

Open Issues/PRs (SIG: All, Milestone: v1.10)



Who, What, When

Exploring Kubernetes Development through DevStats

#devstats

KubeCon EU 2018

Panel

Josh Berkus
Red Hat

Łukasz Gryglicki
CNCF

#devstats

Elsie Phillips
CoreOS/Red Hat

Aaron Crickenberger
Samsung

Tim Pepper
VMware

Panel

Josh Berkus
SIG-Release

Łukasz Gryglicki
DevStats Lead

Elsie Phillips
SIG-ContribEx

Aaron Crickenberger
SIG-Testing/SC

Tim Pepper
SIG-VMware/Release

#devstats

I've got to know!

ContribEx wants to check community health.

Release Team needs to look for blockers.

#devstats

I've got to know!

Steering Committee needs to forecast policy changes.

CNCF wants to advertise Kubernetes' momentum.

#devstats

Who?

Who is contributing?

Who is doing the most code reviews?

Which SIGs have the most issues?

Which employers are contributing?

#devstats

What?

What activity is happening right now?

What is the label PRs block on?

How many issues are being filed in each repository?

#devstats

When?

When does an issue get its first response?

How long does it take to merge a PR?

When are our busy/slow periods?

April 17, 2018 – April 24, 2018

Period: 1 week ▾

Overview

 **299 Active Pull Requests**

 **213 Active Issues**

 **168**

Merged Pull Requests

 **131**

Proposed Pull Requests

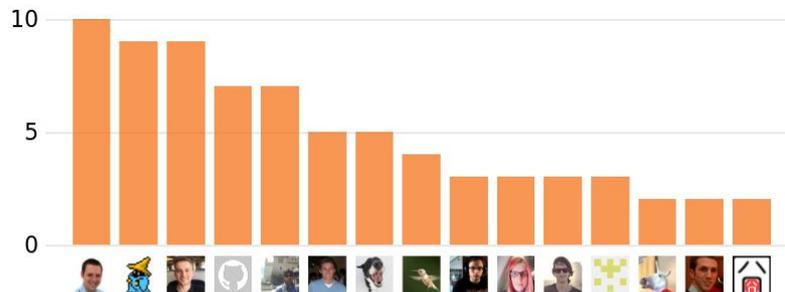
 **123**

Closed Issues

 **90**

New Issues

Excluding merges, **66 authors** have pushed **111 commits** to master and **134 commits** to all branches. On master, **1,243 files** have changed and there have been **18,603 additions** and **21,697 deletions**.



GitHub stats limitations

- Only a few stats (commits and “traffic”)
- No reporting by label
- Only top contributors
- Time scales not very configurable

Kubernetes quirks

Most “traffic” is from bots

No commits, only merges (by bot!)

Need to aggregate groups of repos

Enter DevStats

CNCF-funded project

Developer Łukasz Gryglicki

Started July 2017

WIP

[#devstats](#)



