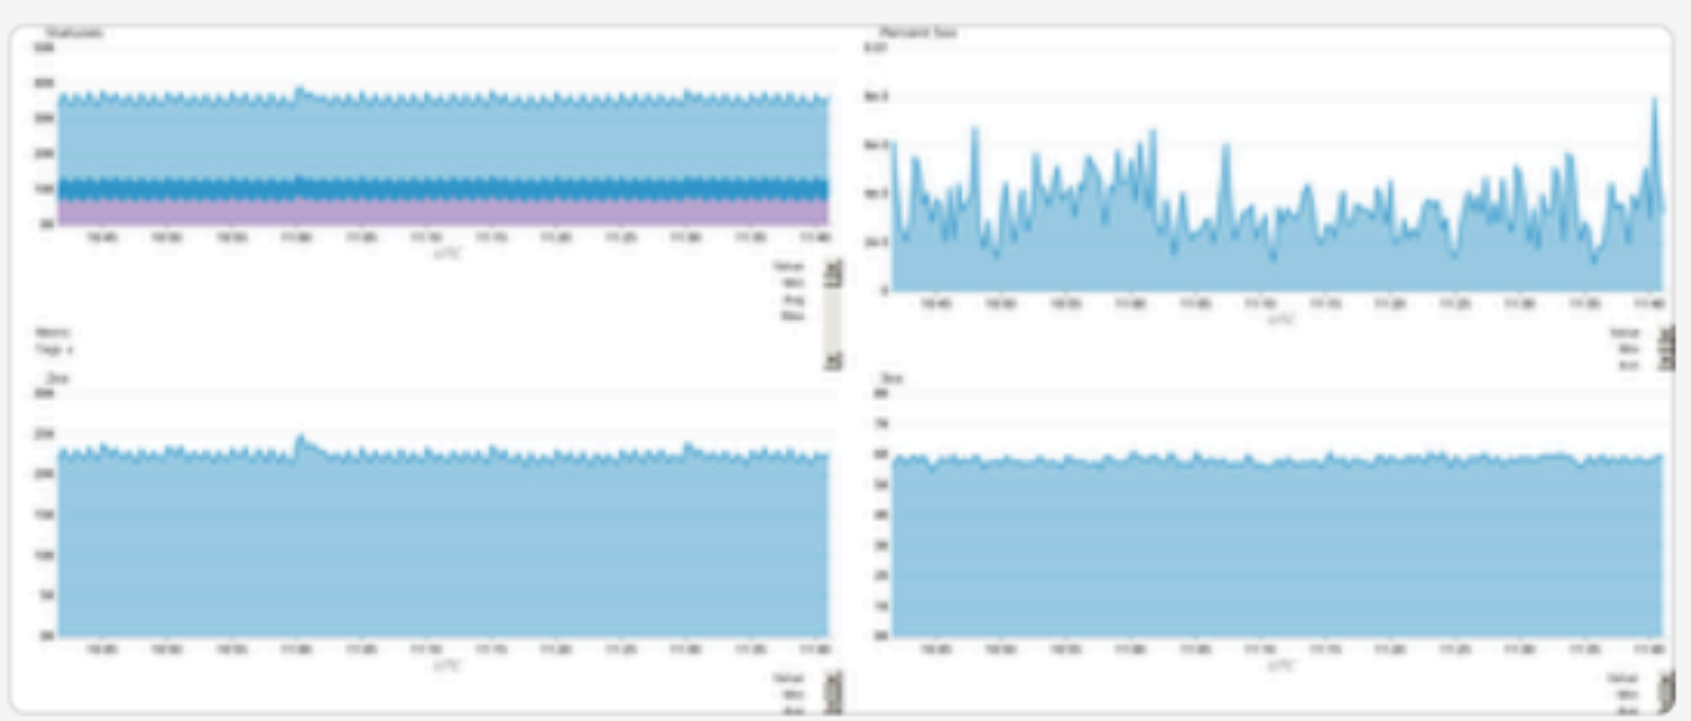
hubot APP 11:45 AM
@carlosmn: authentication required for `.maint enable`. Check your mobile device.

hubot APP 11:45 AM
`gitbackups-maintenance-[redacted] - maintenance mode enabled successfully! Window will expire in 60 minutes.`

carlosmn 11:47 AM
`.instance destroy hostname=gitbackups-maintenance-[redacted]`

brntbeer 1:41 PM
`.dash me -1h glb.services --glb-service dotcom`

hubot APP 1:41 PM
`glb.services (117 kB)`



The dashboard displays four graphs for 'glb.services' metrics over a 1-hour period. The top-left graph shows a stable metric around 100. The top-right graph shows a highly volatile metric fluctuating between 0 and 100. The bottom-left graph shows a stable metric around 100. The bottom-right graph shows a stable metric around 100.

