KubeCon | CloudNativeCon

North America 2026

AI-DDOS



How AI-generated contributions are breaching OSS communities, and what we can do about it

Zixuan Steve Feng

Assistant Professor,
NOODLE Lab,
Virginia Commonwealth University

Salt Lake City, November 2026

Read more



KubeCon CloudNativeCon North America 2026

NOODLE LAB @ VCU

Navigating Open-source, Organization, Developer, and Learning Ecosystems

About me and my lab

Carnegie Mellon University

Oregon State University

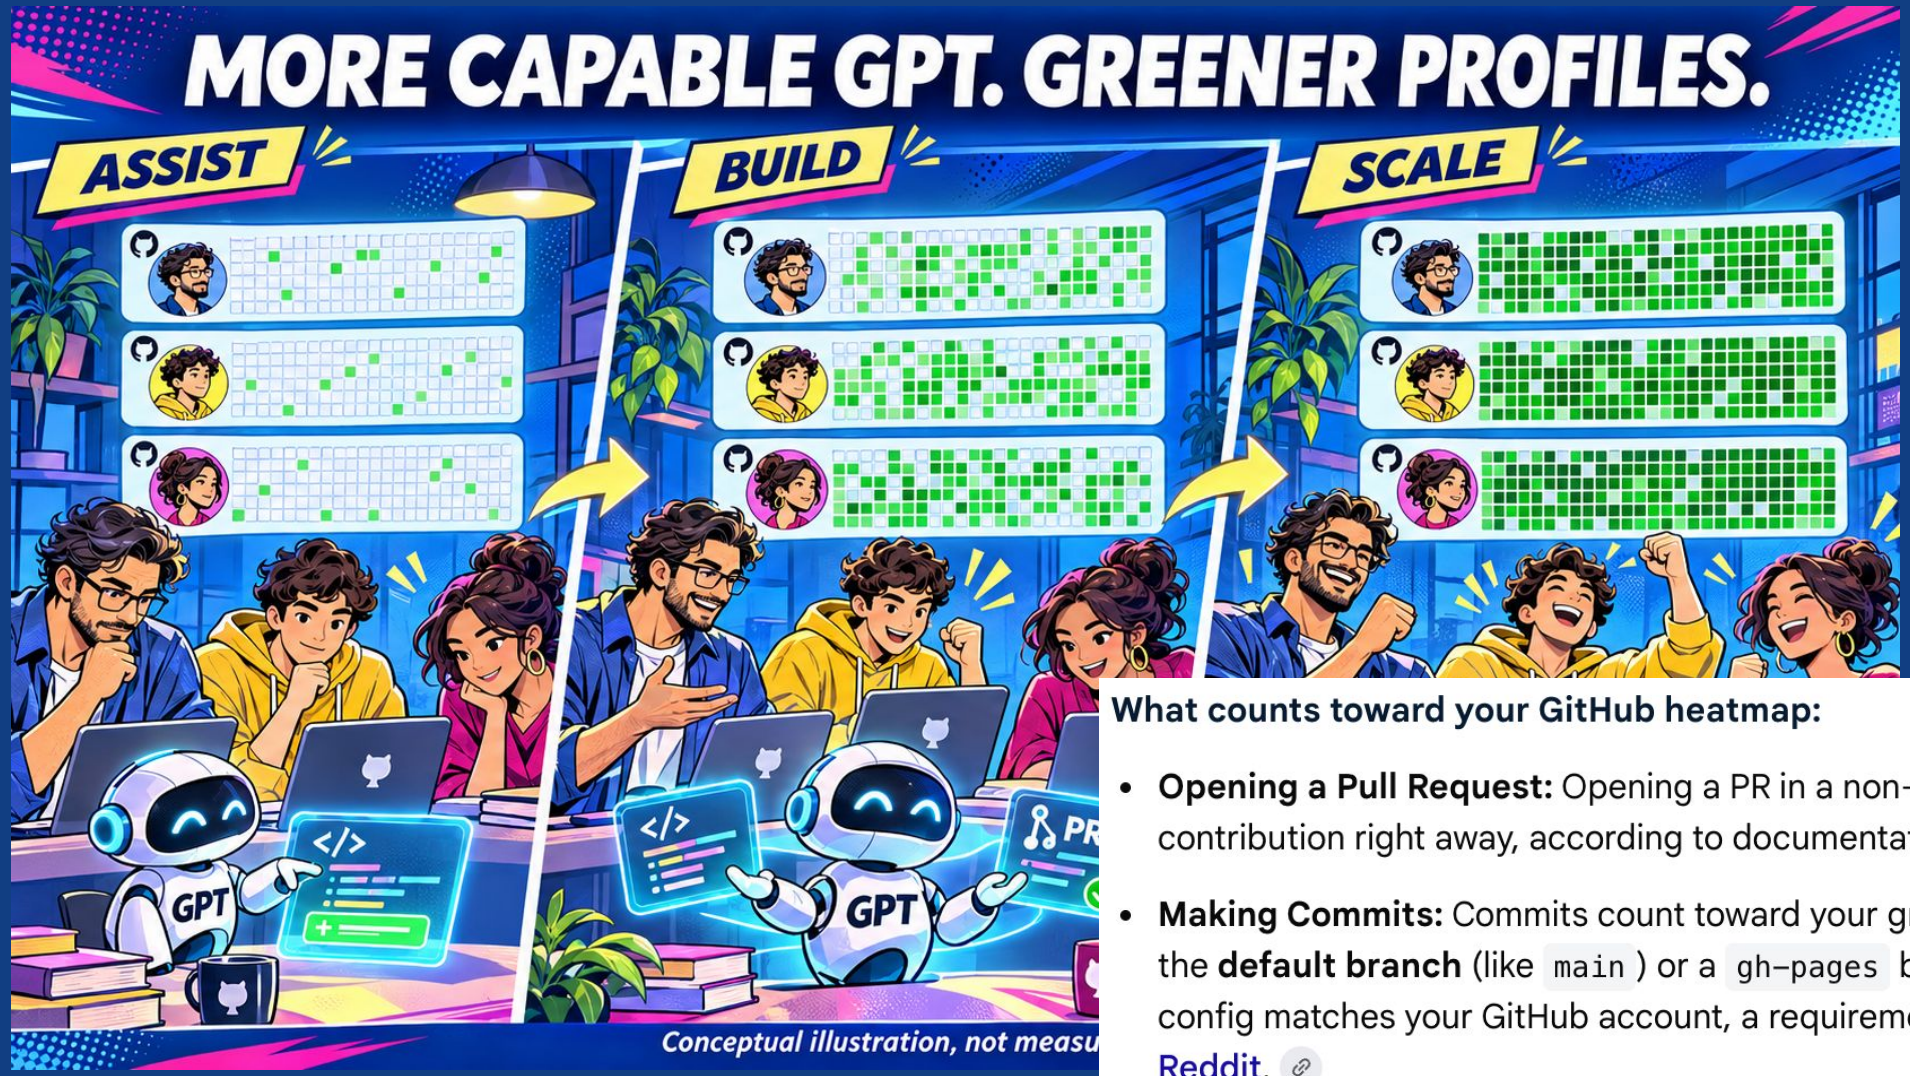
UC Irvine

Northern Arizona University

Colorado State University

The George Washington University

Microsoft Google Sandia National Laboratories THE ASF CLOUD NATIVE COMPUTING FOUNDATION CodeDay



What counts toward your GitHub heatmap:

- **Opening a Pull Request:** Opening a PR in a non-fork repository counts as a contribution right away, according to documentation discussed on GitHub.
- **Making Commits:** Commits count toward your graph as long as they are pushed to the **default branch** (like `main`) or a `gh-pages` branch, and the email in your Git config matches your GitHub account, a requirement highlighted by developers on Reddit.
- **Other Activities:** Creating repositories, opening or closing issues, submitting PR reviews, and participating in discussions also turn squares green. GitHub Docs +2

A COMMUNITY CAN CLOSE TO SURVIVE

- curl ended its six-year bug bounty in January 2026 after AI-generated submissions overwhelmed triage.
- Ghostty requires AI-use disclosure and human understanding of submissions, and can ban contributors who submit low-quality AI-generated work.
- GitHub lets maintainers restrict pull requests to collaborators or disable them entirely.
- …

MAINTAINER ATTENTION IS THE BOTTLENECK



“The deluge of PRs is just wasting our time” [Mailing List]. Even “rejecting them takes more time than just writing the code” [Blog].