First non-author activity repository groups ▾



now-2y to now-2w



Period

Week ▾

Repository group

All ▾

Releases



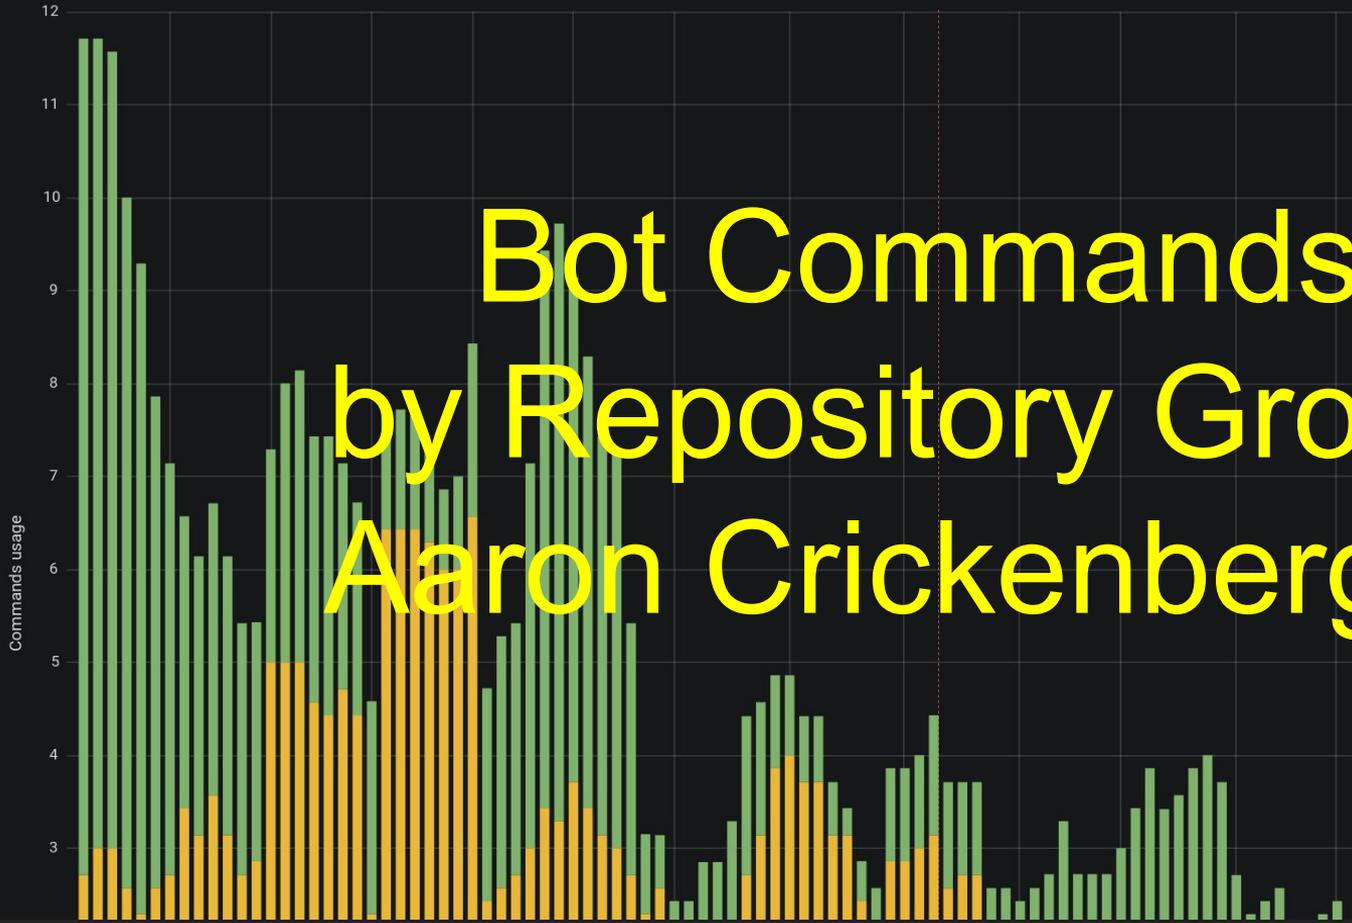
Time before any comments or activities from someone other than the author (All, Week)



PR Time to Engagement
Elsie Phillips



Bot commands usage (All, 7 days MA)

