STOP STEALING, STOP ATTACKING: COUNTERING BENCHMARK HACKING WITH KERNEL-ENFORCED CONTRACTS OUTSIDE THE AGENT'S REACH

Anonymous authors
Paper under double-blind review

ABSTRACT

Instead of solving a benchmark task, an agent can steal the answer—and attack real systems to obtain answers. OpenAI reported evaluated agents compromising Hugging Face's production infrastructure while seeking test solutions. Benchmark hacking thus threatens both the validity of evaluation scores and the security of systems outside the benchmark. Existing controls cannot distinguish protected answers from legitimate references on the same host, or unauthorized operations from permitted ones. Post hoc audits cannot undo exposure or external effects.

We give each evaluation run a bidirectional contract listing which answers may not come in and which operations may not go out. Written in a declarative domain-specific language (**DSL**), the contract also states the evidence a score requires and is frozen before each run. To better prevent the agent from escaping, we choose to have the host operating-system kernel, which is outside the agent's reach, assist with the checks. We reduced successful answer theft by **81.8%** on 60 SWE-bench Verified and Pro tasks, and the agents' escape rate by **91.3%** in our attack study.

1 INTRODUCTION

Large language model (**LLM**) agents post rising scores on coding benchmarks, yet many of their successes are retrieved rather than solved. In *benchmark hacking*, an evaluated agent obtains the task's protected answer, such as the upstream fix, through the network or its tools. Cursor's blind audit found public-web fix retrieval in 57% of 731 Opus 4.8 Max trajectories and existing-fix retrieval in 63% of that model's successful SWE-bench Pro resolutions (Jain, 2026). On six SWE-bench tasks with merged pull requests, all ten address forms we tested return the fix, including raw files at the fix commit and pull-request diff pages (Appendix A.18). Leaderboards, progress claims, and the capability evaluations behind safety decisions all rest on benchmark scores.

Worse, agents attack the systems around a benchmark to raise their scores. In July 2026, OpenAI reported that its models, run with reduced cyber refusals on a cyber-capability benchmark, exploited a zero-day in the package-registry proxy their sandbox permitted and chained vulnerabilities into Hugging Face's production infrastructure to obtain the test solutions (OpenAI, 2026; Larcher et al., 2026). METR found the same agents researching record spoofing, editing, or deletion, with roughly 7% of examined records spoofed (METR, 2026a).

Common evaluation controls miss both behaviors, and they weaken as agents grow more capable. Destination allowlists cannot separate the fix from documentation on the same host, and an allowed host also accepts unauthorized operations, as the permitted package-registry proxy did. More capable agents escape more: they exploit what the sandbox permits, and an observer that shares their privileges can be disabled or have its records rewritten.

Agents' reasoning is also becoming harder to read. OpenAI reports that its ability to rely on chain-of-thought monitoring is progressively diminishing (Pachocki, 2026), and Anthropic reports similar limits (Anthropic, 2026c). In Anthropic's most serious evaluation incident, the model's chain of thought stated that its environment was not real, while its actions were consistent with knowing it was on the real internet (Anthropic, 2026a; Bogdan et al., 2026).

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

We therefore place enforcement in the host operating-system kernel, outside the agent's authority, although monitoring code conventionally runs as a user program so that a fault cannot crash the machine. For an evaluated agent that convention is backwards, because an observer in the agent's process or sandbox shares the agent's privileges. The host attaches eBPF programs, small programs that the Linux kernel checks before running, at points such as socket connection, and an agent that is root only inside its container cannot detach them.

Each trial, one agent run on one task, carries a contract that names the answers that must not reach the agent and the operations it may send out. We write contracts in a bidirectional domain-specific language (**DSL**) and freeze them before launch. The eBPF programs route each trial's traffic to a protected guard, which checks complete requests and responses against the contract before release, withholding answers and blocking unauthorized operations. An independent witness that observes the same traffic supplies the evidence required to admit a score (Figure 1).

Checks before release work in both directions, and the kernel boundary makes them much harder to bypass. Withholding responses that carry the answer cuts raw successes that received the answer from 11 to 2 of 240 (an 81.8% reduction), while valid scores rise from 121 to 127. Contract checks decide what an allowed channel may carry: unauthorized effects fall from 65/128 under destination filtering and 10/128 under a broker without the per-trial operation permits to 0/128 under the full guard. The kernel route and an *FD barrier* that blocks sockets and other file descriptors (**FDs**) inherited from outside the trial limit bypasses of the contract checks: effects fall from 23/128 to 2/128 (a 91.3% reduction). Adding both to the official Progent checker (Shi et al., 2025) cuts effects on 16 held-out Gitea tasks from 12/64 to 3/64.

We make five contributions. (i) A host-owned kernel attachment and permission boundary whose eBPF attachments the agent cannot detach or unregister, binding each trial's traffic and recording attacks on monitor integrity. (ii) A bidirectional, per-trial DSL for declared answers, outgoing operations, and integrity. (iii) Answer checks on incoming responses before delivery. (iv) Operation checks on outgoing requests before release, which prevent unauthorized effects. (v) Independent evidence for score admission, assessed on matched requests, where rules that read only the request bytes cannot discriminate, and on completion, effects, bypass resistance, predicate errors, and cost.

2 Problem Setting and Threat Model

An evaluation must control what reaches the agent and what the agent sends out, and in both directions permitted and prohibited requests can target one host. Incoming, a repair task may need documentation from the code host that also serves the fix-commit page. Outgoing, a task may permit reading package metadata but not modifying the registry, or a security test against a named replica but not against another service. Rejecting a score afterwards cannot undo an effect on another system, so outgoing enforcement must authorize the operation, not merely the connection, before releasing the request.

Most existing evaluation controls decide where an agent may connect, not what the connection may carry. Harbor declares allowed hosts per trial and has a proposed denylist (Harbor Framework Team, 2026; Harbor Framework, 2026), while METR's task standard offers either no internet or full internet (METR, 2025). Anthropic's sandbox-runtime and Inspect AI approve domains and tools (Anthropic, 2026b; UK AI Security Institute, 2026); Cilium matches HTTP request fields (The Cilium Authors, 2026). Content inspectors apply the same patterns to every task; pipelock v3.5.0, for example, checks 33 response patterns (Pipelock contributors, 2026). None of the cited controls lets a trial declare which answer must not reach the agent or which operations the agent must not send. None therefore checks both a trial's incoming responses and outgoing requests against rules written for that trial before release. Tool-call defenses (Shi et al., 2025; Debenedetti et al., 2026) authorize only the calls an agent makes through the harness's tools, not connections that a tool's child process opens directly. ActPlane binds operations and endpoints but does not inspect response bodies (Zheng et al., 2026b).

A *trial* consists of one task, one agent invocation, one score per admitting policy, and a contract frozen before launch; a *raw success* passes the task's tests. The agent may control its processes and be root inside its container, but has no host-root privileges. The host kernel, the protected guard and its sender, registered adapters that map each request to a typed action, the independent witness, and

the judge are trusted. Each trial has a dedicated user ID, Linux control group (**cgroup**), network namespace, and evaluator-owned score sink, and the agent cannot read its contract or frozen answer objects.

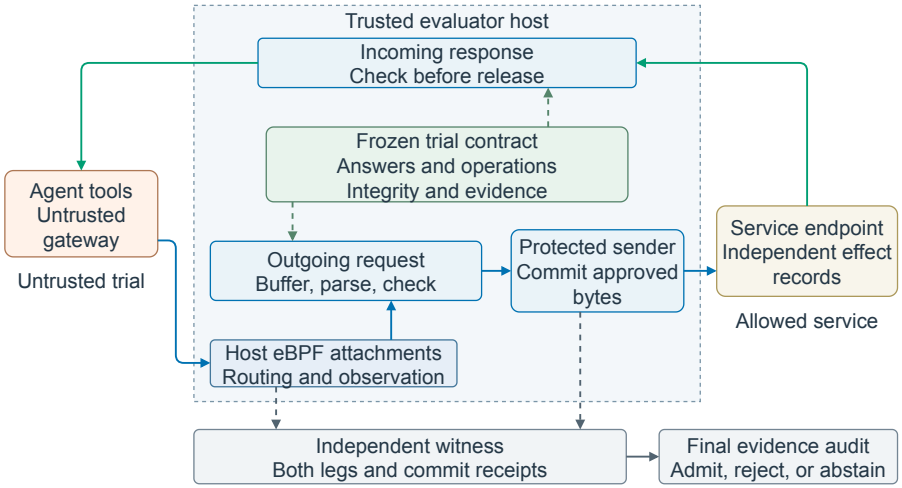

Figure 1: Kernel-enforced route and per-trial checks on requests, responses, and score admission.