WHAT WE MEAN BY AI-DDOS

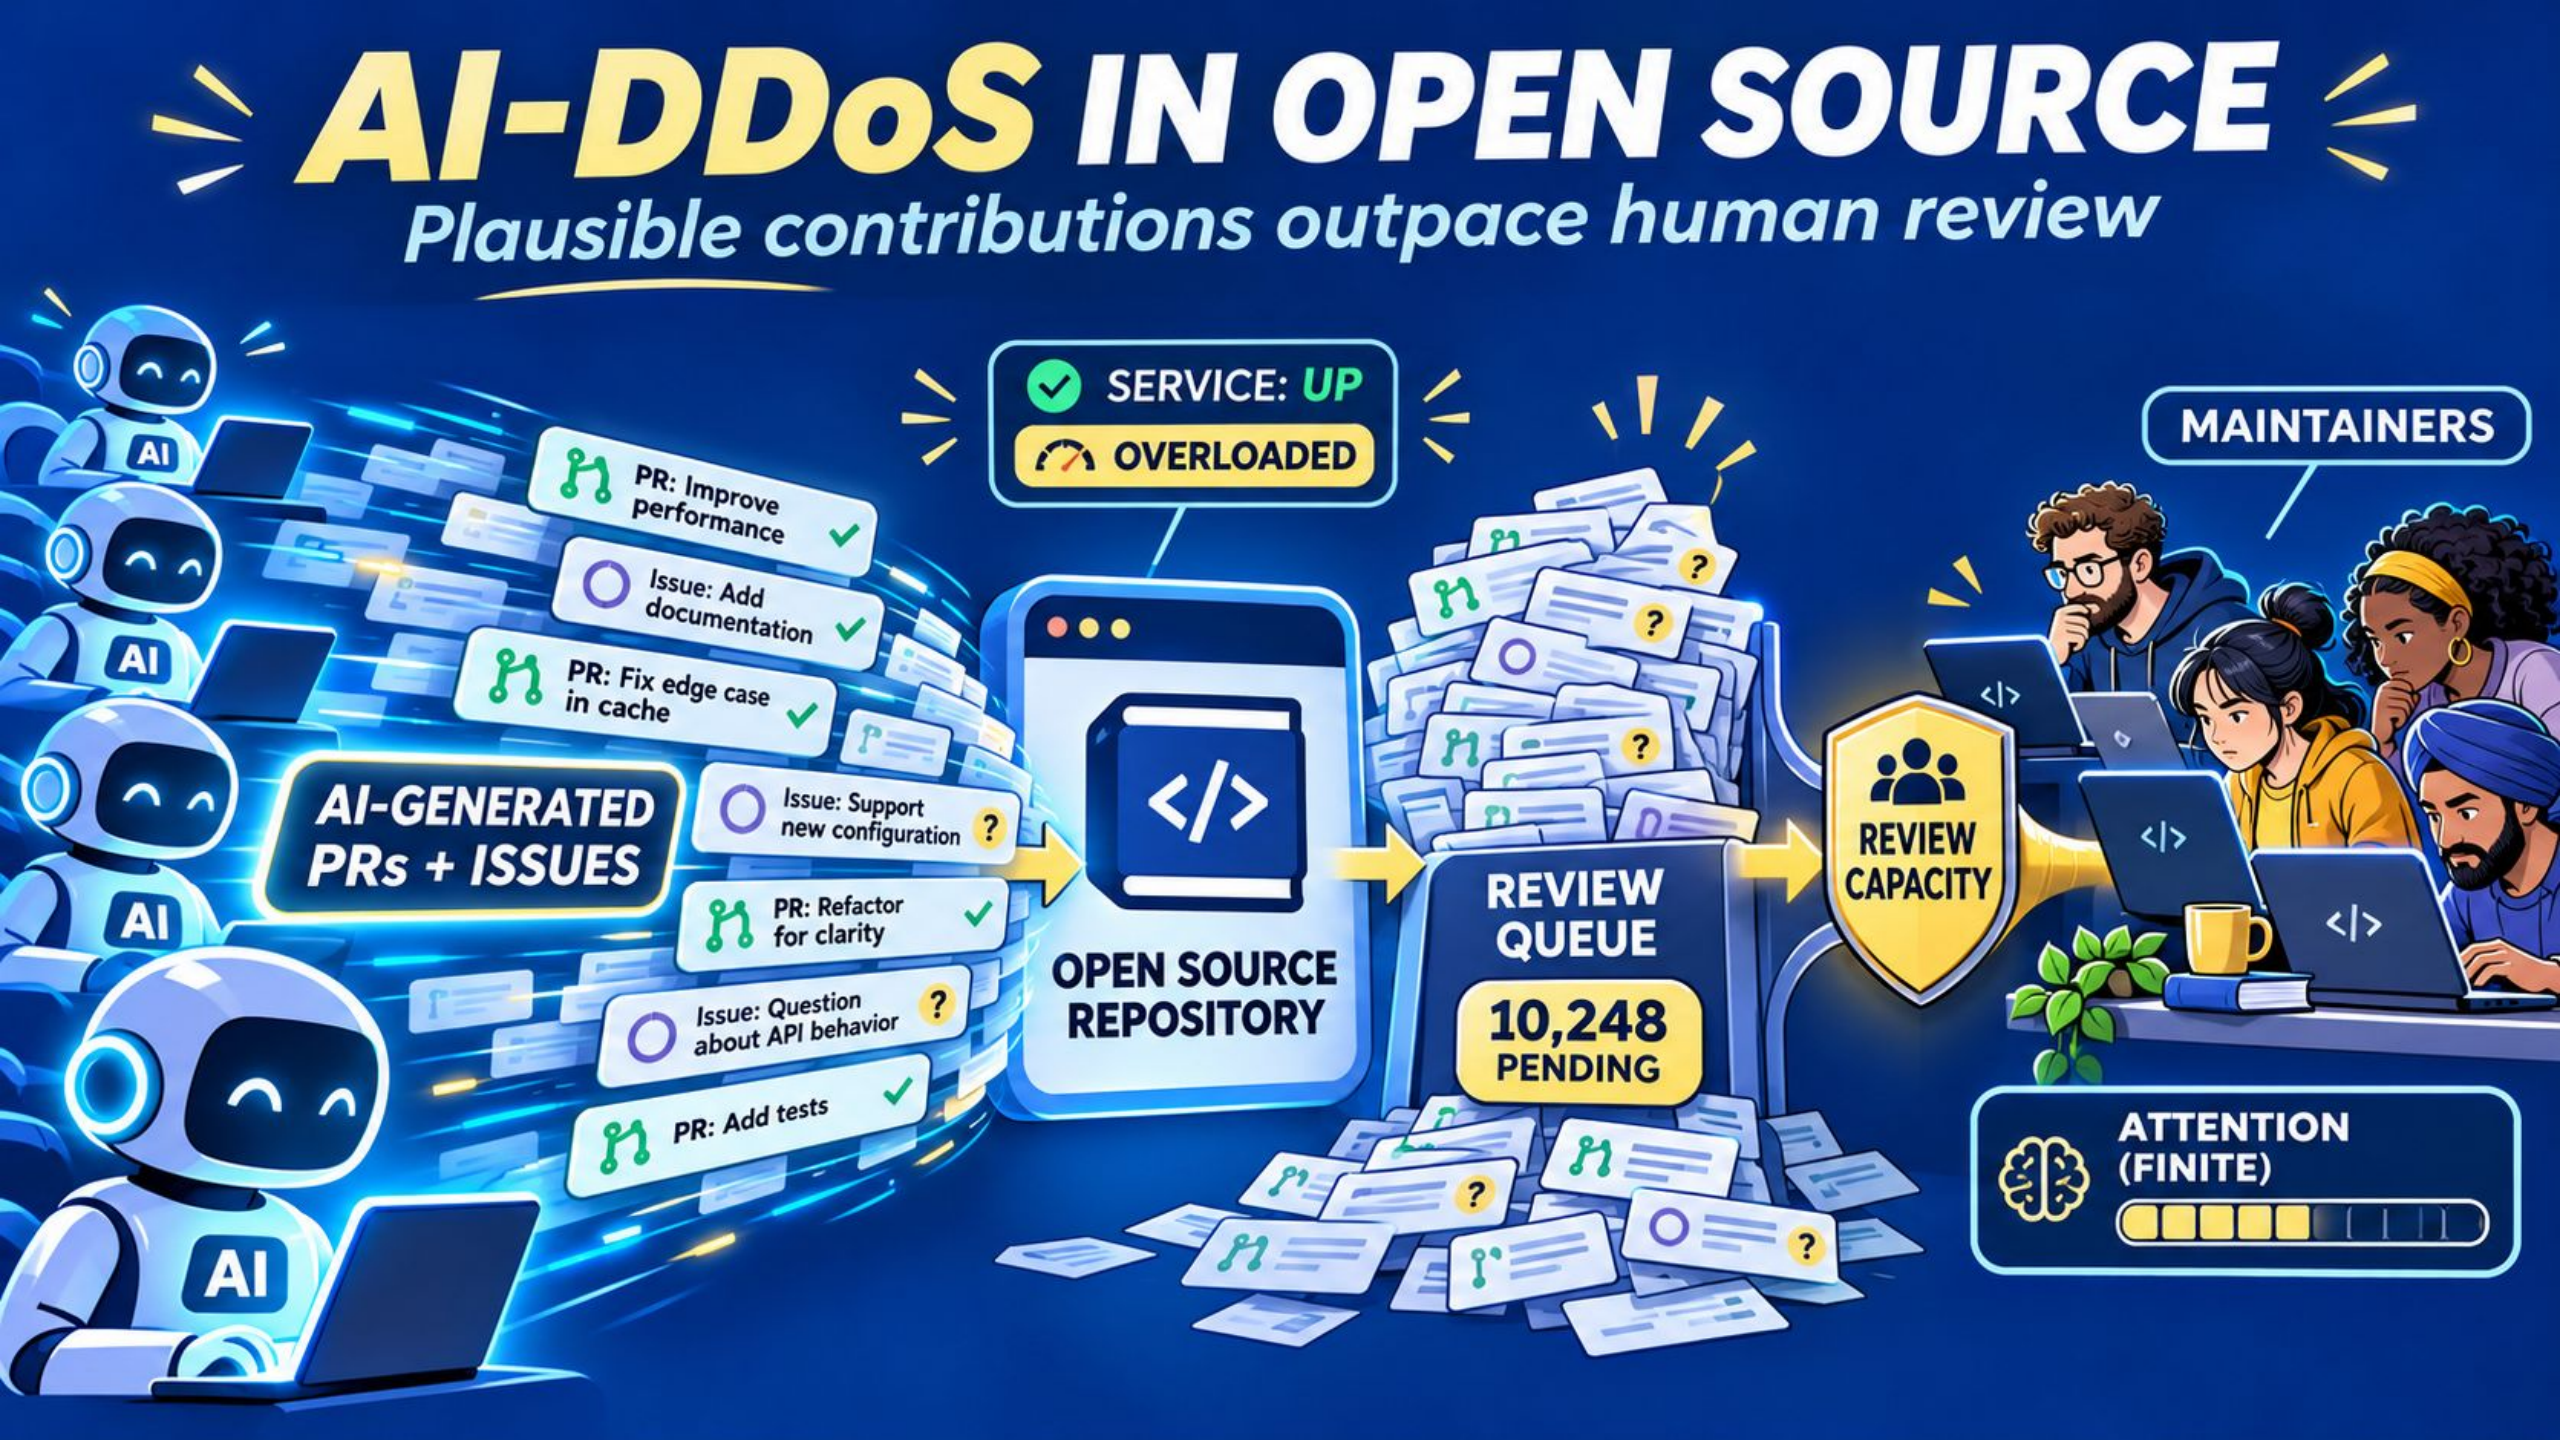
In computing, a **denial-of-service attack** (DoS attack /dɒs/ doss) is a cyberattack in which the perpetrator seeks to make a machine or network resource unavailable to its intended users by **temporarily or indefinitely disrupting services** of a host connected to a network [1].

