AI Tool and Agent Sprawl 2026 Media Brief

Public AI Adoption Is Easy to See. Governed Use Is Much Harder to Prove.

- Run ID: sprawl-v2-full-20260312b
- Scope: 890 completed public GitHub targets from a frozen 1,000-target publication cohort
- Study design: deterministic baseline, one repository per owner, public clone-sourced artifacts

The Short Version

CAISI measured public AI tool and agent governance posture across an 890-target publication subset drawn from a frozen 1,000-target GitHub cohort. The strongest result is not visible autonomous abuse. It is that public repositories often make AI adoption visible before they make approval, deployment, and governance evidence visible.

In the completed subset, 98.2% of targets declared at least one agent. But only 4.04% showed any deployable-agent signal, and 100% of detected agents were missing at least one declared binding. Nearly half of the targets, 47.08%, lacked verifiable governance evidence in the repository. On the tool side, non-source AI tools outside the baseline-approved set outnumbered baseline-approved tools 11:1.

The report is explicit about what it does and does not prove. It is a public-repository visibility study, not a production-runtime exploit census. It does not claim that public zeroes on write-capable or exec-capable agents mean internal runtime risk is low. It shows that public AI and agent use is easier to detect than public proof of governed use.

Headline Findings

- 874 of 890 targets (98.2%) declared at least one agent.
- Average declared agents per target: 8.97.
- 36 of 890 targets (4.04%) showed deployable-agent signal.
- 7984 of 7984 detected agents were missing at least one declared binding.
- 419 of 890 targets (47.08%) lacked verifiable governance evidence.
- Non-source AI tools outside the baseline-approved set outnumbered baseline-approved tools 11:1 (715 to 65).
- 153 of 890 targets (17.19%) showed an EU AI Act Article 50 transparency-proxy gap.

Why This Matters

For AppSec teams, the report shows that the first measurable control problem is often evidence posture rather than a fully exposed exploit path. For CISOs, the approval ratio is best read as a proof gap: visible AI use without durable machine-readable approval evidence. For platform leaders, the report shows that declaration volume and deployable-agent evidence are not the same thing.

The practical governance lesson is simple: discovery matters, but approval records, binding completeness, and evidence continuity have to mature with adoption. Otherwise organizations can detect AI use without being able to defend, review, or explain it cleanly.

Scope And Limits

- This is a completed subset report, not the full frozen 1,000-target cohort.
- The unfinished 110 targets are all AI-native, so some AI-native-heavy findings are likely understated in the published subset.
- Public repositories underexpose private runtime privilege, credentials, and internal approval workflows.
- Regulatory outputs are deterministic readiness proxies, not legal conclusions.

Links

- Full report PDF
- Media brief PDF
- Report page
- Report package
- Run artifacts
- research@caisi.dev