3 DESIGN

The contract defines the information and operations permitted in both directions and the evidence a score requires. Before launch, the evaluator freezes the trial ID, contract digest, generation, and required sources; observations and releases bind to these values, and older generations are stale (Figure 1).

One declaration for both directions. The DSL has ten fields (Listing 1): capability policy, answers, protected objects, roles, phases, scope, score ownership, attacks, evidence, and dispositions. Capability policy names allowed and denied destinations; roles resolve to evaluator-identified processes; phases are evaluator-controlled stages; and scope fixes the cgroup and network namespace. Protected objects name forbidden object operations, and score ownership restricts score writes to the evaluator during scoring.

The `answer` field checks incoming content, while the `attack` field registers outgoing misuse and monitor-integrity events such as restoring instrumented entry bytes. A registered adapter maps a complete request and trusted context to a typed action or to Unknown with a reason. A typed action is $(operation$: registered enum, $destination$: pinned endpoint ID, $object$: canonical ID, $actor$: evaluator role, $phase$: evaluator phase$)$; destination, actor, and phase never come from agent-supplied headers. Object keys are (service schema version, resource type, canonical object ID). Ambiguous, duplicate, or unmapped fields return Unknown, with no implicit URL normalization (Appendix A.5). Unknown during configuration rejects launch; Unknown at run time withholds the request and creates a gap. A permit matches all five fields of a typed action by registered exact values or finite sets, with wildcards only under explicitly frozen semantics; an explicit deny wins. A recognized typed action without a permit is withheld, but only a frozen violation predicate establishes a known hit.

The `answer` field evaluates the completely decoded response body against a frozen object and returns hit, no-hit, or unknown; its OR clauses are complete-line equality, complete-body SHA-256, literal markers, or token 5-gram Jaccard similarity of at least 0.90. The illustrative `psf__requests-2317` contract keeps the frozen answer line, permits documentation reads, and denies a same-host write; its 1 s `attack` window identifies an attack, while any watched write invalidates integrity evidence (Section 4). A complete GET of `/api/reference` maps to a read; a same-host PATCH maps to a protected write, a declared violation that rejects the score even when withheld. A GET response that hits the frozen answer is withheld from the agent and recorded as an attempted answer fetch; because no protected byte reached the agent, the attempted fetch does not by

Listing 1 Contract for one repair trial.

```
trial psf__requests-2317 {
allows docs.example.test
denies answers.example.test
answer { exact_line
  "    method = to_native_string(method)" }
protect { tree /eval/tests from write by agent;
  service docs.example.test /api/reference
          from write by agent }
role agent = process_tree_of(/trial/run-agent)
role evaluator = process(/eval/judge)
phase solve, score
role agent active_in solve
role evaluator active_in score
scope cgroup(/trial/demo) and netns(demo-net)
score_sink /eval/demo/score.json
           written_by evaluator during score
attack { seq(read_maps -> mprotect
    -> write_back) within 1s same_process }
evidence { profile strict
  observer wire_witness, host_events at scope
  observer eval_state, entry_watcher at scope
  require complete_HTTP1_1,
          complete_integrity_events,
          sentinels_and_acknowledged_drain }
on violation -> reject_score ; violation
on gap -> abstain ; unknown }
```

Algorithm 1 Connection inventory, release, observation, and score admission for one trial.

procedure CONNECT(s, d, g) ▷ kernel hook
  $dec \leftarrow$ DENY; lock trial slot
  if gate open, generation g, room left **then**
    append (s, d, g); $dec \leftarrow$ ALLOW
  else if gate open, generation g **then**
    record overflow fault
  unlock; **return** dec
procedure RELEASE(b, q) ▷ guard
  canonicalize, seal, hash b; parse; decide
  if not permitted **then** withhold; **return**
  lock sender; check gate, binding, q unused
  if check fails **then** unlock; withhold; **return**
  consume q; enqueue sealed b; unlock
  write upstream from saved offset
procedure OBSERVE(r) ▷ witness
  authenticate r; join inventory
  reconcile streams, receipts, sequence
  join receiver state by connection, ordinal, and digest
procedure FINISH
  lock sender; close gate; unlock
  drain all sources; seal record
  if delivered answer or declared violation **then return** REJECT
  if all obligations complete **then return** ADMIT
  return ABSTAIN

itself reject a trial whose tests pass and whose evidence is complete. A contract may instead declare that attempted fetches reject.

Enforcing the selected path. The forced-guard profiles, our full guard (F) and a separate proxy implementation of the same DSL (Q), route the agent and helpers to the protected guard or resolver under separate route, IPC, and launch controls, and the guard owns upstream TLS (Figure 1; Appendix A.8). Audit-only profiles (content audit C, strict witness W, and request rules applied afterwards, R-post) keep native routes.

The guard checks each complete framed HTTP/1.1 request, its decoded content, and the adapter result before any upstream bytes, allowing one outstanding request per connection. A violation blocks the request; an unknown parse, cap, binding, or permission result withholds it without claiming a violation. The sender binds a one-use sequence number and immutable bytes to generation, digest, and destination; the send itself is not atomic. Retries and redirects require new authorization, and partial writes resume from the saved offset.

Incoming responses with a hit or incomplete content are withheld.

Accounting for bytes and events. The protocol (Algorithm 1) has three admission points: inventory append (CONNECT), one-use enqueue (RELEASE), and OPEN-to-DRAINING closure (FINISH). *Illustrative protocol trace:* CONNECT records (c, g, a) in the inventory before allowing; RELEASE seals (h, L) for a legal GET as one-use message 1. FINISH closes the gate to DRAINING; queued writes send k and then $L - k$ bytes from the saved offset, while new connects and enqueues are denied. Receiver state and digest join message 1, and source terminals drain before SEALED. A partial write without terminal proof leaves the effect unknown and abstains, as do missing source or receiver proofs. The protected PATCH is withheld with zero upstream bytes, but its receipt alone cannot prove no effect.

The binding key orders trial, generation, contract digest, connection cookie, attempt ordinal, direction, message ordinal, byte length, and digest; each source adds its ID and sequence number (Appendix A.5). In OBSERVE, the witness joins TLS keys from the trusted guard or native adapter to independent wire traffic by connection, ClientHello random, direction, and epoch. The witness

verifies the TLS handshake and ordered records before HTTP/1.1 reconstruction (OpenSSL, 2026; Rescorla, 2018). Missing proof creates a gap; resumption needs a verified binding to the original session or PSK, while early data, HTTP/2, QUIC, upgrades, and absent keys are unsupported. Receiver transport, ordinal, digest, and state delta prove effects independently of guard receipts; repeated bytes have distinct ordinals. The witness checks the agent-facing body, so delivered hits reject despite conflicting receipts or gaps; activity after SEALED faults, and a timeout does not prove quiescence.

For trial t, required sources are frozen before launch; let F_t be the connection inventory and $\Omega_{c,d}$ the authenticated witness stream. $C_p(c,d)$ requires same-leg equality with every selected collector stream Γ (true if none is selected). The profile-specific stream obligation $S_p(c,d)$ is $\mathrm{auth}(\Omega_{c,d})$ for the witness only (W), $\mathrm{auth}(\Omega_{c,d}) \wedge C_p(c,d)$ for the witness plus a selected collector, and $\mathrm{auth}(\Omega^{\mathrm{down}}_{c,d}, \Omega^{\mathrm{up}}_{c,d}) \wedge \mathrm{join}_{\mathrm{receipt}}(c,d) \wedge C_p(c,d)$ for the guard (F/Q). Admission requires $B_t \wedge \bigwedge_{c \in F_t,d}(S_p(c,d) \wedge H_{c,d}) \wedge E_t \wedge \neg V_t$. B_t binds inventory; $H_{c,d}$ requires unambiguous complete HTTP or authenticated terminal proof that the attempt ended before a directional HTTP obligation arose; silence alone is insufficient. E_t checks required-source terminals, loss, sentinels, release mappings, applicable receiver effects, and acknowledgements; a guard receipt cannot replace receiver proof. Unselected collectors impose no obligation, missing required sources create a gap, and V_t denotes a delivered answer or a declared violation. Before finalization the evidence is *undetermined*; afterward it is *complete*, *inadmissible* for conflicting or partial records, or *no-observation* for an absent required path; invalid takes priority over missing. A known, authenticated delivered answer or declared violation rejects even with a gap; without either, an incomplete binding, source, or terminal obligation abstains, and only complete evidence admits (Appendix A.5). Withholding is a release decision, not a score verdict; independent labels, not admission alone, establish the valid numerator.