[1] NIST. “Denial of service (DoS).” Computer Security Resource Center Glossary. https://csrc.nist.gov/glossary/term/denial_of_service

WHAT WE MEAN BY AI-DDOS

In OSS, AI-DDoS is a denial-of-service effect in which plausible AI-generated contributions flood open-source repositories, overwhelming communities' capacity to review them.

AI assistance alone does not make a contribution slop.


AI-DDoS IN OPEN SOURCE
Plausible contributions outpace human review
AI
AI
AI
AI
AI-GENERATED PRs + ISSUES
PR: Improve performance
Issue: Add documentation
PR: Fix edge case in cache
Issue: Support new configuration
PR: Refactor for clarity
Issue: Question about API behavior
PR: Add tests
SERVICE: UP
OVERLOADED
OPEN SOURCE REPOSITORY
REVIEW QUEUE
10,248 PENDING
REVIEW CAPACITY
MAINTAINERS
ATTENTION (FINITE)

HOW WE STUDIED AI-DDOS in OSS

KubeCon CloudNativeCon North America 2026

A phenomenon-based mixed-methods study

CHARACTERIZE THE PROBLEM



Practitioner accounts

blog posts
Reddit threads
mentor-list threads

Reflexive thematic analysis

6 themes + hypotheses



Repository traces

294 repositories

2.05M PRs + issues

Weekly BSTS analysis

Train: 2023–2024
Compare: 2025

Informs follow-up

IDENTIFY COMMUNITY RESPONSES



Maintainer interviews

11 maintainers

Organizational Resilience Theory

11 remediation strategies

7 member-ch



Practitioner survey

229 valid responses

Each of 11 strategies rated

Likelihood of adoption

-point scale

1 SIGNALS

WHAT PRACTITIONERS REPORTED

KubeCon CloudNativeCon North America 2026

Six themes recur across all three source types

Theme	What practitioners described	Reddit	Blogs	Mailing List	Total
T1. Volume and submission flood	AI-generated contributions exceed review capacity.	16	20	41	**77**
T2. Quality collapse of submissions	Submissions look plausible but fail quality expectations.	11	7	17	**35**
T3. Transient contributor	Contributors mass-generate drive-by submissions.	8	4	10	**22**
T4. Incentive and platform structure	Bounties, heatmaps and pay-to-play structures magnify submissions.	9	9	1	**19**
T5. Maintainer and mentor wellbeing	Review overload drives burnout and withdrawal.	4	6	9	**19**
T6. Defensive contraction	Communities close channels and restrict newcomer access.	7	15	11	**33**