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: <https://github.com/github/github/issues>

sev3: [#pe-notifications](#)



**WORKED FINE IN
DEV**

OPS PROBLEM NOW

[3]

No one
talks
to each
other

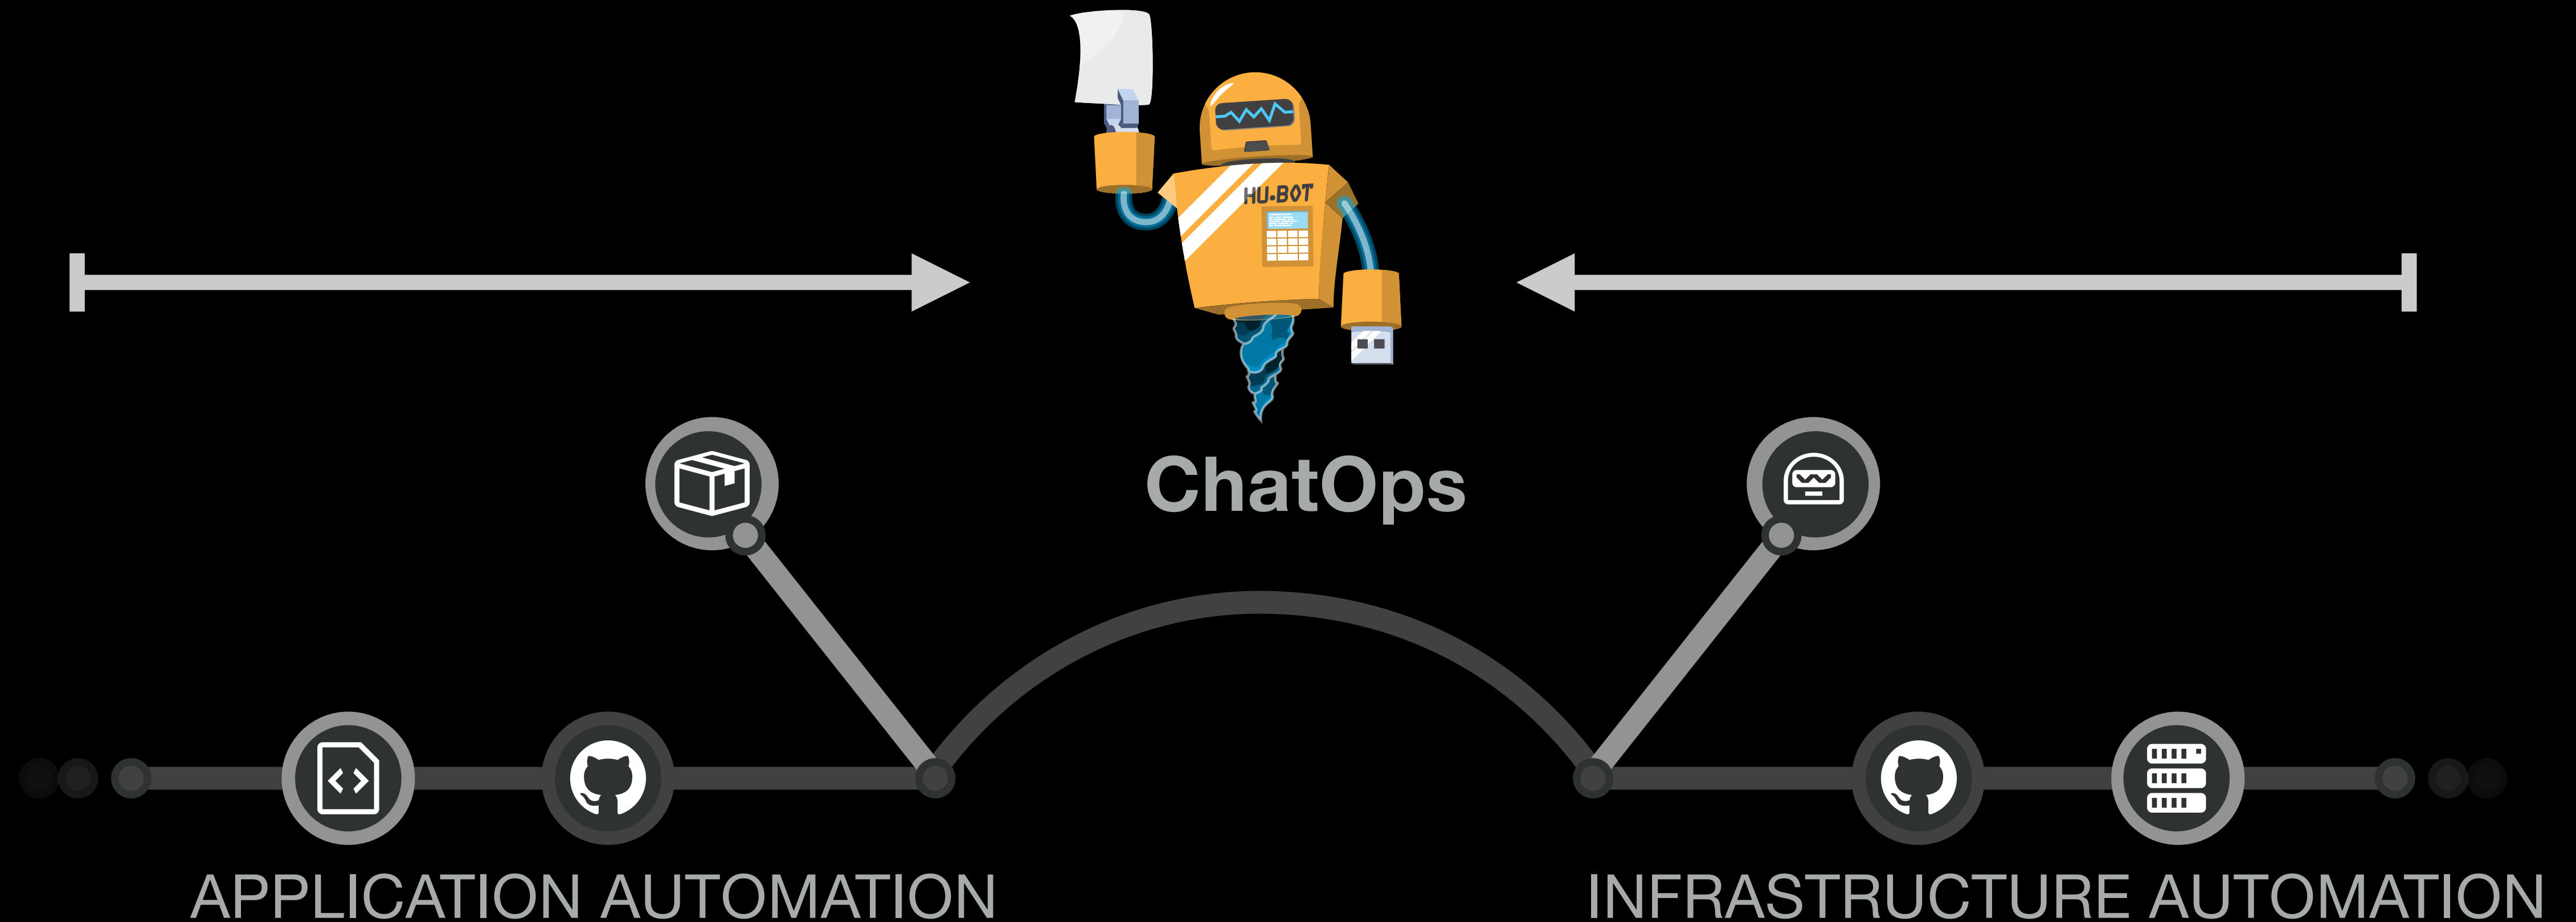


WELCOME TO
OPERATIONS
PLACE NEW
WORK HERE



eWEEK

ChatOps as the cultural glue



jonahberquist 6:46 PM
.mysqlproxy disable github-
[redacted]

hubot BOT 6:46 PM
disabling github-
[redacted]
github-
[redacted]::
disabled server mysql/github-
[redacted] Show more...