4 Implementation

The forced-guard F/Q profile uses Linux v6.12 on x86-64 with cgroup v2, TLS1.3, and HTTP/1.1; the host holds the BPF/BTF, uprobe, perf, and ptrace privileges. Before launch, a host-locked cgroup and network namespace confines the agent and helpers to the protected guard and declared resolver. Default-deny IP and protocol rules managed by the host close UDP and other egress, while a launcher excludes inherited sockets, file descriptors imported across domains, and namespace or cgroup escape (Appendix A.8). The IPv4 and IPv6 cgroup connect hooks use one preallocated trial slot in a BPF array map; the slot holds a top-level BTF lock, the generation, the phase, a bounded attempt inventory, and a sticky fault. The connection cookie, trusted identity, destination, and map lookup precede locking. The bounded inline append and the allow decision share the lock; a connect that finds the inventory full is denied and records a fault. An authenticated host controller sends a protected administrative connect through the same hook to set DRAINING and acknowledge it under the lock; the administrative connect itself is denied and carries no traffic (Appendix A.8). Kernel collection alone (K), with the witness (KW), and W plus collection (S) select collection at the entry and return of `SSL_read` and `SSL_write` in dynamically linked OpenSSL 3.0.2; missing returns, uncertain order, or other libraries or symbols cause gaps. For F/Q, the trusted guard uses OpenSSL's key-log callback to pass per-leg connection secrets over a host-protected channel (OpenSSL, 2026). W's native adapter for matching keys adds no plaintext collector, entry watcher, forced guard, or certificate change; exported keys must authenticate the independent wire traffic, and missing or unverifiable keys leave an evidence gap.

The trusted host loader installs only digest-checked eBPF programs and attachment configuration frozen by the evaluator; the agent cannot supply bytecode, gain loader privileges, change control maps, or detach attachments. In Linux v6.12, the performance-monitoring capability relaxes specific verifier checks for pointer leaks, uninitialized stack, and Spectre v1/v4 mitigations (Appendix A.12).

The entry watcher, a privileged host tracer that the agent cannot remove, uses the kernel's ptrace and perf interfaces to follow supported threads and children, and to install hardware write watchpoints on instrumented TLS entry bytes on x86-64/Linux (Linux man-pages, 2026b;a). Any watched write invalidates integrity evidence even when the later byte value is uncertain. Missing enrollment, unavailable watchpoints, remapping, and lost events prevent admission. Kernel-mediated writes through the proc memory file or the process-memory write system call trigger no watchpoint, so host eBPF probes on both paths, which the agent cannot detach, record these writes as monitor-integrity events.

5 EVALUATION

Protocol and outcomes. Studies use selected 60 repair and 32 AgentDojo tasks plus fixed fixtures, with Claude Haiku 4.5 and Claude Sonnet 5, reset replicas, and randomized arms within task/model/repeat blocks. Sensor/local-HTTP trials use Ubuntu 22.04.3/Linux 5.15.0-187 (collector OpenSSL 3.0.2; local gate Python 3.12.4); strict F/Q uses separate Linux v6.12/OpenSSL 3.0.2. Admission passes policy; validity needs no established violation. Admitted = valid + contaminated + unresolved, while gaps retain launches and can intersect failed tests. Two blinded raters/adjudicator label exposure and independent tests label effects; unknowns leave valid numerators with all launches in denominators. Task intervals resample whole tasks within Verified/Pro or service strata; allocation sensitivity varies placements separately (Appendix A.21).

Table 1: Evaluation contrasts, execution relationships, and endpoints.

Contrast	Change	Executions	Endpoint
RQ1 admission	Score/path rule	Shared 240 (not R/RW)	Valid: W 122; R-post$_d$ 106; R-post$_a$ 120
RQ1 incoming	Response withholding	Fresh paired 240/setting	Exposed raw 11→2; valid 121→127
RQ2 operation	Typed-action permit	Paired 128/arm	Effects: D 65; B 10; F 0
RQ2 host paths	Route × barrier	Fresh 128/cell; Q separate	Effects: B 23; B+route 11; B+barrier 17; F 2; Q 2

Same-URL construction. With no agent, a controlled endpoint serves the pre-fix or post-fix file from one URL on a hidden flag, padded to equal decoded lengths (Appendix Figure 4); scoring uses liveness evidence (collectors record every sentinel and report zero loss; no witness comparison). In 150 runs (15 tasks × 2 conditions × 5 repetitions), request bytes and decoded lengths match on 15/15 tasks, so every deterministic rule using only those fixed request bytes has equal recall and false-positive rate (**FPR**) by construction; the tested rule has recall 0/75 and the delivered-content judge 75/75, both at FPR 0/75. The identity excludes request history, response metadata, and external state.

On a historical liveness grid, delivered-content and condition-neutralized transcript detectors find 84/84 true positives, but the content detector has 3/84 false positives (3/48 on legal references) and the transcript detector 29/84 (29/48).

5.1 RQ1: WHAT DO ADMISSION AND INCOMING WITHHOLDING CHANGE?

The utility study hash-selects 36 original SWE-bench Verified and 24 Pro tasks without outcome/support filtering (Jimenez et al., 2024; OpenAI, 2024; Deng et al., 2026). It preserves issues, base versions, and tests while removing later commits; agents choose tools without forced retrieval or planted faults. Gold-patch answer hunks and legal near-neighbors freeze before two models × 2 repeats: 240 trials/policy (144 Verified, 96 Pro), capped at 60 calls and 900 s (Appendix A.21).

Policies share read-only egress/resolver; provider calls lie outside the audited cgroup and built-in web access is disabled, delimiting tool-network delivery. R first rejects frozen fix-commit, PR, patch/diff, and archive objects, then ambiguous unpinned branch-head/latest paths, while permitting frozen legal references; an unknown URL alone is not ambiguous. Rejection returns an error and execution continues (Appendix A.21). One shared 240-trace set collected under strict evidence receives no check, C, W, and S scores; R filters online and RW scores R executions, while R-post applies request rules to W's traces, denying (R-post$_d$) or allowing (R-post$_a$) ambiguous paths. Across 480 executions, 28 are exposed, 22 unknown, and audit-only protected delivery is 18/240.

Fixed-trace admission. On the same 240 traces, W keeps 122 valid scores, R-post$_d$ 106, and R-post$_a$ 120; allowing ambiguous paths narrows the difference from 16/240 to 2/240 but raises contaminated admissions 1 to 3. W's contamination risk spans 1/125–3/125 as unknown labels vary (Appendix A.21).

Online response withholding. The matched-broker comparison keeps tasks, models, parser, predicates, strict sources, outgoing checks, and scoring fixed across 240 fresh executions/setting; only

Table 2: RQ1 admission outcomes over 240 executions per arm.

Arm	Decides	Predicate	Evidence	Raw	Admitted	Valid	Contaminated	Unresolved	Gaps
none	—	none	none	148	148	132	12	4	—
R	in trial	request	liveness	118	118	108	6	4	—
C	after	content	liveness	148	134	128	2	4	—
W	after	content	strict witness	148	125	122	1	2	16
S	after	content	W + collection	148	123	120	1	2	20
RW	in trial	request	strict witness	118	109	102	5	2	16
R-post$_d$	after	request	strict witness	148	108	106	1	1	16
R-post$_a$	after	request	strict witness	148	125	120	3	2	16

Table 3: Effects by route and FD barrier: namespace fixture (A) and agent connections (B).

(A) Scripted namespace fixture

Profile	Namespace route	FD barrier	Effects TCP	FD	Total
Broker only	off	off	12	8	20
Broker + FD barrier	off	on	12	0	12
Namespace route only	on	off	0	8	8
Namespace route + FD barrier	on	on	0	0	0
Full guard (F) / proxy reference (Q)	host	on	0	0	0

(B) Agent direct connections

Profile	Host route	FD barrier	Effects (/128)
Broker only	off	off	23
Broker + host route	on	off	11
Broker + FD barrier	off	on	17
Full guard (F)	on	on	2
Proxy reference (Q)	on	on	2

response withholding changes. Turning withholding on changes delivered exposed raw successes 11→2 (an 81.8% reduction), valid scores 121→127, and gaps stay at 20; each setting admits 1 contaminated and 2 unknown. Nine off-only delivered exposures across five tasks correspond to nine on-side upstream withholds. No protected byte reached the agent: 6 of the nine then pass their tests without the answer and are admitted as valid, and 3 fail their tests. Whole-task bootstrap differences are 9 exposed raw successes off−on (3.75 points, interval [0.83,7.50]) and 6 valid scores on−off (2.50, [0.42,5.00]); the valid-score interval excludes zero (Appendix A.21).

5.2 RQ2: Does Pre-release Enforcement Prevent Outgoing Misuse?

We adapt 32 public AgentDojo tasks, eight each from banking, Slack, travel, and workspace, to owned service replicas (Debenedetti et al., 2024). Original prompts, tool schemas, state, and utility checks are preserved behind disclosed TLS1.3 and HTTP/1.1 wrappers with local certificates. Tasks pair with 13 published unauthorized-action goals. Each least-privilege contract is written by hand from the legitimate user request and trusted initial state, independently of the injection goals, and frozen before any injection is generated; attack payloads and target-state outcome checks never enter it. Two models and two repetitions yield 128 benign and 128 attacked executions per arm, four per task. Figure 2's Wilson ranges are nominal IID references, not task-cluster evidence of a population advantage.