THE REPOSITORY DATASET

KubeCon CloudNativeCon North America 2026

Activity window: January 2023–December 2025

"Markets are tough, so there's incentive to farm GitHub contributions for resumes" [Reddit]

01 **Curated starting pool**
goodfirstissue.dev, accessed March 2026

02 **Popularity threshold**
At least 5,000 GitHub stars yielded 355 candidates

03 **Sustained contribution traffic**
≥100 contributions/year, no inactivity gap >6 months

04 **Observation coverage**
Created no later than 2023, continuously active through 2025

05 **Available contribution channels**
Nonarchived, with issues and pull requests enabled

294

repositories

after all five criteria

Including Kubernetes and Grafana

GitHub REST API
Creation time, contributor ID and outcome status

1,224,374
pull requests

822,133
issues

356,972
unique contributors

Excluded accounts with bot-name markers: [bot], -bot, _bot

HYPOTHESES FOR CONTRIBUTION ACTIVITY BECAUSE OF AI:

H1

All contributors

H1a PR volume ↑
H1b PR merge ratio ↓
H1c Issue volume ↑
H1d Issue completion ratio ↓

H2

Recurring contributors

H2a PR volume ↑
H2b PR merge ratio ↓
H2c Issue volume ↑
H2d Issue completion ratio ↓

H3

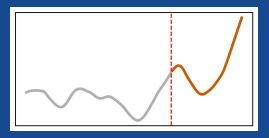
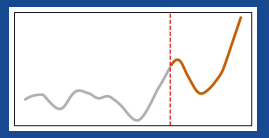
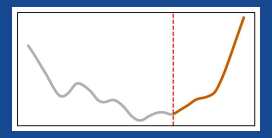
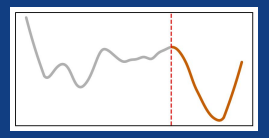
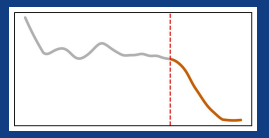
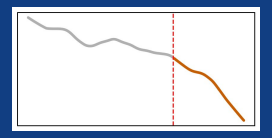
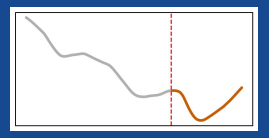
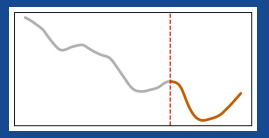
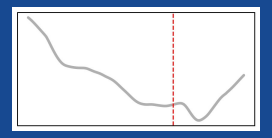
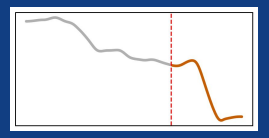
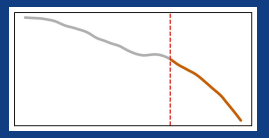
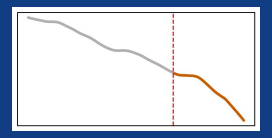
One-time contributors

H3a PR volume
H3b PR merge ratio
H3c Issue volume
H3d Issue completion ratio

CONTRIBUTION TRENDS IN 2025

KubeCon | CloudNativeCon North America 2026

Mean weekly differences versus the modeled counterfactual

Metric	All contributors	Recurring contributors	One-time contributors
PR volume	 **+519.86/wk** +6.80% relative P 99.9% |d| 0.69	 **+497.30/wk** +6.91% relative P 99.9% |d| 0.69	 **+24.67/wk** +5.84% relative P 99.4% |d| 0.43
PR merge ratio	 **≈−1 pp** −1.06% relative P 99.9% |d| 0.68	 **≈−3 pp** −3.45% relative P 99.9% |d| 2.14	 **≈−9 pp** −18.18% relative P 99.1% |d| 2.84
Issue volume	 **−431.15/wk** −8.12% relative P 99.1% |d| 0.65	 **−453.11/wk** −12.40% relative P 99.9% |d| 0.97	 **+58.20/wk†** +4.00% relative P 82.8% |d| 0.27
Issue completion ratio	 **≈−3 pp** −3.40% relative P 99.1% |d| 1.57	 **≈−5 pp** −8.09% relative P 99.9% |d| 1.60	 **≈−5 pp** −10.30% relative P 99.9% |d| 1.22

