

CONSOLE.LOG()

The Generative AI Policy Landscape in Open Source

By [kate holterhoff](#) | February 26, 2026

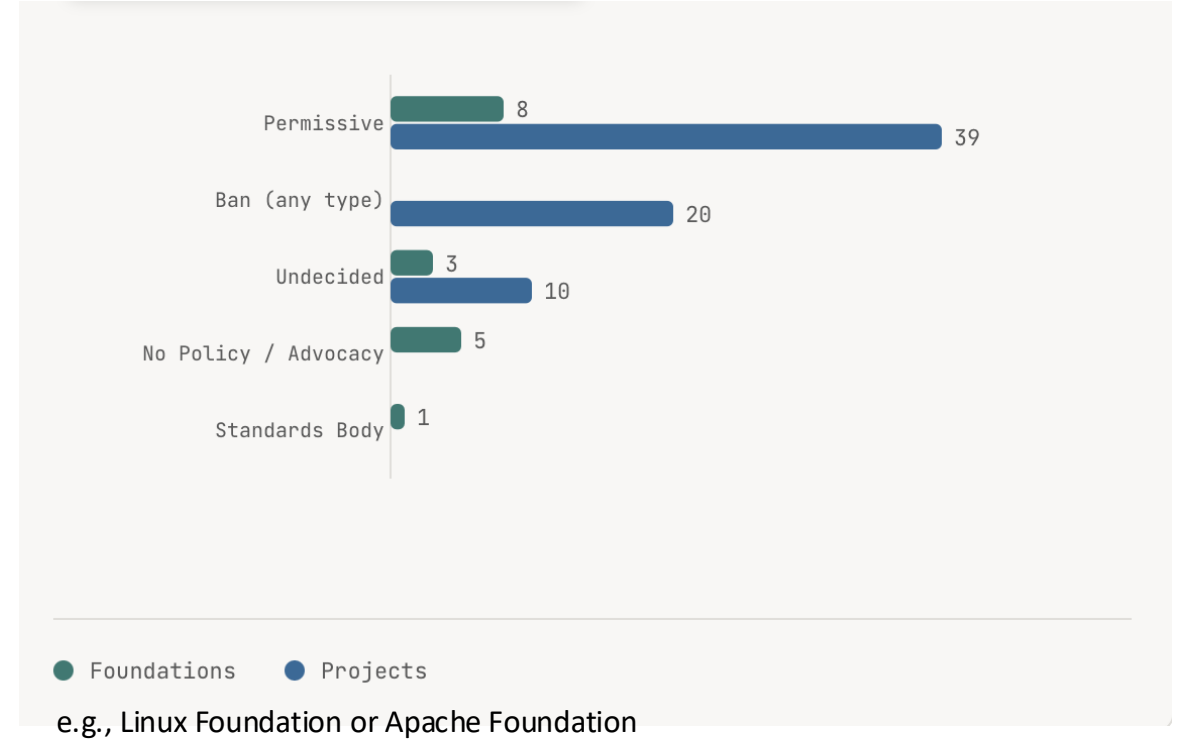

In a [RedMonk article](#), Kate Holterhoff maps and clarifies how open source communities are responding to the rise of generative AI, especially in terms of governance and contribution policies.

Rather than arguing for a single position, the author documents policies across 87 organizations and highlights emerging patterns, tensions, and open questions.

Using an empirical, survey-based approach, the article analyzes how open source foundations and projects are responding to generative AI, categorizing their policies and identifying key concerns such as legal risk, quality, and accountability, while emphasizing that the ecosystem remains fragmented and in transition.

	QUALITY	COPYRIGHT	ETHICS
Apache Arrow	M	M	L
Apache DataFusion	M	M	L
Apache Kvrocks	M	M	L
Apache Software Foundation	M	H	L
Asahi Linux	H	H	H
Blender	H	M	M
Castle Game Engine	H	L	L
Chromium	H	H	L
Cloud Hypervisor	H	H	L
CloudNativePG	H	M	M
CPython	H	L	L
curl	H	L	L
Debian	H	H	M
Django	M	M	L
Eclipse Foundation	H	H	M
Electronic Frontier Founda...	H	M	H
elementary OS	H	H	H
FastAPI	H	L	L
Fedora	H	H	M
Firefox	H	M	L
Forgejo	H	H	H
FreeBSD	M	H	L
FreeCAD	H	M	L

Policy Stance



No shared standard yet; governance is evolving independently across projects.

Policies Are Driven by Risk Concerns

Licensing and provenance are unclear for AI-generated code.

Maintainers report increased low-quality contributions (“AI slop”).

Accountability is ambiguous when issues arise.

Security and trust concerns are harder to manage.

Most policies are reactive responses to these risks rather than proactive design.

Three Emerging Policy Approaches

Restrictive: Limit or prohibit AI-generated contributions.

Permissive: Allow use with minimal guidance.

Conditional: Allow use with disclosure and human accountability requirements.

These categories show early patterns, but no dominant model has emerged.

Policies Are Driven by Risk Concerns

Traditional assumptions about authorship and licensing are breaking down.

Maintainers face increased workload and review pressure.

The “anyone can contribute” model is becoming harder to sustain.

Generative AI is leading to a reassessment of open-source governance practices.

The Landscape is Still Evolving

Policies are experimental and frequently changing.

Organizations are learning through trial and error.

Community norms around AI use are not yet established .

The ecosystem is in a transitional phase.

Takeaway

- Generative AI is forcing open source to rethink its governance model.
- Moving from implicit trust to explicit verification.
- Shifting from purely human authorship to hybrid contribution models.
- Requiring clearer policies where norms once sufficed.
- This is not just a tooling shift. It is a fundamental change in how open source operates.

What Does This Mean for ePIC?

- We aim for a vibrant developer community. AI tools can lower barriers to contribution, but they introduce ethical, legal, and technical risks.
- These risks should be addressed through a lightweight policy rather than strict restrictions.
- This is an evolving area, and we will coordinate with collaboration leadership and host labs on how to approach this.

Initial direction for software development workflows

• Submitters

- We do not plan to restrict which tools contributors use in PRs.
- We encourage transparency: adding an optional field in the PR template to disclose AI-assisted contributions.
- Contributors remain responsible for understanding and validating all submitted changes.

• Reviewers

- The role of AI tools in code review is still under discussion.
- We will evaluate this further today, discussing lessons learned CodeRabbit.