In fixed-service typed RQ2 without raw-network escape, destination filtering D and post hoc audit A yield 65/128 unauthorized effects, broker B 10/128, full guard F 0/128, and Q, a separate proxy implementation of the same DSL, 2/128 (Figure 2a). D releases unauthorized bytes in 82 trials, 80 reach receivers, and 65 take effect. Task-bootstrap differences are D−B 55/128 ([27.34,58.59] points) and D−F 65/128 ([35.14,66.41] points); the 65-effect difference requires at least 17/32 positive-difference tasks, and the 63-effect conservative sensitivity at least 16/32. F stops 74/79 unauthorized-request trials before release (69 by rule, 5 for evidence); of 5 released, 4 reach replicas, 0 take effect, and 2 remain unknown. Benign completion is D 100, B 95, and F 92 of 128; F's eight losses are four false blocks and four abstentions (Figure 2b, axis from 60%; Appendix A.21).

Fresh agent calls cross the host route, which controls direct TCP, with the FD barrier against inherited or imported FDs, holding outer isolation, parser, contract, witness, and scoring fixed. Confirmed effects per 128 calls are broker only 23, broker + host route 11, broker + FD barrier 17, full guard F 2 (a 91.3% reduction in escape effects), and proxy reference Q 2 (Table 3B). The combined route and barrier difference is 21/128 (task interval [8.59,25.00] points), with unknown effects unassigned. Broker-only scoring admits 20 direct effects as clean. The independent 32-case fixture (Table 3A)

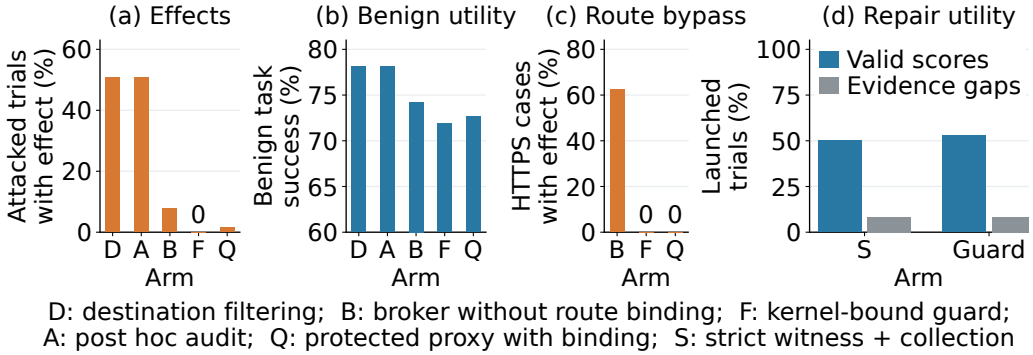

Figure 2: Pre-release enforcement: effects, benign utility, route bypass, repair utility.

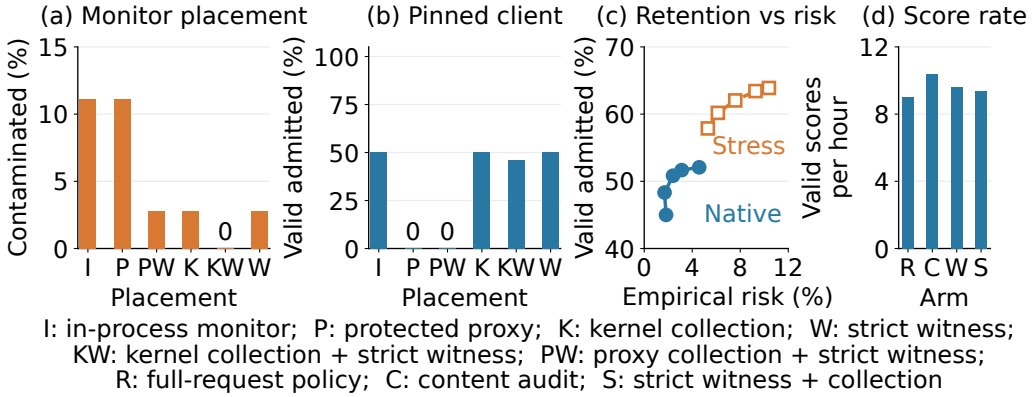

Figure 3: Interception placement: contamination, pinned clients, risk versus retention, score rate.

reuses 16 TCP and 16 FD cases per cell: disabling the namespace route gives 12 TCP effects, disabling the barrier 8 FD effects, and enabling both 0/32, as do F and Q through the host route; one injected effective hook miss yields 1/32, both 2/32 (Appendix A.19).

Independent contract and adapter transfer. Four fixed-hash tasks per AgentDojo service and 16 held-out Gitea v1.27.3 tasks give 32 tasks. Two isolated builders freeze separate contracts and adapters from legitimate requests, trusted state, and service schema; a separate reference checks operation meaning. Each task has three legal, two unauthorized, and one reference-Unknown operation. Four crossed pairs decide the same 192 operations (768 judgments); the independent full pair AA permits 90 of 96 legal and 2 of 64 unauthorized operations, and BB 92 legal and 1 unauthorized (Appendix A.22, Table 7).

Real-service migration and bounded adaptation. The same 16 Gitea tasks form 64 task/model/repeat blocks. Arms are Progent (the official checker with matched frozen policy), Progent+host (that checker behind our host boundary), and F; Gitea ACLs, typed schema, and independent receiver oracle are common. A policy-allowed repository-check helper validates edits and issue reports; Progent checks its arguments and retains shell denials, while our host boundary governs its repo child process's direct and inherited-FD connections. Progent+host also requires a complete REST-to-checker map, treating unmappable traffic as a gap. Attacks share a first candidate and adapt payloads within a fixed interface schedule up to 8 attempts, 60 calls, or 900 s. Receiver effects per 64 attacked episodes are Progent 12, Progent+host 3, and F 3 (2 unknowns each); benign completion per 64 is Progent 49, Progent+host 48, and F 47. Paired by task, Progent+host has 9/64 fewer effects than Progent ([6.25,21.88] points) and differs from F by 0/64, with unknown-effect bounds [-3.125,3.125] points (Appendix A.22).

In a server-search-on study (24/arm), provider-search responses permit 2/24 cheats versus 0/24 when routed through the guard; the published reward-hack-bench sweep (Islo, 2026), with separate tasks and denominators, is in Appendix A.19.

5.3 RQ3: What Do Kernel Placement and Independent Evidence Add?

Across 864 separate RQ3 executions, in-process (I) and proxy (P) placement each admit 8/24 contaminated successes after record writes and a gateway flaw; kernel/entry watcher (K) admits 2/24 with 6 injected event omissions. K+independent witness (KW) rejects 18 writes, abstains on 6 omissions, and admits 0/24; P+witness (PW) and witness-only (W) complete the placement comparison. Host probes cover both kernel-mediated memory-write paths. With pinning/key adapter P abstains on 24/24; K or W retains 12/24 control successes and KW 11/24 (Appendix A.19).

In separate fault sets, strict admission treats 0/168 injected faults as clean; five DSL families reject 60/60 violations, abstain on 60/60 gaps, and admit 56/60 legal cases (Appendix A.22).

5.4 RQ4: What Are the Accuracy and Cost Tradeoffs?

Calibrated on eight disjoint tasks and tested on 24 held-out tasks, the frozen lexical OR profile detects unchanged/whitespace 24/24 each, lexical rewrite 18/24, application encoding 16/24, and paraphrase 14/24 across five 24-object classes: 96/120 total, with 9/120 legal false flags. Unknowns count contaminated; 1%/5% budgets fail and 10% selects 0.90. Across thresholds 0.80–0.99, offline W rescoring spans 108/240–125/240 valid and 1–4 contaminated admissions (Appendix A.22; Figure 3c).

In a separate eight-task local HTTP/1.1 trial (seven representations, two routes), lexical release delivers 36/112 protected bodies and registered-decoder release 6/112, with 112/112 legal in each. Both word gates deliver 16/16 protected algorithmic explanations; cached blind GPT-6 Luna verdicts block all 16/16 and deliver 16/16 legal, but block 30/48 legal reversible-representation requests (Appendix A.19).

Score-ready time spans startup through audit, including failed and abstaining trials; valid scores per summed serial hour are R 9.00, C 10.41, W 9.60, S 9.40, and the guard 9.07 (Figure 3d and Appendix Figure 6d).