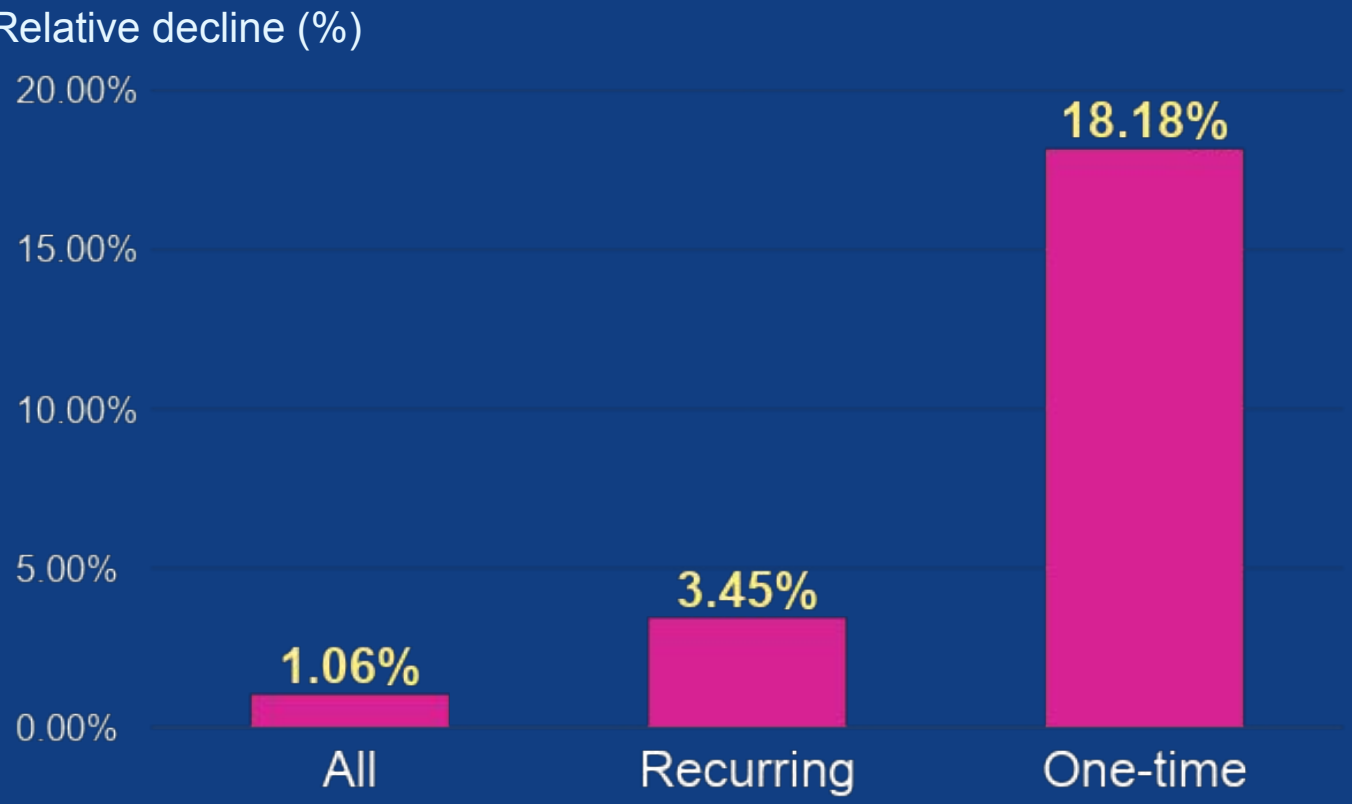
Original plots: gray before 2025 or nonsignificant, orange after 2025 when significant. Red dash: January 2025.

pp = percentage points, approximated from rounded ratios. P = posterior probability. |d| = absolute Cohen's d.

† Below the paper's 95% threshold. Issue results require caution because of significant placebo effects.

ONE-TIME CONTRIBUTORS SAW THE SHARPEST DECLINE

merge rate



18.18%

relative decline

"One-time" means one contribution to that repository in the observation period.

THREE WAYS COMMUNITIES RESPOND (11 interviews, 229 survey respondents, 11 responded strategies)

PRESERVATIVE

Protect current capacity

CI quality gates
No-review cutoffs

ADAPTIVE

Change the workflow

AI-use disclosure
Shared triage

TRANSFORMATIVE

Rethink participation

Demonstrated understanding
Review intent before code

11 RESPONSES TO AI-DDOS

KubeCon | CloudNativeCon North America 2026

Likely or very likely to adopt (N = 229). Effort classifications come from interviews.



ABSORB THE PRESSURE

Preservative

Strategy	Effort	Likely
SP1 CI quality gates	Low	73%
SP2 No-review cutoffs	Low	71%

Low-effort defenses drew the strongest interest.



ADAPT THE WORKFLOW

Adaptive

Strategy	Effort	Likely
SA1 AI-use disclosure	Low	63%
SA2 Pre-review gates	Medium	62%
SA3 Behavioral risk flags	Medium	58%
SA4 AI review filters	High	48%
SA5 Rotate triage + community review	High	34%



RETHINK CONTRIBUTION TRUST

Transformative

Strategy	Effort	Likely
ST1 Verify understanding	Medium	67%
ST2 Align AI policy with project mission	Medium	60%
ST3 Review intent and plans	High	43%
ST4 Verify identity or get endorsement	High	39%

Adoption intentions, not measured effectiveness. Source: Afroz et al. (2026), Table III arxiv.org/abs/2607.04003

RESILIENCE UNDER AI-DDOS

Keeping OSS open while protecting maintainer capacity

KubeCon CloudNativeCon North America 2026

RESILIENCE FOUNDATIONS

Resilient behavior

Stay open under load vs. avoidance

Resilience resources

Maintainer attention, engagement, wellbeing

Resilience capabilities

Anticipate, make sense and coordinate

Resilient response

How can OSS stay open under AI-DDoS?