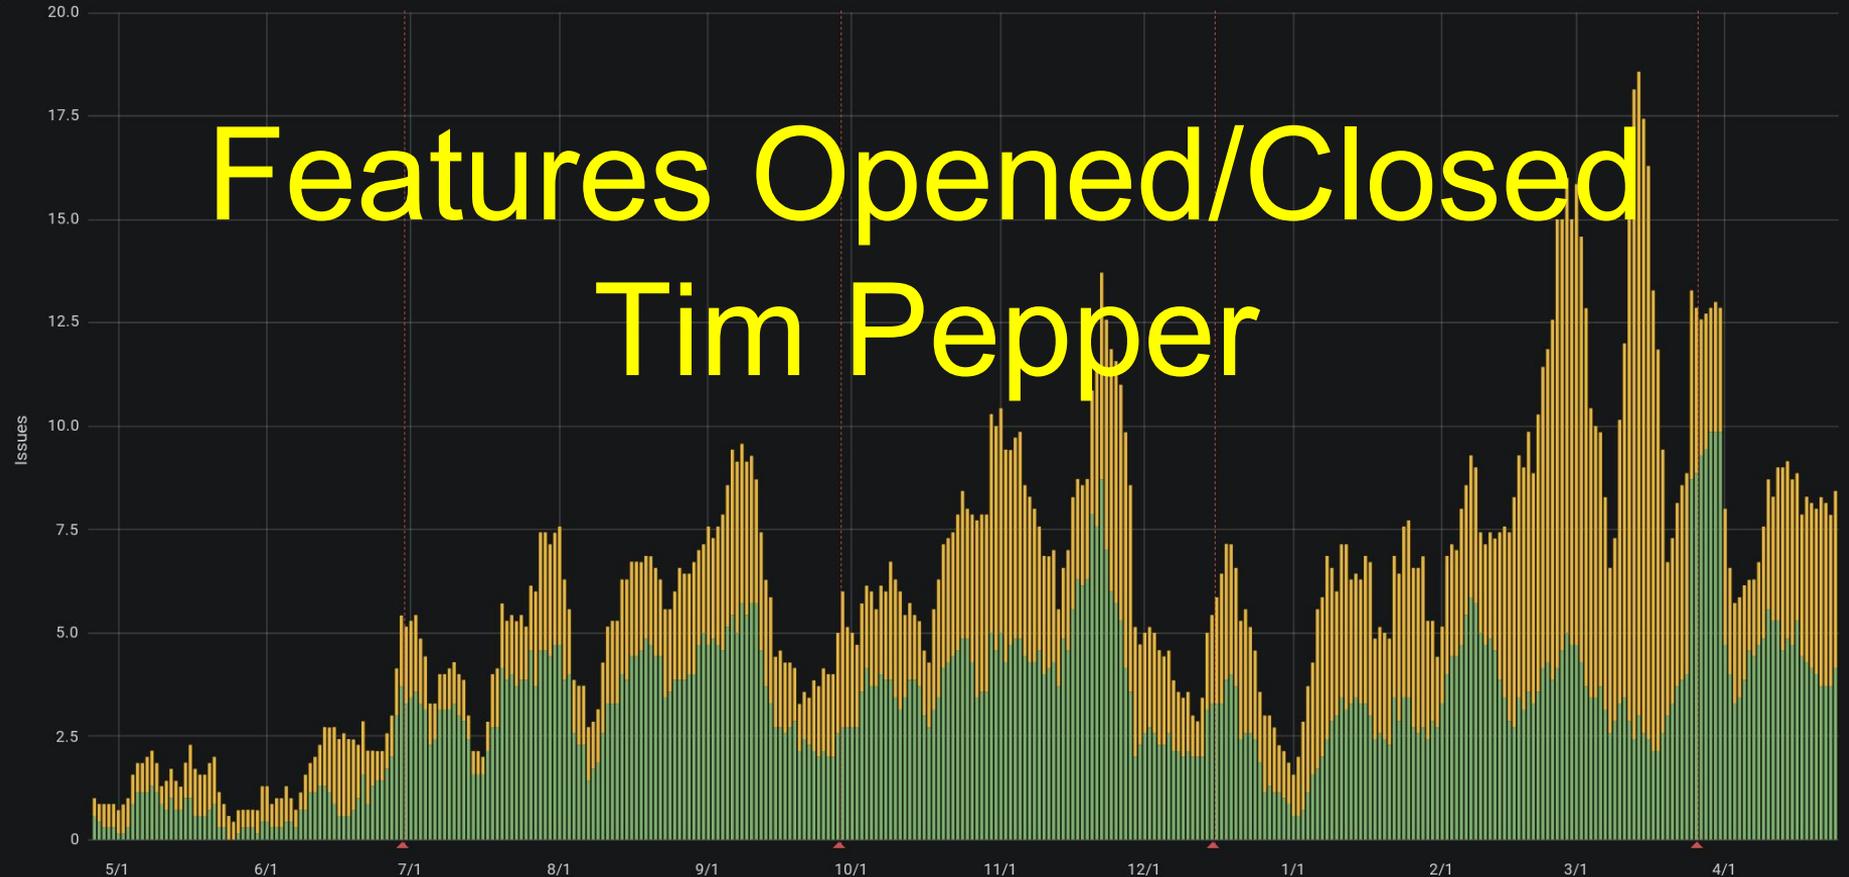


	min	max	avg	current	total
/release-note-none	1	7	3	1	252
/approve no-issue	0	9	2	1	205

Bot Commands by Repository Group Aaron Crickenberger

i

Issues SIG: All Kind: feature (7 Days MA)