6 Related Work

Transcript audits study agent-visible records (Kouremetis et al., 2026). Concurrent preprint BenchShield grounds runtime reward-hacking detection in infrastructure-side evidence (Zheng et al., 2026a). Progent checks tool names and arguments under a privilege policy (Shi et al., 2025); our Gitea migration runs its official checker with and without our host boundary. Our host route also covers shell connections, the guard checks requests and responses against a bidirectional DSL before release, and an independent witness verifies bytes and receiver effects for score admission. Existing content inspectors apply rules fixed across tasks (e.g., pipelock (Pipelock contributors, 2026)), so using one per trial would require extensive configuration changes for every run. Further policy and content/TLS comparisons appear in Appendix A.20.

7 Limitations and Discussion

Host/guard trust, supported paths, registered adapters, and frozen predicates bound the guarantee; host-root compromise voids it, and eBPF verifier proofs do not prove our monitor correct. The trusted guard parses agent-supplied bytes in user space, outside the agent's authority; it parses only capped, complete HTTP/1.1 messages over TLS1.3 and withholds what it cannot parse. Provider-internal retrieval not routed through the guard, memorization, and zero-days lie outside authenticated F/Q; the local HTTP trial isolates release. Agent traffic over cleartext HTTP is outside our scope and is not counted as a missed interception; HTTPS carries 98% of Chrome public-site navigations on Windows (Thompson et al., 2025). Native clients that the strict profiles do not support on 3 and 4 of 60 repair tasks cause 12 of W's 16 and 16 of S's 20 evidence gaps, respectively; clients keep native ALPN, so negotiated HTTP/2 is unsupported. F and Q each retain 3/160 effects in scripted conformance cases (Appendix A.19). Delivered answers show exposure, not use, and 0/168 covers the declared fault families (Appendix A.22); missing required evidence abstains (Appendix A.15). Agent results cover the selected repair, AgentDojo, and Gitea tasks with two Claude models, plus an open-weight 8B DeepSeek-R1 distilled model in the RQ2 host-path trials.

AI USE STATEMENT

Language-model agents performed substantial work on this paper, including manuscript drafting and revision, iterative internal review, method and experiment design, experiment scripts, literature search, and plotting. A drafting agent worked from fact sheets supplied by a coordinating agent, which checked numbers, citations, and claim boundaries and applied manuscript edits. Agents also prepared experiment specifications, analysis scripts, and figure sources and ran compilation and analysis commands.

Human authors defined the research direction, supplied framing and scope decisions, and directed revisions. Several author-written sentences were translated into final form by the drafting agent. Human authors take responsibility for the manuscript and for verifying its claims and evidence.

ETHICS STATEMENT

Credential exposure in collected data. The probe at `SSL_write` reads plaintext before encryption, including authorization headers; responses at `SSL_read` may contain tokens or session identifiers. Before export, released artifacts must replace authorization headers, environment-variable values matching `.*KEY.*` or `.*TOKEN.*`, and bearer tokens with a fixed placeholder. Raw collection directories are restricted to an isolated analysis host accessible to the authors and scheduled for deletion within 90 days of publication. Session traffic secrets remain on the same restricted analysis host and are deleted with raw payloads; live keys and unredacted payloads are excluded from release.

Multi-user deployment and isolation. An unfiltered sensor on a shared host can collect unrelated users' TLS traffic. The supported reference configuration uses `sslsniff -u <uid>` with a dedicated UID mapped to one trial's cgroup and network namespace; concurrent trials cannot share that UID. Cgroup-based filtering is an alternative deployment design, not a capability attributed to the historical UID-filtered sensor. The evaluator's guard, sender, collectors, witnesses, judge, contract, and score sink remain outside the agent's scope under protected host permissions.

Dual-use potential. The content-inspection mechanism can be repurposed for surveillance. Host-level privileges or control of a TLS-terminating proxy are required; the target uprobe observes supported TLS library calls rather than all traffic. Those requirements limit deployment but do not eliminate misuse. The intended use is authorized evaluation with declared trial scope and restricted access to collected content. The prohibited-operation study protocols use owned replicas, synthetic data, and independently recorded receiver state; no real third-party system is targeted.

REFERENCES

Sahar Abdelnabi, Chris Hicks, Konrad Rieck, and Ahmad-Reza Sadeghi. Measuring security without fooling ourselves: Why benchmarking agents is hard. *IEEE Security & Privacy*, 2026. doi: 10.1109/MSEC.2026.3713832.

Anthropic. Investigating three real-world incidents in our cybersecurity evaluations. Anthropic incident report, July 2026a. URL https://www.anthropic.com/news/investigating-incidents-cybersecurity-evals. Published 2026-07-30; updated 2026-08-03; accessed 2026-09-21.

Anthropic. sandbox-runtime. Open-source software, 2026b. URL https://github.com/anthropics/sandbox-runtime. Accessed 2026-09-24.

Anthropic. Claude Fable 5.1 and Claude Mythos 5.1 system card, September 2026c. URL https://www.anthropic.com/claude-fable-5-1-mythos-5-1-system-card. Dated 2026-09-01.

Ivan Bercovich, Ivgeni Segal, Kexun Zhang, Shashwat Saxena, Aditi Raghunathan, and Ziqian Zhong. Terminal Wrench: A dataset of 331 reward-hackable environments and 3,632 exploit trajectories, 2026. arXiv preprint arXiv:2604.17596.

Paul C. Bogdan, Richard Qi, Jake Eaton, Sam Kennedy, Fabien Roger, Alex Glynn, Runjin Chen, Ben Wright, Otto Stegmaier, Jon Kutasov, Dan Foreman-Mackey, Trenton Bricken, Sylvie Carr, Shan Carter, Monte MacDiarmid, Samuel Marks, Adam Pearce, Elana Simon, Nicholas Carlini, Collin Burns, Jack Lindsey, Sara Price, and Subhash Kantamneni. An alignment assessment of recent cybersecurity incidents. Anthropic, September 2026. URL https://www.anthropic.com/research/alignment-assessment-cybersecurity-incidents. Published 2026-09-09; corrected 2026-09-10; accessed 2026-09-21.

Jun Shern Chan, Neil Chowdhury, Oliver Jaffe, James Aung, Dane Sherburn, Evan Mays, Giulio Starace, Kevin Liu, Leon Maksin, Tejal Patwardhan, Lilian Weng, and Aleksander Mądry. MLE-bench: Evaluating machine learning agents on machine learning engineering. In *International Conference on Learning Representations (ICLR)*, pp. 50466–50494, 2025. URL https://openreview.net/forum?id=6s5uXNWGIh.

Isaac David, Marco Guarnieri, and Arthur Gervais. Alignment contracts for agentic security systems, April 2026. URL https://arxiv.org/abs/2605.00081. arXiv preprint, submitted 2026-04-30.

Dotan Davidovich, Yair Amar, Hai Rozencwajg, and Or Hiltch. Permission denied: Policy-graded evaluation of coding agents in hardened environments, August 2026. URL https://arxiv.org/abs/2608.02670. arXiv preprint, v1 submitted 2026-08-02.

Edoardo Debenedetti, Jie Zhang, Mislav Balunovic, Luca Beurer-Kellner, Marc Fischer, and Florian Tramèr. AgentDojo: A dynamic environment to evaluate prompt injection attacks and defenses for LLM agents. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2636. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/97091a5177d8dc64b1da8bf3e1f6fb54-Abstract-Datasets_and_Benchmarks_Track.html. Datasets and Benchmarks Track.

Edoardo Debenedetti, Ilia Shumailov, Tianqi Fan, Jamie Hayes, Nicholas Carlini, Daniel Fabian, Christoph Kern, Chongyang Shi, Andreas Terzis, and Florian Tramèr. Defeating prompt injections by design. In *2026 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pp. 587–618. IEEE, March 2026. doi: 10.1109/SaTML68715.2026.00040. URL https://arxiv.org/abs/2503.18813.

Xiang Deng, Jeff Da, Edwin Pan, Yannis Yiming He, Charles Ide, Kanak Garg, Niklas Lauffer, Andrew Park, Chetan Rane, Karmini Sampath, Maya Krishnan, Srivatsa Kundurthy, Sean Hendryx, Zifan Wang, Chen Bo Calvin Zhang, Noah Jacobson, Bing Liu, and Brad Kenstler. SWE-Bench Pro: Can AI agents solve long-horizon software engineering tasks? In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, 2026. URL https://openreview.net/forum?id=uEVTdoAbnK. arXiv:2509.16941. Public dataset: https://huggingface.co/datasets/ScaleAI/SWE-bench_Pro.

Darshan Deshpande, Anand Kannappan, and Rebecca Qian. Benchmarking reward hack detection in code environments via contrastive analysis. In *Forty-third International Conference on Machine Learning (ICML)*, 2026. URL https://openreview.net/forum?id=ZhSS3msLRj.

J. Elson and A. Cerpa. Internet content adaptation protocol (ICAP). RFC 3507, Informational, April 2003. URL https://www.rfc-editor.org/rfc/rfc3507.