OPEN UNDER LOAD

Adaptive & transformative remediation

Proof-of-engagement

Mitigation upstream

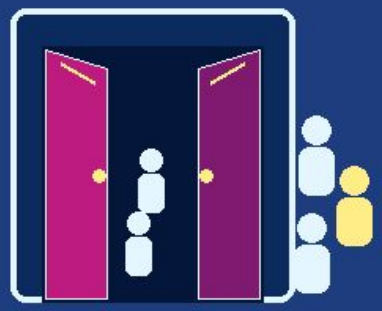
Organizational growth

DEFENSIVE CLOSURE

Rigidity trap

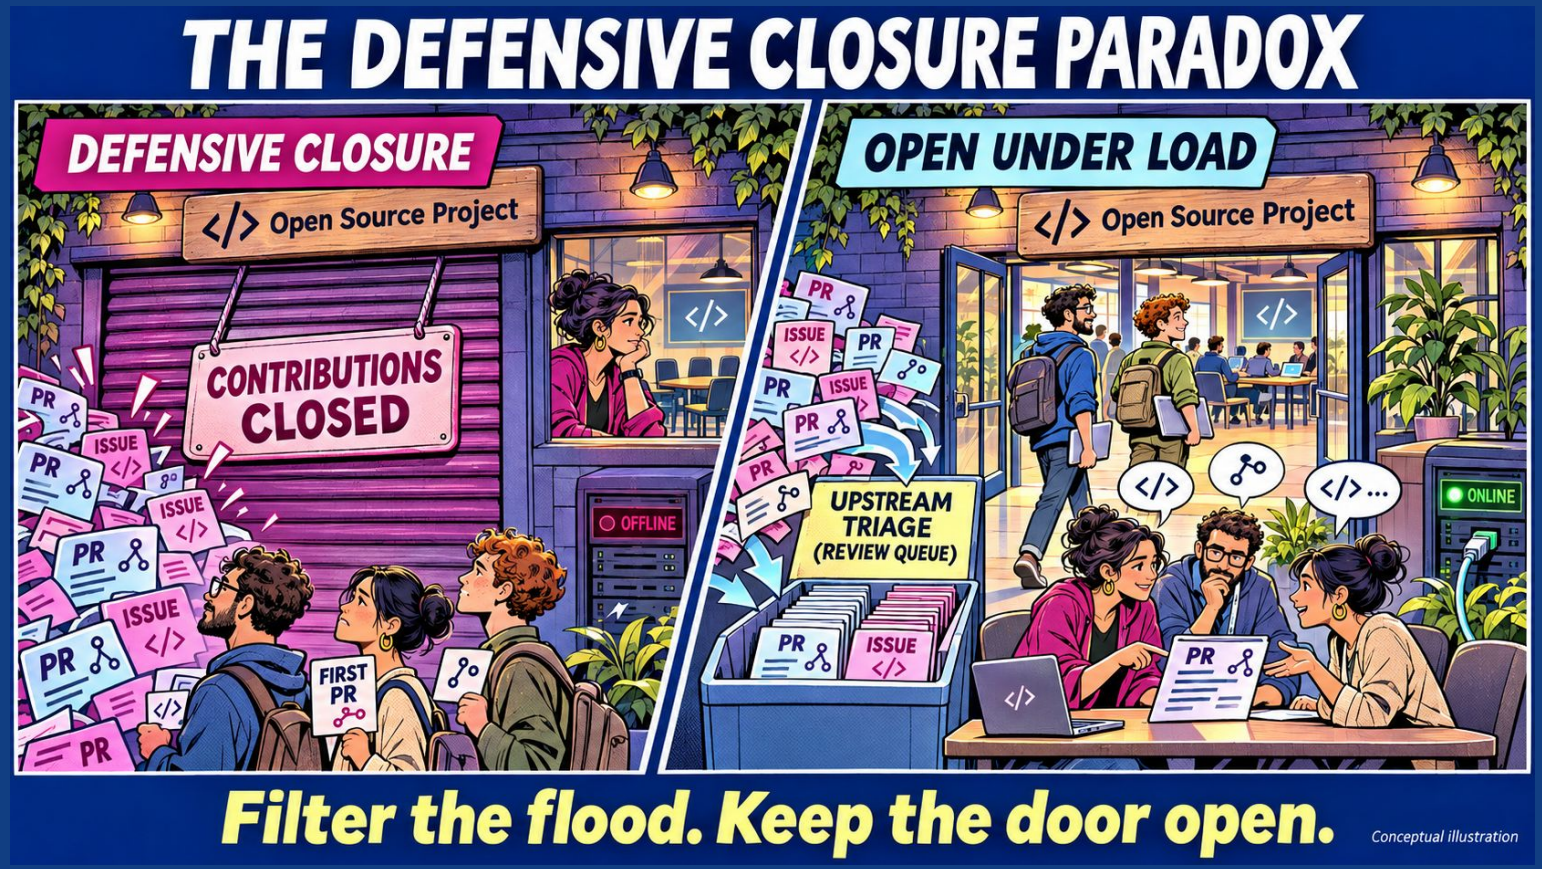
The defensive closure paradox

THE OSS DOOR



THE SUSTAINABILITY TRAP

If we call this contribution flood “AI-DDoS,” we should remember the lesson from security research: do not unplug the service. Yet our observations suggest that we are getting close to closing the door.



AI REVIEWING AI LEAVES A POLICY DECISION

What should the model flag?