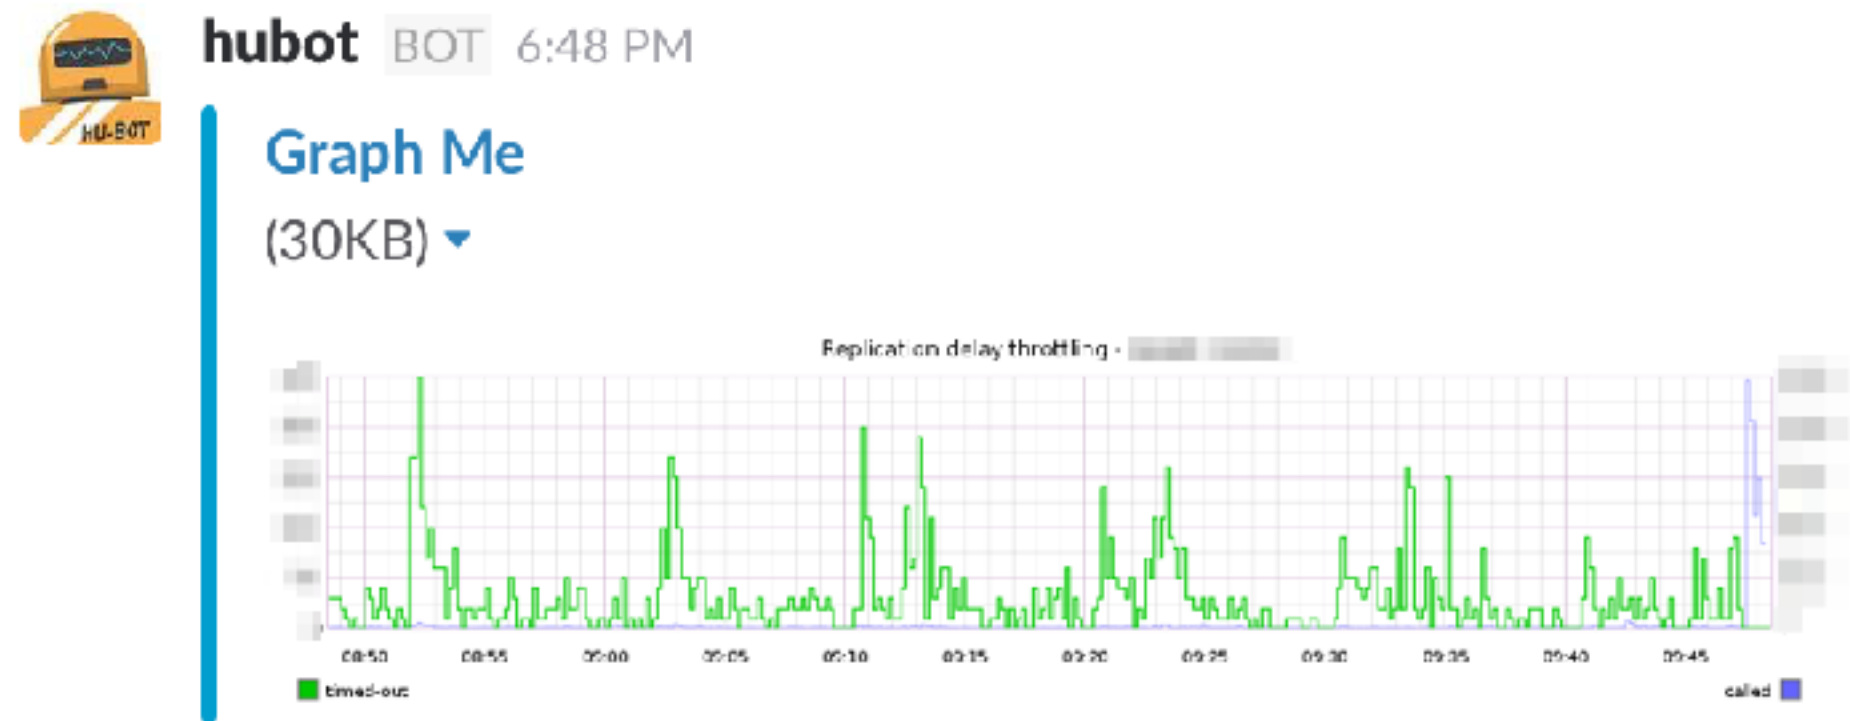
tomkrouper 6:47 PM
Do we want to re-up durability on a couple replicas to make sure that is fine?

samlambert 6:47 PM
.status green

hubot BOT 6:47 PM
samlambert: Status updated to green
samlambert: I updated GitHub Chat's status, too.

zerowidth 6:48 PM
.graph me -1h alias(alpha(secondYAxis(github-
[redacted]),0.5),"called") +
aliasByNode(maximumAbove(github-
[redacted] timed-out,0),3) +
title=Replication+delay+throttling+-
+mysql1+cluster&hideLegend=false&template=plain&lineMode=staircase

jnunemaker 6:48 PM
deploying doubling the threshold for user session writes





PASTE EMBLEM HERE

HUBOT

*(note: it's pronounced
hew-bot)*

A CUSTOMIZABLE,
LIFE EMBETTERMENT ROBOT

COMMISSIONED BY GITHUB

4. AFFILIATED COMPANY
(INCLUDE AFFILIATED COMPANY LOGO (IF APPLICABLE))

5. SIGNATURE OF INVENTOR
PLEASE DO NOT ATTEMPT TO SIGN WHILE INEBRIATED

VIEW HUBOT'S DOCUMENTATION
(LEARN ABOUT GETTING STARTED, ETC.)

VIEW HUBOT'S SOURCE CODE
(VIA [HTTP://GITHUB.COM/GITHUB/HUBOT/](http://github.com/github/hubot/))





FIG. 1 — HUBOT SCHEMATICS




Summary

Focus For Today



Architecture
Languages



Flaky Test
Detection



ChatOps
Hubot



Unwrap your GitHub Secret Link

<https://github.com/winterfest-2020/unwrap-parcel/issues/new/choose>





Thank you