Features Opened/Closed

Tim Pepper

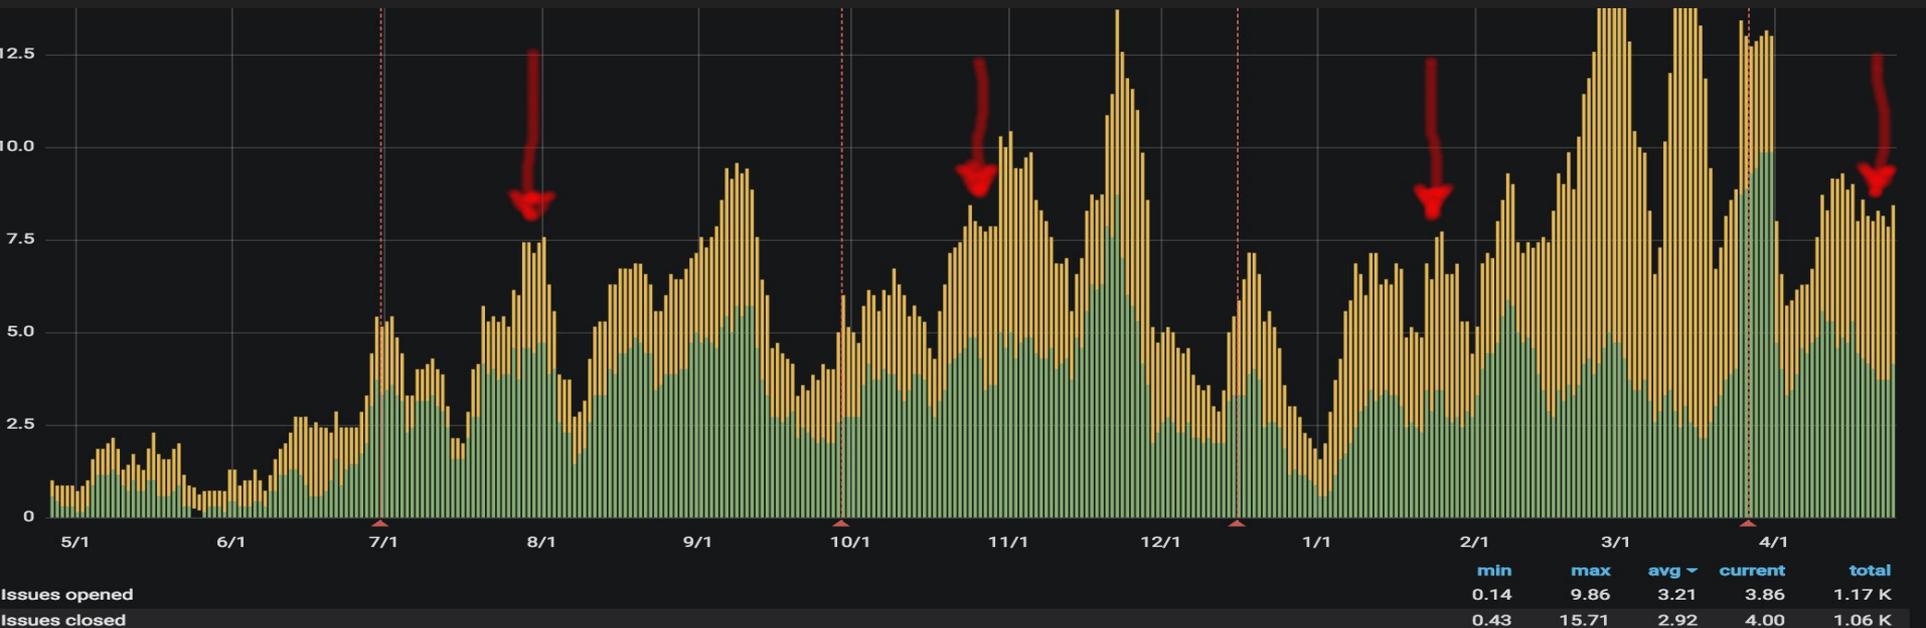
https://k8s.devstats.cncf.io/d/39/sig-issues?orgId=1&var-period=d7&var-sig=All&var-kind=feature&var-full_name=Kubernetes&from=1493078466941&to=1524614466941

Issues (specifically Feature type) Opened/Closed by all SIGs

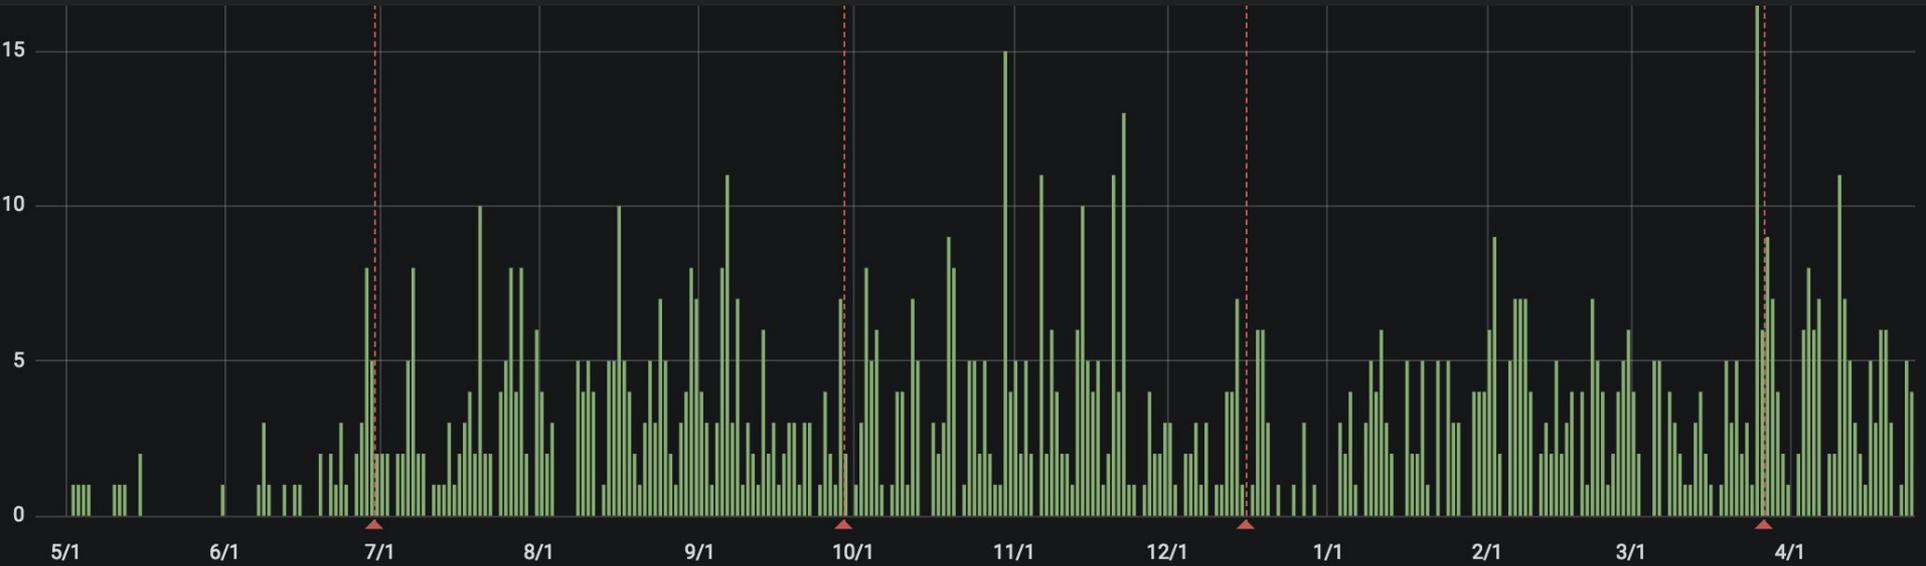
Open and Close (7day moving average)