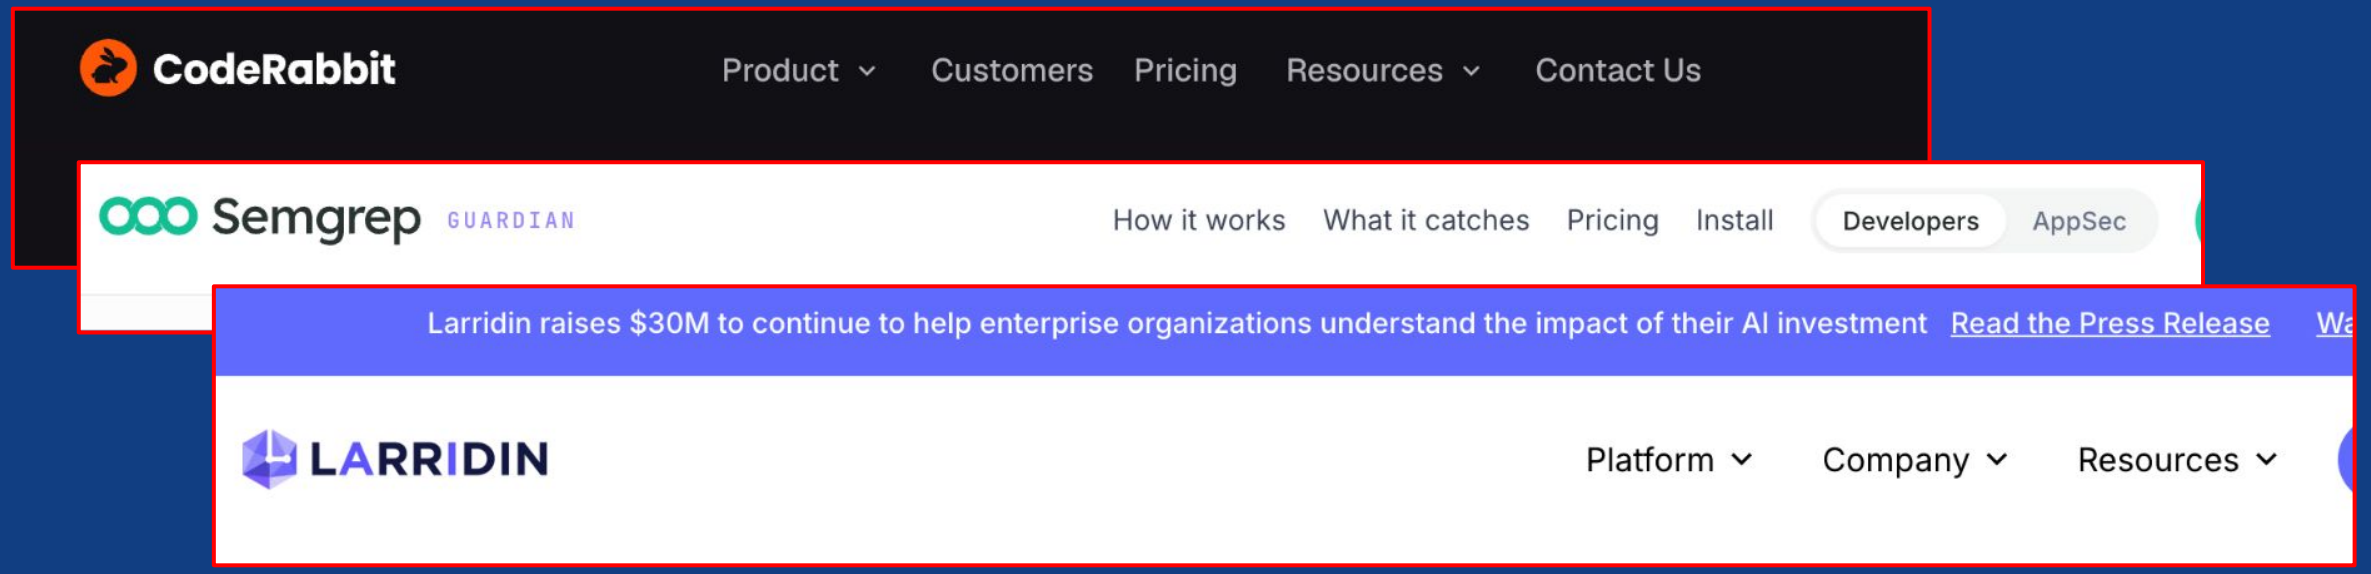
What should it be allowed to reject?

Who can question the decision?

A confidence score does not settle the community’s responsibility.

Speaker perspective, informed by Afroz et al., 2026, Section VI. AI screening is a proposed response, not a proven solution.

EXPLICIT CHECKS AND HUMAN JUDGMENT



Do we really want to use Bot to police human?

RANDOM-FOREST FEATURE COMPARISON

"Some people are trying to build up realistic looking GitHub profiles so that they can do supply line attacks" [R8]

Model	Accuracy	Precision	Recall
History alone	**85.2%**	**80.9%**	**92.3%**
Full text + metadata	81.6%	79.2%	85.7%
Text + metadata without history	59.3%	57.9%	68.1%
Embeddings alone	56.0%	56.6%	51.5%
Embeddings + metadata without history	59.1%	58.0%	66.5%
Embeddings + full metadata	81.4%	79.6%	84.3%

Do we really want to use Bot to police human?

Videos



How One Founder Manages 20 AI Coding Agents at Once

YouTube · AI Engineer
21 hours ago

Learn how Untangle uses manager **agents** to oversee up to 20 builder **agents** concurrently, automating tasks & streamlining the **development** process.<ctrl100>

enthusiasts, and aspiring coders. Whether you're exploring the depths of Python, diving into data science, mastering web development, or staying updated on the latest trends in AI, Coding Nexus has something for you.

Follow publication

They Replaced 30 Engineers With 30 AI Agents and Bu… Codebase in 12 Days

Tattva Tarang Follow 5 min read · Mar 5, 2026



Bargan Constantin
Posted on Oct 22 · Edited on Oct 22 · AI-assisted

❤️ 4 🦄 2 🤯 2 🙌 2 🔥 2

Anthropic Runs 30,000 Agents at Once. What Keeps Them Under Control?

Mentorship, Interrupted: How GenAI Disrupts the GSoC Pipeline and Why Mentors Matter More
R Xing, Z Feng, M Radomile, S Taylor, S Vargas, A Sarma
arXiv preprint arXiv:2610.38167

The EPIC Framework for Spec-Driven Development
B Claudino, R Xing, Y Zuo, T Menezes, A Sarma, K Damevski, Z Feng
arXiv preprint arXiv:2610.07534

Technostress in the Age of AI: A Preliminary Study with Software Professionals
RS Santos, I Santos, C Magalhaes, Z Feng
arXiv preprint arXiv:2609.15686

"You Can't Open an LLM With a Screwdriver": The De-Democratization of Software
Z Feng, I Santos, K Damevski, A Sarma
arXiv preprint arXiv:2608.24720

Charting uncertain waters: A socio-technical roadmap for sustaining open source communities in the age of genai
Z Feng, R Milewicz, E Murphy-Hill, T Menezes, A Serebrenik, ...
ACM Transactions on Software Engineering and Methodology 35 (8), 1-21

" AI Slop is DDoSing Open Source": Understanding the Impact of AI-Generated Contributions on Open Source Sustainability
S Afroz, C Miller, T Menezes, E Gilmour, A Sarma, Z Feng
arXiv preprint arXiv:2607.04003



Thank you!

About me and my lab

KubeCon | CloudNativeCon

North America 2026