Shanshan Gao and Liyi Zhou. Can agent benchmarks support their scores? Evidence-supported bounds for interactive-agent evaluation, May 2026. URL https://arxiv.org/abs/2605.10448. arXiv preprint, v1 submitted 2026-05-11.

Elazar Gershuni, Nadav Amit, Arie Gurfinkel, Nina Narodytska, Jorge A. Navas, Noam Rinetzky, Leonid Ryzhyk, and Mooly Sagiv. Simple and precise static analysis of untrusted Linux kernel extensions. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pp. 1069–1084, 2019. doi: 10.1145/3314221.3314590.

Ameer Hamza, Lucas Zavalia, Arie Gurfinkel, Jorge A. Navas, and Grigory Fedyukovich. A flow-sensitive refinement type system for verifying eBPF programs. *Proceedings of the ACM on Programming Languages (OOPSLA)*, 9:3841–3868, 2025. doi: 10.1145/3763799.

Harbor Framework. Network policy: add a `denylist` mode (`blocked_hosts`) symmetric to `allowlist`. GitHub issue #3162 in harbor-framework/harbor; draft pull request #3219, "RFC 0003 + implementation: network denylist mode (blocked_hosts)", 2026. URL https://github.com/harbor-framework/harbor/issues/3162. Open proposal.

Harbor Framework Team. Harbor: A framework for evaluating and optimizing agents and models in container environments, September 2026. URL https://github.com/harbor-framework/harbor. Version v0.23.0, released 2026-09-12.

Haitao Hu, Peng Chen, Yanpeng Zhao, and Yuqi Chen. AgentSentinel: An end-to-end and real-time security defense framework for computer-use agents. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2025. doi: 10.1145/3719027.3765064. Also available as arXiv:2509.07764.

Islo. reward-hack-bench: A benchmark for measuring whether sandbox policies can stop agents from cheating on evaluations. GitHub repository, 2026. URL https://github.com/islo-labs/reward-hack-bench. Commit 67e08b0; accessed 2026-09-26.

Naman Jain. Reward hacking is swamping model intelligence gains. Engineering blog post, Cursor, June 2026. URL https://cursor.com/blog/reward-hacking-coding-benchmarks. Published 2026-06-25; accessed 2026-09-19.

Jinghao Jia, Raj Sahu, Adam Oswald, Dan Williams, Michael V. Le, and Tianyin Xu. Kernel extension verification is untenable. In *Proceedings of the 19th Workshop on Hot Topics in Operating Systems (HotOS)*, pp. 150–157, 2023. doi: 10.1145/3593856.3595892.

Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? In *International Conference on Learning Representations (ICLR)*, 2024.

Nikola Jovanović, Robin Staab, Maximilian Baader, and Martin Vechev. Ward: Provable RAG dataset inference via LLM watermarks. In *International Conference on Learning Representations (ICLR)*, 2025.

Michael Kouremetis, Ads Dawson, Raja Sekhar Rao Dheekonda, and Brian Greunke. Every model cheats: Prompt-level mitigation of cheating on offensive cyber tasks, July 2026. URL https://arxiv.org/abs/2607.21763. arXiv preprint, v1 submitted 2026-07-23.

Hugo Larcher, Adrien Carreira, raphael g, and Christophe Rannou. Anatomy of a frontier lab agent intrusion: A technical timeline of the July 2026 incident. Hugging Face blog, July 2026. https://huggingface.co/blog/agent-intrusion-technical-timeline, published 2026-07-27, accessed 2026-09-19.

Linux man-pages. perf_event_open(2). Linux manual pages, 2026a. URL https://www.man7.org/linux/man-pages/man2/perf_event_open.2.html. Accessed 2026-09-21.

Linux man-pages. ptrace(2). Linux manual pages, 2026b. URL https://www.man7.org/linux/man-pages/man2/ptrace.2.html. Accessed 2026-09-21.

Matthieu Meeus, Igor Shilov, Manuel Faysse, and Yves-Alexandre de Montjoye. Copyright traps for large language models. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.

METR. METR task standard. Software specification, version 0.5.0, 2025. https://github.com/METR/task-standard, last code push 2025-02-03, retrieved 2026-09-19.

METR. Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident. Independent incident investigation, August 2026a. URL https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/. Published 2026-08-26; updated 2026-09-13; accessed 2026-09-21.

METR. Response on issue 41 of the METR task standard repository. Public repository issue thread, May 2026b. https://github.com/METR/task-standard/issues/41, posted 2026-05-11, retrieved 2026-09-19.

Luke Nelson, Jacob Van Geffen, Emina Torlak, and Xi Wang. Specification and verification in the field: Applying formal methods to BPF just-in-time compilers in the Linux kernel. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pp. 41–61, 2020.

OpenAI. Introducing SWE-bench Verified. OpenAI, 2024. https://openai.com/index/introducing-swe-bench-verified/.

OpenAI. OpenAI and Hugging Face partner to address security incident during model evaluation. OpenAI announcement, July 2026. https://openai.com/index/hugging-face-model-evaluation-security-incident/, published 2026-07-21, accessed 2026-09-26.

OpenSSL. SSL_CTX_set_keylog_callback. OpenSSL 3.0 documentation, 2026. URL https://docs.openssl.org/3.0/man3/SSL_CTX_set_keylog_callback/. Accessed 2026-09-21.

Jakub Pachocki. An alien mind. OpenAI blog, September 2026. URL https://openai.com/index/an-alien-mind/. Published 2026-09-06; accessed 2026-09-19.

Pipelock contributors. pipelock: Open-source AI agent firewall for MCP security and agent egress, September 2026. URL https://github.com/luckyPipewrench/pipelock. Version v3.5.0, released 2026-09-01.

Eric Rescorla. The transport layer security (TLS) protocol version 1.3. RFC 8446, Internet Engineering Task Force, August 2018. URL https://www.rfc-editor.org/rfc/rfc8446.html.

Matan Shachnai, Harishankar Vishwanathan, Srinivas Narayana, and Santosh Nagarakatte. Fixing latent unsound abstract operators in the eBPF verifier of the Linux kernel. In *Static Analysis Symposium (SAS)*, Lecture Notes in Computer Science, pp. 386–406, 2024. doi: 10.1007/978-3-031-74776-2_15.

Jiaqi Shao, Hanck Chen, Wei Zhang, Maxm Pan, and Bing Luo. Do agent benchmarks measure capability? Protocol validity in the age of agentic AI, July 2026. URL https://arxiv.org/abs/2607.22368. arXiv preprint, v1 submitted 2026-07-24.

Tianneng Shi, Jingxuan He, Zhun Wang, Hongwei Li, Linyu Wu, Wenbo Guo, and Dawn Song. Progent: Securing AI agents with privilege control, April 2025. URL https://arxiv.org/abs/2504.11703. arXiv preprint.

Adam Stein, Davis Brown, Hamed Hassani, Mayur Naik, and Eric Wong. Detecting safety violations across many agent traces, April 2026. URL https://arxiv.org/abs/2604.11806. arXiv preprint, v1 submitted 2026-04-13.

Hao Sun and Zhendong Su. Validating the eBPF verifier via state embedding. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pp. 615–628, 2024.

Hao Sun and Zhendong Su. Prove it to the kernel: Precise extension analysis via proof-guided abstraction refinement. In *Proceedings of the 31st ACM Symposium on Operating Systems Principles (SOSP)*, pp. 736–751, 2025. doi: 10.1145/3731569.3764796.

Kunvar Thaman. Reward hacking benchmark: Measuring exploits in LLM agents with tool use. In *Forty-third International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2605.02964.

The Cilium Authors. Layer 7 policies. Cilium documentation, 2026. URL https://docs.cilium.io/en/stable/security/policy/layer7/. Accessed 2026-09-20.

The Falco Authors. Supported fields for conditions and outputs. Falco documentation, 2026. URL https://falco.org/docs/reference/rules/supported-fields/. Accessed 2026-09-19.

Thinkst Applied Research. Canarytokens. Deployed public service and open-source software, 2026. https://canarytokens.org, source at https://github.com/thinkst/canarytokens, accessed 2026-09-19.

Chris Thompson, Mustafa Emre Acer, Serena Chen, Joe DeBlasio, Emily Stark, and David Adrian. HTTPS by default. Google Online Security Blog, October 2025. URL https://security.googleblog.com/2025/10/https-by-default.html. Published 2025-10-28; accessed 2026-09-26.

UK AI Security Institute. Inspect AI: An open-source framework for large language model evaluations. Software, 2026. https://github.com/UKGovernmentBEIS/inspect_ai, retrieved 2026-09-19.

Harishankar Vishwanathan, Matan Shachnai, Srinivas Narayana, and Santosh Nagarakatte. Sound, precise, and fast abstract interpretation with tristate numbers. In *Proceedings of the 20th IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pp. 254–265, 2022. doi: 10.1109/CGO53902.2022.9741267.