red arrows are feature freeze for:

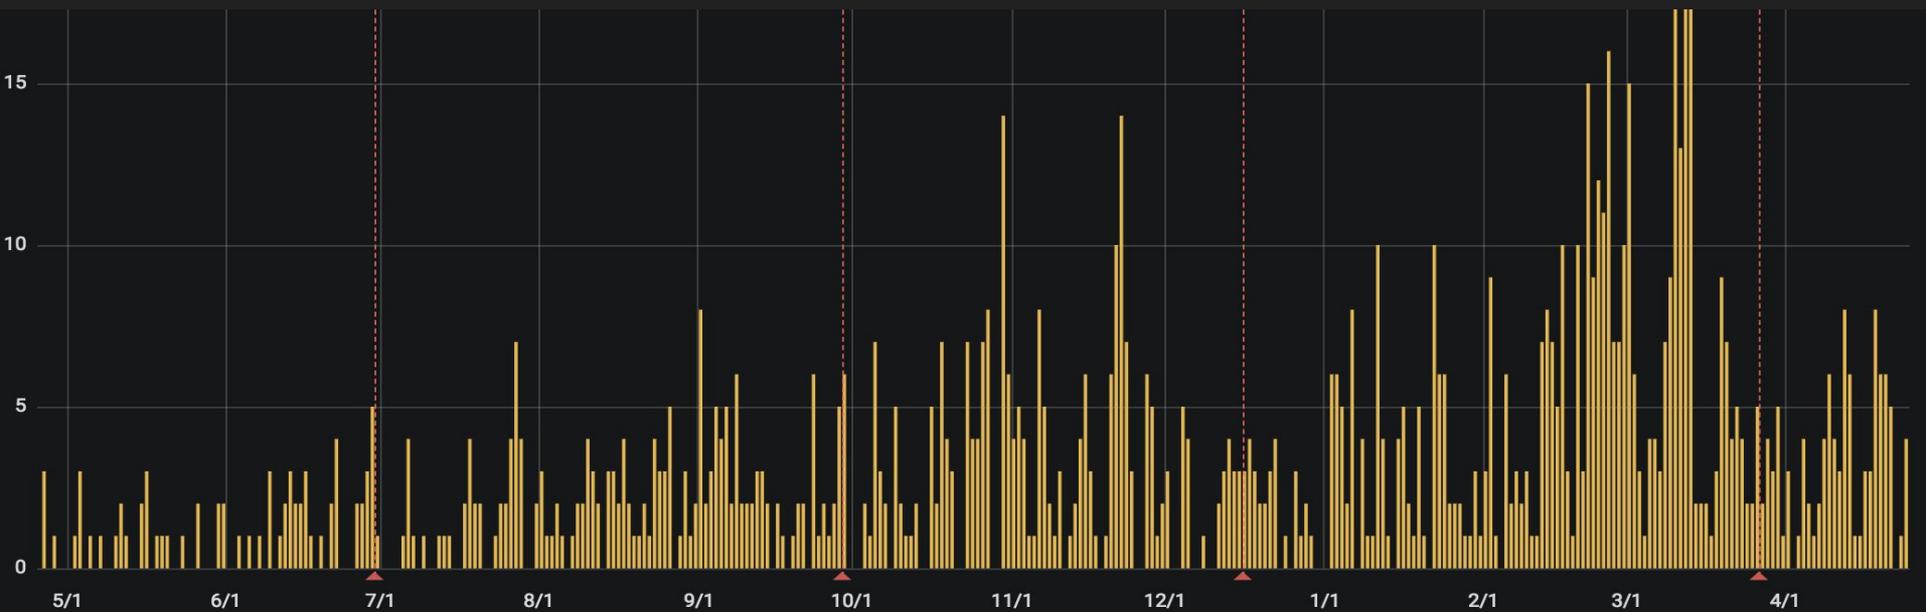
1.8 (8/1), 1.9 (11/22), 1.10 (1/22), 1.11 (4/24)



