PRODUCT TEARDOWN

GitHub product teardown

When AI makes pull requests free to write, who pays the cost of reviewing them?

Existing product · Lens: maintainer review load · Data: Octoverse 2025 and 2026 press reports

01 · PRODUCT

An existing platform at huge scale

GitHub hosts code and the pull request (PR) workflow: propose a change, get it reviewed, merge it. Microsoft has owned it since 2018.

Why GitHub

The default home of open source, so changes in contribution behaviour show up here first.

Why the maintainer lens

GitHub publicly acknowledged low-quality PR volume in 2026 and is weighing new controls.

WHY IT MATTERS

518.7M

pull requests merged in a year, up 29%.

Writing code is getting cheaper with AI. Review is the part that has not.

02 · SCALE

A huge and growing developer base

180M+
developers on GitHub

36M+
new developers in the year, up 23%

630M
repositories, 395M of them public

21.9M
developers in India, second largest country

~80%
of new developers use Copilot in week one

4.7M
paid Copilot subscribers, up 75% (Jan 2026)

Who competes, and how GitHub earns

Competitors

GitLab: FY2026 revenue $955.2M, up 26%

GitLab: 50M+ registered users

Codeberg: Gentoo is moving there over Copilot

Bitbucket: named, not measured

Revenue sources

Team and Enterprise subscriptions

Copilot plans, from $10 to $100 a month

$2B run rate in July 2024 (latest verified)

No current revenue figure found.

03 · PROBLEM

Writing a PR is free. Reviewing one is not.

Maintainers report a flood of low-quality AI-generated PRs, and today's main controls are blunt: limit PRs to collaborators, or turn them off.

HOW MIGHT WE

cut the time maintainers spend on PRs that will never merge, without shutting out real newcomers?

04 · USERS

Four people, one to serve first

SERVE FIRST · 19 OF 20

Overloaded Maintainer

Volunteer who reviews every incoming PR.

SECOND · 17 OF 20

AI-Assisted Newcomer

First-time contributor using Copilot.

THIRD · 15 OF 20

Enterprise Reviewer

Lead reviewing a flood of internal PRs.

LATER · 11 OF 20

Foundation Admin

Sets policy across many projects.

Scored on size, headroom, fit with the fix and strategic value (my judgement, to validate with interviews).

The journey breaks at triage

Stage	What happens	Feeling	Where it breaks
Notice	A surge of PR alerts	Uneasy	Cannot tell signal from noise
Triage	Opens each PR to judge it	Drained	No ranked queue
Review	Reads code, leaves feedback	Frustrated	No effort signal from author
Respond	Explains why a PR is closed	Repetitive	Same reply again and again
Decide	Merge, close or restrict	Torn	Only blunt controls
Aftermath	Considers burnout or leaving	Disengaged	Burnout risk (my inference)

Seven pain points, two to fix first

#	Pain point	Evidence	Score
1	Volume of low-quality PRs	Maintainer reports; GitHub acknowledges	17
3	Blunt controls	Collaborator-only or off are the main levers	17
4	Weak triage tooling	No ranked queue	15
2	No effort or AI signal	Reviewers cannot see how a PR was made	14
5	Low-quality report spam	curl closed its bounty over low-quality reports	14
6	Unread contribution rules	Newcomers skip templates	14
7	Trust in incentives	Some say GitHub promotes AI, not fights the problem	11

Score out of 20 on reach, severity, evidence and fixability (my judgement).

07 · SOLUTIONS

Five ideas to ease the review load

A

PR readiness gate

Rules a PR must meet before review: linked issue, template, checks

C

AI-use disclosure

A field saying whether and how AI helped

E

Pre-submit nudge

A checklist and warning before the author submits

B

Ranked triage queue

PRs sorted by completeness and history

D

Contributor trust tiers

Gentle limits for first-timers that lift as PRs merge

A readiness gate ranks first on RICE

Rank	Solution	Reach	Impact	Confidence	Effort	RICE
1	A. PR readiness gate	300	2	0.8	4	120
2	C. AI-use disclosure	300	1	0.8	2.5	96
3	E. Pre-submit nudge	500	1	0.5	3	83
4	B. Ranked triage queue	400	2	0.5	8	50
5	D. Contributor trust tiers	250	2	0.5	5	50

RICE = (Reach × Impact × Confidence) ÷ Effort. Reach is thousands of repos per quarter. All inputs are illustrative estimates.

09 · MVP

MVP: an opt-in readiness gate

A repo setting with three switches (linked issue, completed template, passing checks) and a checklist that tells the author what is missing.

Fewer
review hours per merged PR

≤5%
relative drop in first-timer merges

A/B
volunteer repos vs control

10 · TEST

An A/B test across volunteer repos

Item	Definition
Hypothesis	A readiness gate cuts review effort without blocking genuine contributors
Success	Fewer review hours per merged PR; fewer PRs closed unreviewed
Guardrails	First-timer merges drop no more than 5%; total merged PRs hold
Stop rule	Pause if first-timer merges fall more than 10% against control

Four gated phases

Phase	What happens	Gate
1. Pilot	Maintainers who raised the problem plus mid-size projects	Review hours fall, no guardrail breach
2. Expand	Repos with many PRs closed unreviewed	Stable across sizes and languages
3. Launch	All public repos, off by default; changelog and guide	Adoption up, complaints flat
4. Iterate	Add disclosure, then trust tiers and queue	Evidence supports each

Free for public repos; paid org-wide analytics could be tested later.

12 · RISKS

What could go wrong

Newcomers discouraged

Gates may put off real first-timers. The nudge and the guardrail exist for this.

Trust gap

Maintainers may distrust a platform that also sells Copilot. Involve them early.

Overlap

GitHub may already be building similar features. Check before committing.

TAKEAWAYS

Three things to remember

1. GitHub merged 518.7M PRs in a year, and one maintainer says review burden is higher than before AI.
2. Today's controls are blunt: limit PRs to collaborators, or switch them off.
3. Test an opt-in readiness gate first, and watch first-timer merges as the guardrail.

Would you gate PRs on your repo?

SOURCES

Sources and caveats

Sources

GitHub Octoverse 2025; GitHub changelog

The Register; DevClass (Feb 2026)

GitLab FY2026 8-K; Storyboard18; Siasat

Tech Insider (Copilot prices)

Caveats

No current revenue or market share

Copilot prices are secondary

Scores are my own estimates

Thank you

Questions and feedback welcome