


# Gomboc AI Publishes First Open Benchmark for AI Code Remediation

*15 cloud scenarios. 43 merge-ready fixes. 100% loop closure. 12 minutes and \$17 to author once; seconds and zero-cost apply everywhere after.*



NEW YORK, NY, UNITED STATES, June 2, 2026

/EINPresswire.com/ -- [Gomboc AI](#) today published the first open benchmark for AI code remediation,

documenting the result of its deterministic remediation platform across 15 real-world cloud scenarios spanning AWS, GCP, and Azure. The benchmark is publicly available at <https://www.gomboc.ai/show-your-work>.

The methodology is open, and the scenarios are published on GitHub. Any team can run the same benchmark and publish their own numbers.

PR review time is up 91%. Developers are merging 98% more code. AI-generated changes are reaching production faster than any team can validate them. And every time a fix needs to be redone, organizations pay the token cost twice. Gomboc AI sits on top of the tools teams are already using - Cursor, Claude Code, Copilot - to make every fix accurate, governed, and cost-optimized. The benchmark is the proof.

The benchmark covers 15 production cloud scenarios across security, reliability, and cost.

- The security findings cover misconfigured IAM policies, open network security groups, and unencrypted storage across AWS, GCP, and Azure.
- The reliability findings include a production database handling 50,000 orders a day with no backups and no failover.
- The cost findings include dollar amounts: \$2,050 per month for duplicate CloudTrail configurations, \$279 per month for redundant NAT gateway routing, and \$870 per month for oversized EC2 instances.

Every fix generated by Gomboc is idempotent, tested, and traceable to a policy. None requires a human to interpret the output before it can be applied.

The benchmark numbers represent the worst case, not the steady state. The 12 minutes and \$17

in tokens cover the one-time work of authoring policies for scenarios Gomboc had never encountered before. Once a policy exists, applying it across every repository, every pipeline, and every future occurrence of the same issue takes seconds and effectively incurs zero in token cost.

This is the core of the Gomboc economics: enterprises pay once to codify a fix, then apply it infinitely. Every other AI coding tool pays the full generation cost every time a similar issue appears, because there is no memory, no policy layer, and no governance to carry the work forward.

"The era of vibes-based AI tooling is over. Every team evaluating AI for code remediation deserves to know five things: Is the output idempotent? Is it governed? Does it align to a policy? Can it be reproduced? Is there an audit trail? We can answer yes to all five. We are inviting every other tool in the stack to do the same." [Ian Amit, CEO and Co-Founder, Gomboc AI](#)

The benchmark methodology is fully documented and reproducible. The 15 scenarios are published on GitHub as open ORL files, alongside the rule sets, test cases, and expected outputs for each fix. Any team can clone the repository, run the benchmark against the same scenarios, and publish their own results. This is intentional. Gomboc is not asking the industry to take its word for it. We are asking the industry to run the same test and show its work.

"AI that suggests a fix is not the same as AI that executes safely in production. The real question is whether those changes are deterministic, policy-aligned, and repeatable across an enterprise. That is the problem we built Gomboc to solve, and the benchmark is how we prove it." [Matthew Sweeney, Co-Founder and CTO, Gomboc AI](#)

The benchmark is available now at <https://www.gomboc.ai/show-your-work>. The scenarios are open-sourced at <https://github.com/Gomboc-AI/rattleback/tree/main/scenarios>.

Gomboc AI's Community Edition is available free at <https://www.gomboc.ai/community>

About Gomboc AI:

Gomboc AI is the only AI Code Security Assistants (ACSA) platform delivering 100% accurate, deterministic fixes sitting above the generative AI tools teams already use. Built for DevOps and platform teams, Gomboc automatically converts cloud and Infrastructure-as-Code risks into safe, merge-ready code changes that scale across modern engineering environments without slowing delivery. For more information, visit [gomboc.ai](https://gomboc.ai).

Sonia Awan

Outbloom Public Relations

[soniaawan@outbloompr.net](mailto:soniaawan@outbloompr.net)

Visit us on social media:

[LinkedIn](#)

---

This press release can be viewed online at: <https://www.einpresswire.com/article/916564444>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.