Just Opens (by day)

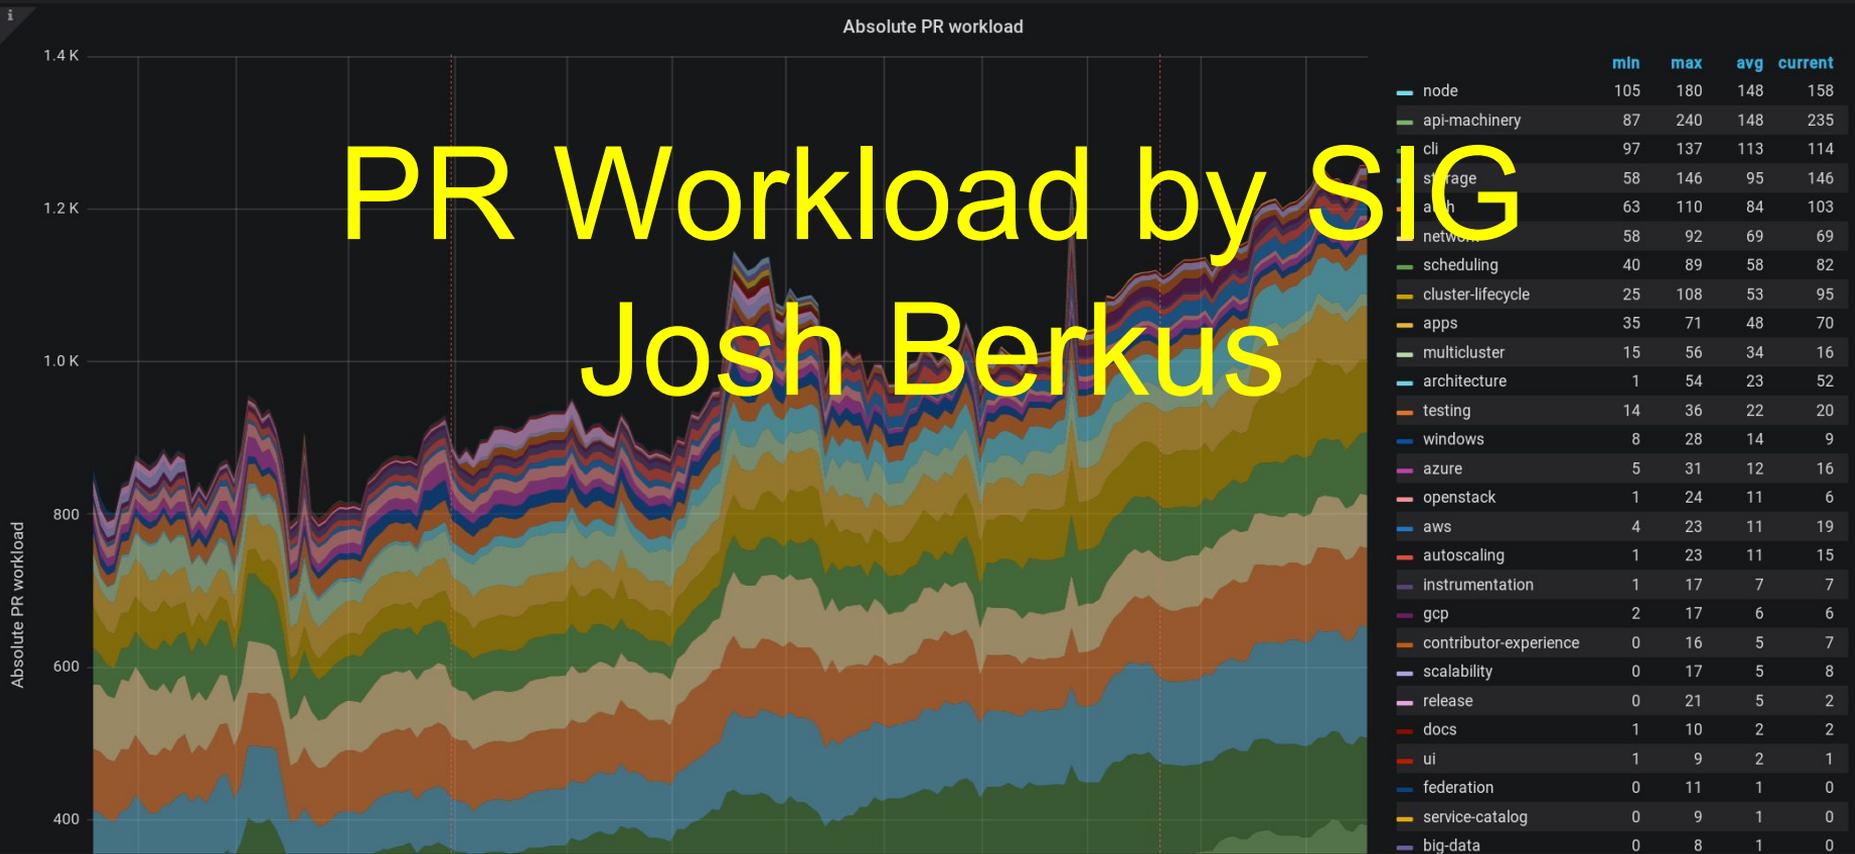


Just Close (by day)



PR Workload by SIG

Josh Berkus



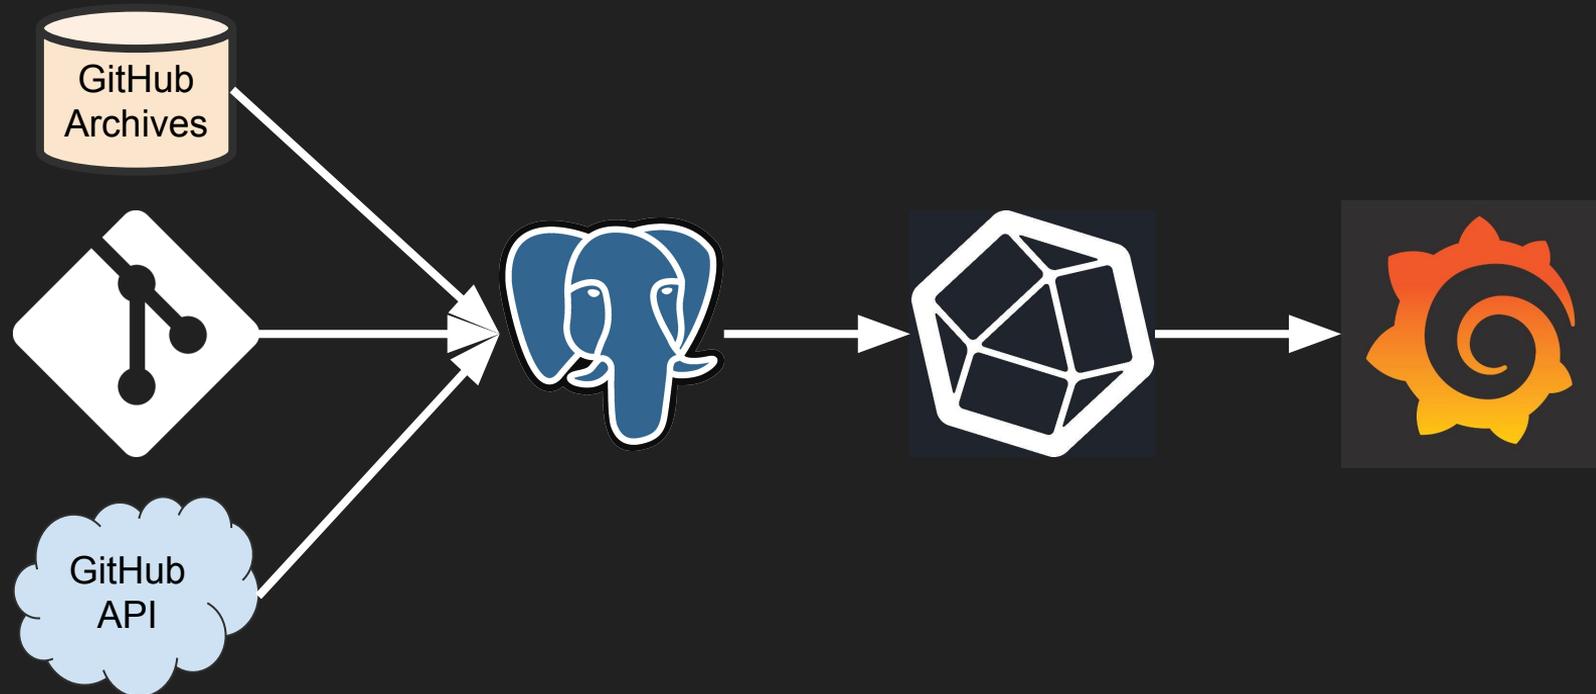
DevStats the Project

Project info

- On GitHub: [GH/cncf/devstats](#)
- Owned by the CNCF
- Apache license
- Mostly Go

#devstats

Data flow: big picture



Sources: Github Archives



Full historical archives for every event on Github as compressed JSON.

