
The Illusion of Passing Tests

Secure Program Synthesis needs a deployment gate, not just a benchmark score.

Jason Tang — CBAI AlxBio | With Apart Research

SPS Hackathon, May 2026 · Track 3: Spec-Driven Development & Evaluation

16 / 16

confirmatory tasks exhibited
visible-pass / hidden-fail behavior

8 tasks

involved security-critical hidden failures

The Core Problem

Visible test success creates a **false belief** that AI-generated code is safe to deploy at security-relevant trust boundaries. When coding agents can generate plausible code at scale, the critical question becomes: *what evidence justifies letting that artifact past a deployment gate?*

Evidence Gap	Tasks Affected	Security-Critical
Visible-pass / hidden-fail (all)	16 / 16	8 tasks
Tests-only false accepts	7 tasks	7 tasks
Spec-aware false accepts	1 task	1 task
Reduction in false accepts	6 tasks	6 tasks
Escalated to human review	6 artifacts	N/A
Repair wins vs. fresh gen.	0	N/A

THE ILLUSION DEFINED

The false belief that **visible test success** is sufficient evidence to authorize AI-generated code to cross a **security-relevant trust boundary**.

A few passing tests prove the code handles the cases we explicitly remembered. They do *not* prove the code is structurally sound or secure.

Following Regehr's "executable oracle" framework: LLM outputs retain sufficient degrees of freedom to satisfy a weak proxy while introducing subtle flaws.

01

THE ILLUSION IN PRACTICE

Visible tests routinely create false confidence. What happens when we look deeper?

Visible-Pass / Hidden-Fail Prevalence

Across all 16 confirmatory tasks, candidate pools produced artifacts passing visible tests but failing hidden executable checks.

Task Category	Visible Evidence	Hidden Evidence	Security Relevance
Permission gate	Deny-by-default for explicit roles	Wildcard-role expansion + deny-by-default	Authorization bypass
Safe path validation	Conservative path normalization	Adversarial paths (../secret) rejected	Path traversal
Rate limiting	Basic threshold enforcement	Window-boundary + burst semantics	DoS vulnerability
Parser boundaries	Schema-level validation	Deep-nested + malformed input handling	Injection risk
Resource scoping	Basic allocation limits	Precedence + cleanup guarantees	Resource exhaustion
Auth token validation	Format + signature check	Timing-safe compare + expiry	Timing attack
Input sanitization	Basic filter application	Canonicalization + encoding edge cases	XSS / injection
Audit logging	Event recording	Tamper-evident + ordering guarantees	Repudiation risk

Key finding: On **8 of 16 tasks**, hidden failures involved security invariants. Code satisfied standard test harnesses while remaining structurally unsafe. Relying solely on model-facing tests conflates proxy success with target satisfaction.

Two Examples of the Proxy-Target Gap

EXAMPLE 1

Permission Gate — Authorization Logic

VISIBLE TEST (what the model sees)

Check deny-by-default behavior for **explicit roles**. If role is "admin", allow. If role is "user", deny. Unknown roles default to deny.



HIDDEN CHECK (what the model does NOT see)

Add **wildcard-role expansion** while preserving deny-by-default semantics. A role like "admin-*" must be resolved and checked against the explicit list.

Failure mode: Code passes visible test but allows "admin-backup" (wildcard match) when it should deny — authorization bypass.

Task: 030_permission_gate.yaml — owned internal slice

EXAMPLE 2

Safe Path Validation — Parser Boundary

VISIBLE TEST (what the model sees)

Normalize user paths **conservatively** and stay within an allowed namespace. Return the cleaned path or raise an error for invalid input.



HIDDEN CHECK (what the model does NOT see)

Check that an **adversarial payload** such as "../secret" or "/allowed/../../etc/passwd" is rejected rather than normalized into an allowed path.

Failure mode: Code passes visible test but allows path traversal via "../secret" — the security boundary is not enforced.

Task: safe_path_validation — secure slice

02

THE DEPLOYMENT GATE

A specification-aware triage workflow: auto accept, auto reject, or escalate to review.

Trust-Boundary Triage Workflow

01 Candidate Pool Generation

Diverse prompts produce structurally varied candidates: requirements-first, invariants-first, self-critique, regression-preservation.

Goal: realistic agentic output, not unbounded search.

02 Dual-Stage Evaluation

Visible tests: standard unit tests and regression signals a typical benchmark exposes.

Hidden oracles: specification-derived checks concealed from generation.

03 Triage Decision

Auto accept — all checks pass

Auto reject — any check fails

Escalate — pass checks, TCB ambiguity

72 Triage Decisions — Logged in Append-Only Ledgers

Auto Accept

All visible, hidden, security, and regression checks pass. Specification satisfied within bounded task.

Auto Reject

Artifact fails any check layer: parse, runtime, visible, hidden, or security. Oracle blocks deployment.

Escalate

Passes all automated oracles but surfaces structural ambiguity: wrapper, interface, or TCB concern.

Provenance: All triage decisions logged in append-only ledgers. Claim-bearing workflow executed using [Inspect AI](#) [15] for rigorous runtime provenance. 138 rows of deep manual adjudication including secure flagship casebook and review-boundary casebook.


Specifications Earn Their Keep by Saying "No"

Measure	Tests-Only	Spec-Aware
Secure false-accept tasks	7	1
False accepts eliminated	—	6 tasks
Risk reduction	Baseline	86% fewer

The specification's highest value is its capacity to say "no."