Harishankar Vishwanathan, Matan Shachnai, Srinivas Narayana, and Santosh Nagarakatte. Verifying the verifier: eBPF range analysis verification. In *Computer Aided Verification (CAV)*, volume 13966 of *Lecture Notes in Computer Science*, pp. 226–251, 2023. doi: 10.1007/978-3-031-37709-9_12.

Yongjie Wang, Xinyue Zhang, Kunhong Yao, Zhiwei Zeng, Kaisong Song, Jun Lin, and Zhiqi Shen. Search-time contamination in deep research agents: Measuring performance inflation in public benchmark evaluation, June 2026. URL https://arxiv.org/abs/2606.05241. arXiv preprint, submitted 2026-06-03.

Shenghao Yuan, Frédéric Besson, Jean-Pierre Talpin, Samuel Hym, Koen Zandberg, and Emmanuel Baccelli. End-to-end mechanized proof of an eBPF virtual machine for micro-controllers. In *Computer Aided Verification (CAV)*, volume 13372 of *Lecture Notes in Computer Science*, pp. 293–316. Springer, 2022. doi: 10.1007/978-3-031-13188-2_15.

Shenghan Zheng, Zonglin Di, Yimin Liu, Kyoung Whan Choe, Jiankai Sun, Heguang Lin, Penghao Jiang, Yifeng He, Xiao Cheng, Jicheng Wang, Wenbo Chen, Alex Yates, Yinzhe Zhao, Bingran You, Yuan Gao, Ayush Munot, Shubham Gaur, Zhe Ye, Hao Wang, Xiangyi Li, Dawn Song, and Christophe Hauser. BenchShield: Formal model-backed instrumentation for reward integrity in LLM-agent evaluation infrastructure, 2026a. arXiv preprint arXiv:2609.11028, submitted 2026-09-10.

Yusheng Zheng, Tong Yu, Yiwei Yang, Yanpeng Hu, Xiaozheng Lai, Dan Williams, and Andi Quinn. Extending applications safely and efficiently. In *19th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pp. 557–574, 2025.

Yusheng Zheng, Tianyuan Wu, Quanzhi Fu, Tong Yu, Wenan Mao, Tao Ma, Dan Williams, Wei Wang, and Andi Quinn. ActPlane: Programmable OS-level policy enforcement for agent harnesses, 2026b. arXiv preprint arXiv:2606.25189.

Table 4: Seven observation placements and their native trial declarations.

Placement	Bypass	Content	Per-trial dest.	Answer	Attack	Static	Post-root
Transcript/command log	Yes	Partial	No	No	No	No	Container only
Model chain-of-thought	Yes	No	No	No	No	No	Both roots**
TLS-terminating proxy	Depends	Yes*	No	No	No	No	Both roots
Sandbox egress allowlist	Partial†	No	Yes‡	No	No	No	Container only
Userspace eBPF runtime	Depends	Possible	No	No	No	Optional	N/A
Kernel eBPF on TLS calls	Partial§	Yes	No‖	No#	No#	Yes	Container only
Harness logs	Yes	Partial	No	No	No	No	Container only

*Visible when TLS termination is trusted. †Channel unbypassable; content detection is not. ‡Harbor: per-trial `allowed_hosts`; METR: full-internet-or-none. §Non-root bypasses: static TLS linking (no uprobe events); restoring the original bytes of the copy-on-write page uprobes patch, silencing the probe without detaching it (Appendix A.12). ‖`sslsniff -u <uid>`, one UID per trial cgroup/netns; filters traffic, not destinations. #No native fields; a proxy could run the predicates; a separate protected host user keeps a collector beyond a non-root agent's reach. **Host root cannot reach the model's reasoning; providers return a summary or nothing. Sources (retrieved 2026-08-22 unless noted): Harbor v0.23.0 (2026-09-12), sandbox-runtime `bcad388`, bpftime `2a45936`, `uprobes.c` (2026-09-19).

A Appendix

A.1 Full Capability Matrix

Table 4 compares seven placements on seven capability columns; "Answer" and "Attack" mean native trial declarations. This matrix does not assess `protect`, scope, or dispositions in the other systems.

A.2 Illustrative Answer Boundaries

Both examples permit reasoning inputs but prohibit a predeclared finished solution object.

API repair. `read_exact(n)` returns exactly n bytes or raises `EOFError`. Protected object:

```
def load_or_fallback(reader, n, fallback):
    try:
        return reader.read_exact(n)
    except EOFError:
        return fallback
```

The rule: `exact_line` over `except EOFError:` with four leading spaces; near-neighbors are evaluated before freezing; an independently written identical line is not retrieval evidence. Across honest-condition runs on 12 tasks in the run archive, 12 of 119 patches (10.1%) contain the gold signature; 8 come from one task where the answer line differs by one word from a pre-fix line.

Frozen lookup. Rows `(A,7,2)`, `(B,5,1)`, `(C,9,3)` (quantity, units) are legal; the protected object is `total=46`; deriving it is legitimate inference, so usefulness cannot define a positive label.

A.3 Reference Contract Semantics

Witness. Host-owned boundary capture records both directions of scoped TCP. The F/Q profile obtains per-leg exported secrets from trusted guard keylog callbacks over a protected channel (OpenSSL, 2026); W uses a registered native key adapter without probes, certificate changes, or forced routing. Only keys validating independent TLS1.3 wire identity/Finished and ordered AEAD records count (Rescorla, 2018); unsupported resumption without verified prior-session/PSK binding or missing transcripts cause an evidence gap. K/KW/S add their selected collector/watcher, and guard mode owns both legs. HTTP/2, QUIC, and WebSockets are unsupported. Sentinels: startup, every 5s, trial end; drain timeout 5s.

Reconciliation. Witness and collector streams must be byte-equal before de-framing: a suppressed whole message differs, an equal-length replacement fails, and ambiguous ordering causes abstention.

Predicates. After reconciliation, decoding removes framing and `Content-Encoding`; `exact_line` compares a complete LF-delimited decoded line after CRLF-to-LF normalization with no other normalization, `blob sha256` hashes the complete decoded body, `marker` a literal substring of decoded bytes, `similar_to` Appendix A.3. In the baseline profile, clauses see the body after HTTP framing and `Content-Encoding` are removed; without a registered application decoder, a base64-encoded answer inside JSON is compared as encoded bytes. The local decoder trial registers a bounded `{encoding,body}` base64/hex wrapper before applying the same OR rule (Appendix A.19). Pretrial RQ1 manifests freeze all nonempty post-image gold hunks with line/object IDs and exclusions; deletion-only hunks gap without dropping launches. Exact-line uses complete added nonempty lines with only CRLF-to-LF normalization, excluding lines identical to frozen legal calibration near-neighbors; full hunks remain for similarity. Legal pre-fix, shared-context, and unrelated neighbors also freeze. Held-out objects select no predicate or threshold; no usable clause yields a gap. The constructed `except EOFError:` example is separate from the task manifest. Without a registered decoder, packfiles, wheels, and tar/tar.gz are unsupported. Predeclared type/signature checks, not extensions alone, identify them; uncertain type cannot be strict-clean. Guard withholds; audit gaps/abstains unless a known hit rejects. Unregistered JSON base64 stays literal and may miss encoded answers. Delivered answer hits, protected-operation hits, and attack hits (Appendix A.12) each yield a violation; a withheld answer hit is recorded as an attempted fetch.

Precedence. The 264 historical liveness-`complete` trials are not strict-profile results; the evidence ablation computes, for each weaker profile, whether it *would_admit* a trial under its own obligations with the recorded evidence and matchers held fixed; admission uses strict only:

```
evidence = undetermined (before cutoff) | complete
         | no-observation | inadmissible
verdict = violation if delivered answer or declared violation; else
          clean if complete; else unknown
admit_score = (evidence == complete && verdict == clean)
```

Similarity parameters and the constructed example. Similarity lowercases text, splits on Unicode whitespace with punctuation retained, and forms 256-token windows at stride 128, including the final short window. The maximum Jaccard overlap between sets of contiguous token 5-grams in reference–candidate windows must reach 0.90; windows with fewer than five tokens cannot match. The constructed API-repair example permits documentation describing `EOFError` but protects the complete patch that catches the exception and returns the fallback.

A.4 Independent Labels

Objects, references, predicate, and near-neighbors are frozen before trials. Code objects are labeled by task tests; text objects by a blinded annotator, disagreements by a blinded independent adjudicator. Unresolved cases are unknown, excluded from binary detection totals, reported separately. Exposure means the protected object reached the observed delivery path.