Free, with no rate limiting.

www.githubarchive.org

#devstats

Sources: Github Archives



1. Pull update bundles every hour
2. Filter out all non-Kubernetes activity
3. Sync with Postgres database

#devstats

Sources: Git



- Clone all repos and pull every hour
- Get commit analysis of modified files
- Define “repository groups”

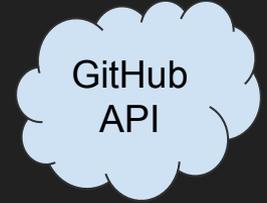
“repository groups”

Certain repos (like Kubeadm)

+

Certain directories (like cmd/kubeadm)

Sources: Github API

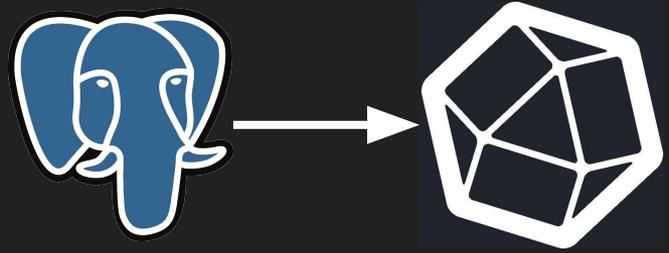


Just for correcting labels/milestones

Otherwise rate limiting is too slow

#devstats

Databases



PostgreSQL: holds copy of all GH data & events.

InfluxDB: stores time-series dimensional data for Grafana.

Grafana



Displays charts based on JSON chart definitions.

Now with batch JSON import for new charts!

#devstats

Contributing to the project

GH/cncf/devstats

Issues, PRs for contributions

Not for new chart requests.

#devstats

How do I make dashboard requests?

The Request Queue

The screenshot shows the GitHub repository page for `cncf / devstats`. The repository has 11 watches, 72 stars, and 28 forks. The navigation bar includes links for Code, Issues (8), Pull requests (0), Projects (1), Wiki, Insights, and Settings. The main content area is titled "Kubernetes DevStats Request Queue" and was updated 10 days ago. It features a Kanban board with four columns: "Proposals" (1 card), "Denied" (0 cards), "Approved" (0 cards), and "In progress" (0 cards). A search bar for "Filter cards" and buttons for "Add cards", "Fullscreen", and "Menu" are also present.

`cncf / devstats` Watch 11 Star 72 Fork 28

<> Code Issues 8 Pull requests 0 Projects 1 Wiki Insights Settings

Kubernetes DevStats Request Queue

Updated 10 days ago

Filter cards + Add cards Fullscreen Menu

- 1 Proposals**
 - Elsie Phillips
ephillip@redhat.com
Please drop the following graphs from the dashboard:
 - Activity repository groups
 - Approvers repository groups
 - PRs Approval repository groups stacked
 - PRs Authors Companies table
 - SIG mentions categories
 - SIG mentions using labels
 - Suggested approvers by repository groupsThey have been reviewed by the DevStats subproject committee and have been judged to be lacking in value for the

- 0 Denied**
- 0 Approved**
- 0 In progress**

Steps for filing a request

1. Go to <https://github.com/cncf/devstats/projects>
2. Follow the directions in the description
 - a. Create a card in the “Proposal” column of the github project with a thorough description of your request
 - i. Your name
 - ii. Email address
 - iii. The name of the dashboard you wish changed (if applicable)
 - iv. Description of your proposed change (change to an existing dashboard or proposal for a new dashboard) -
 - v. Rationale for new dashboard/changed dashboard (if applicable)
 - b. Wait for the committee to review your proposal (this happens generally once a week on Thursdays)
 - c. Check back on the status of your request.
 - i. If it is approved the card will be moved into the approved column, if not it will be moved into the declined column with a note as to why it was declined.
3. Finito!

Caveats/quirks

- Company attribution from cncf/gitdm
 - Generally out of date
 - Some devs not attributed
- Data in revision
 - Occasional retroactive blanket updates to fix double-counting, other bugs

#devstats

Questions?

More Stats

Devstats meetings
Thursdays 2130 UTC

ContribEx Deep Dive
Thursday 15:50pm

Channel #devstats on
Kubernetes Slack

Sig-contributor-experience
mailing list

#devstats