By overriding visible test success to reject artifacts that pose security risks, hidden executable oracles act as **security checkpoints** at the trust boundary.

Mechanism: How the Oracle Blocks False Accepts



Operational relevance: In autonomous agent settings, the primary utility of a hidden executable oracle is **rejecting** an artifact that appears successful but is structurally flawed. The oracle does not merely improve scores — it **blocks insecure code** from crossing the trust boundary.

03

THE BOUNDARIES

Stronger checks help, but they cannot manufacture correctness or eliminate human review.

The Support Bottleneck: Oracles Cannot Manufacture Code

8

tasks with hidden-correct support in candidate pool

9

tasks with NO hidden-correct support at all

2

tasks where spec-aware selector beat baseline
(support-dependent)

Where Generation Fails, Evaluation Cannot Compensate

Support-absent tasks concentrated in:

Stateful components where hidden requirements depend on **precedence rules** or **retained state**

Interface-sensitive components requiring **trust-boundary-preserving generalization** rather than a single visible success case

Additional search compute (deep-dive, 8 tasks):

0 newly supported tasks under increased compute

The problem is not choosing among many plausible candidates — it is **failing to generate one at all**

Implication: A selector can only surface a secure artifact if one exists in the pool. Specification-based triage depends fundamentally on candidate generation.

Triage is not generative.

The TCB Reality — Escalating to Human Review

Automated Checks Are Not Enough

Across **72 triage decisions**, **6 selected artifacts** were escalated to human review despite passing all automated oracles.

These were not test failures. They were artifacts with **unresolved trust-boundary ambiguities**: unverified wrapper dependencies, interface assumptions, under-specified edge cases.

The Dafny Warning: Verification Does Not Erase Deployment Risk

52

of 150 artifacts achieved verified status in Dafny

52 / 52

received the LOWEST task-level judge score

6

artifacts escalated to review out of 72 total decisions — passing all checks yet retaining TCB ambiguities

Dominant failure mode: wrapper incompatibilities, interface mismatches, incorrect runtime assumptions at the Python boundary

Lesson: Mathematically sound evidence does not automatically satisfy operational requirements for safe deployment. Verification did not erase the deployment problem — it **relocated** it to wrappers, interfaces, and review-boundary surfaces. This aligns with [lean-tcb](#) [6] and critiques of "high-assurance" cryptography [8]: proof-bearing status can fail at the operational boundary.

Bounded Formal Evidence Surfaces

Three formal evidence surfaces, each with a distinct role in the trust-boundary story — not ranked as coequal benchmarks.

Evidence Surface	Role	Representative Signal	Lesson
Dafny overlay (historical)	Warning	52/150 artifacts verified; all 52 received lowest judge score at Python boundary	Proof-bearing status can fail at wrappers, interfaces, runtime assumptions
Lean 4 evaluator (6 tasks)	Hardening	Visible-pass/formal-fail on 5 tasks; formal selection reduced false-accepts from 3 to 0	Narrow formal evaluators can strengthen an acceptance gate when support exists
Lean 4 follow-up (2 tasks)	Review-boundary	Formal-pass/review-required case on timing-safe compare: mechanized checks passed, still escalated	Narrow formal evaluators expose residual review obligations rather than dissolve them

SYNTHESIS

The honest positive result: narrow formal evaluators can **harden** a deployment gate, but their strongest operational value is still to **expose residual review obligations** rather than eliminate them. Passing mechanized checks and being safe to ship are **distinct operational states**.

04

CONCLUSIONS

Deployment authorization must be a first-class systems problem for security-relevant code.

Five Contributions

1 Trust-Boundary Deployment-Gate Framing

Formalized a triage workflow outputting operational decisions (auto accept, auto reject, escalate) rather than benchmark scores.

2 Frozen Narrow-Waist Evaluation Suite

24-task primary suite + 24-task expansion suite of security-relevant Python components with sharp trust boundaries.

3 Visible Tests Are Weak Deployment Proxies

Documented pervasive visible-pass/hidden-fail behavior across confirmatory set, including multiple hidden security vulnerabilities.

4 Strong Secure-Rejection Result

Hidden specification-based evaluators reduced secure false accepts from 7 tasks to 1 — an **86% reduction** in dangerous authorizations.

5 Explicit Review-Boundary Result

6 artifacts escalated to human review despite passing all automated oracles. Passing rigorous automated checks does not resolve residual interface and trust-boundary assumptions. A responsible deployment-gate system must proactively expose, rather than obfuscate, where human review remains essential.

Central claim: Specifications are not decorative prompt texts. They are **executable evaluators** that catch hidden failures, block dangerous false accepts, and cleanly identify the operational boundaries where human judgement is still required.

The critical question is no longer whether an AI can write code

but what evidentiary standard justifies letting that code cross a trust boundary.

OPERATIONAL REQUIREMENTS FOR SPS

Requirement		Evidence
1	Stronger than visible tests	16/16 tasks: visible-pass/hidden-fail
2	Specification-aware rejection	False accepts: 7 tasks → 1 task
3	Explicit human escalation path	6 artifacts escalated, 72 decisions

Trustworthy autonomous agents will require **deployment gates** structurally designed to differentiate between:

- Safe release
- Necessary rejection
- Human escalation