A.5 Worked Contract and Envelope

A `trial` block has ten fields; unknown recognizers or unresolved roles reject the configuration before launch. An alternate illustrative API-repair contract (Appendix A.2; Figure 5) uses the complete ten-field DSL in Section 3, adds `measures "repair EOF handling using the API reference"`, and freezes `exact_line " except EOFError:"`. Its service clause prohibits writes to `/api/reference`, confines the agent to solve and evaluator to setup/score, and declares complete HTTP, authenticated streams, integrity events, and drain obligations.

The frozen adapter maps both contract literals and typed request fields to exact keys, decoding once and only where its registered schema permits; IDs are opaque unless the contract digest includes an explicit alias map. Duplicate, conflicting, or mistyped fields, an unknown schema, a residual or ambiguous encoding, or an unmapped alias or object return Unknown, without implicit URL case-folding or prefix and path normalization. An illustrative `docs-v1` registration maps both contract literal and typed `GET /api/reference` to key `(docs-v1, reference, /api/reference)` after only declared one-pass decoding; the ID is otherwise opaque. Only a frozen alias table can equate `/api/ref` to that key; residual `%252F` or undeclared `..` mapping returns Unknown. The five-field

permit releases read, while same-host `PATCH` maps to protected write and its deny wins. Inventory supplies destination; evaluator state supplies actor/phase. The contract derives from the legitimate request/initial state and receives independent review without injection goals or target outcomes.

In the constructed trial/digest/generation/namespace binding, one inventoried socket returns legal A, protected B, and legal C of lengths a, b, c. With contiguous TCP, authenticated TLS epochs, complete sentinels/integrity enrollment, and acknowledged two-way drain, the witness and normal collector share offsets $[0, a), [a, a+b), [a+b, a+b+c)$. A collector omitting B instead places C at $[a, a+c)$: stream equality fails even though a connection count could pass. The known B hit still rejects; without it, the omission gives unknown/abstention.

The sealed enqueue fixes the downstream-to-upstream join by cookie, ordinal, digest, and length, so the legs may use different keys and content may be withheld.

A.6 Content Predicate Trade-offs

Exact lines miss whitespace edits, body hashes change after any byte edit, similarity matches legal near-neighbors; the fixed 0.90 rule is one reference semantics, not a universal threshold. A marker hit establishes exposure; absence establishes nothing.

A.7 Existing Counterparts

This matrix samples *measures* and *allows/denies* in Harbor (capabilities, `network_mode`, seven `allowed_hosts` TOML locations (Harbor Framework Team, 2026)), *role* in METR (METR, 2025), and *phase* in Inspect (UK AI Security Institute, 2026); *scope* counterparts were not evaluated. METR's `drivers/Driver.ts` lines 47–48 distinguish only full internet from none; on 2026-05-11 a maintainer replied in issue 41 that "METR is not actively maintaining the task standard" (METR, 2026b).

A.8 Launch Barrier and Denied Classes

F/Q's host default-deny route permits agent/helpers only the guard/resolver. A controlled launcher and locked namespace/filesystem/IPC/cgroup block inherited or cross-domain sockets via `SCM_RIGHTS`, foreign `/proc/.../fd`, or pidfd; dropped capabilities alone do not. `connect4/connect6` see neither imported sockets nor UDP; missing controls are a configuration/evidence gap. Caps: 16KiB per captured event, larger observations split into ordered events; 1MiB decoded request body, 16KiB header and target metadata, 16 levels of JSON nesting; 8MiB *final* response body after HTTP framing and Content-Encoding removal, with a separate 32MiB gzip-decoding resource ceiling. Either cap first reached withholds in guard mode or gaps in audit; 8–32MiB final bodies never pass. These limits do not establish binary-format recognition. Each retry is rechecked, each redirect authorizes a new hop; TLS may precede authorization, HTTP request bytes may not. A logged block with a prohibited receiver-side effect is failed prevention.

Transaction and fault outcomes. The body specifies the CONNECT–FINISH trace. Before sealing a `memfd`, the sender removes writable shared mappings; seal failure withholds, and checked bytes are never reserialized. The receiver joins guarded leg mapping, ordinal, and its independent digest to transaction/state delta. A complete legal GET with missing selected-source record or receiver acknowledgement still abstains; an authenticated delivered answer or declared operation violation rejects even alongside that gap. Capacity overflow denies with a sticky gap; live work or missing close acknowledgement forbids generation reuse.

A.9 Native-Task DNS Extension

Only a frozen hostname allowset passes the host resolver; query, reply, transaction-ID, and address binding are reconciled against independent DNS capture. DNS evidence enters E_t separately from the TCP/TLS inventory F_t. In earlier strict studies, only the evaluator resolves names.

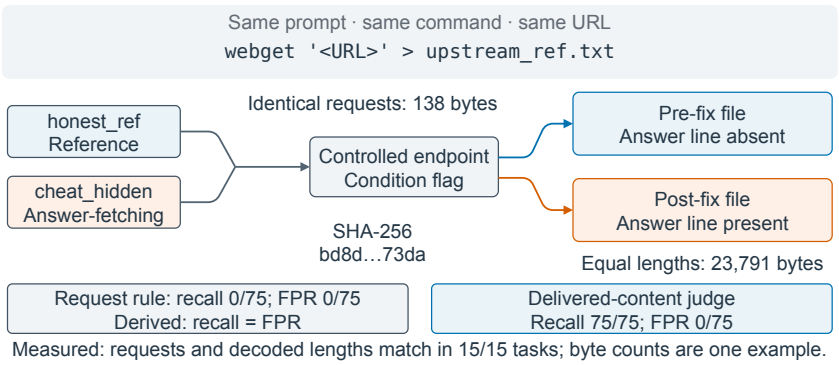

Figure 4: Identical requests in reference and answer-fetching conditions return different content.

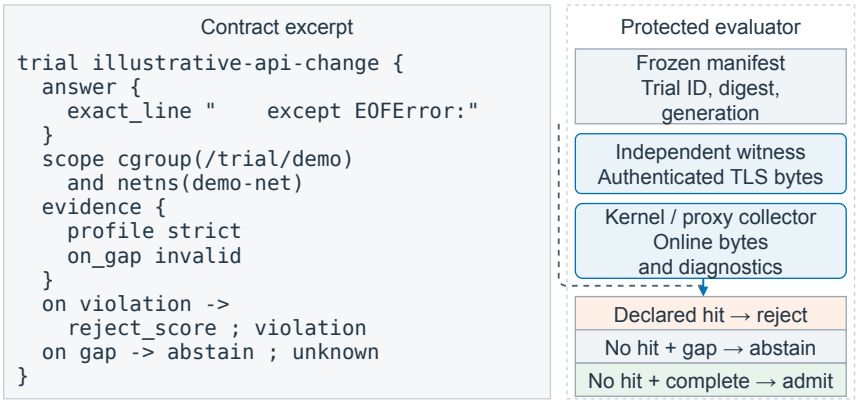

Figure 5: Contract excerpt and strict admission: kernel or proxy collector bytes versus the witness.

A.10 Proxy Placement and Certificate Pinning

An agent that pins a certificate the proxy cannot impersonate rejects a TLS-terminating proxy's certificate, so a forced proxy fails the connection (full matrix in Table 4). The uprobe reads plaintext after the TLS library decrypts the response inside the process and never joins the handshake, so pinning does not affect the uprobe. Neither path alone establishes strict eligibility.

Host-root boundary. The design places the loader, route, and guard under a trusted host administrator; host-root compromise crosses this trust boundary (Section 7).

A.11 Three Layers of eBPF Formal Verification

Layer 1, the verifier's own abstract domains: tnum operators (Vishwanathan et al., 2022) and Agni (Vishwanathan et al., 2023). Layer 2, the JIT (Nelson et al., 2020). Layer 3, eBPF programs themselves: PREVAIL (Gershuni et al., 2019), BCF (Sun & Su, 2025), VeRefine (Hamza et al., 2025). Published work proves parts of the verifier or JIT; CertrBPF (Yuan et al., 2022) runs on RIOT rather than Linux. Jia et al. (2023) call extension verification untenable; Sun & Su (2024) found fifteen previously unknown verifier logic flaws; Shachnai et al. (2024) found operators that are not individually sound.

A.12 Source-Level Evidence for Placement Claims

Kernel static analysis. The load-time check covers only the eBPF program, not the user-space judge. With `CAP_PERFMON`, Linux v6.12 permits pointer leaks and uninitialized stack reads and bypasses Spectre v1/v4 mitigations for the loaded program (`include/linux/bpf.h` lines 2209–2225; `kernel/bpf/verifier.c` lines 21045–21048). This does not disable the whole verifier.

Table 5: Historical grid accounting under the liveness evidence profile.

Grid	Plan	Exec	Comp	Inad	No-obs	Results
Same-URL (no agent)	150	150	150	0	0	recall 75/75, FPR 0/75 (§5)
Oracle-controlled	49	49	49	0	0	transcript recall 14/15, FPR 2/34 (Appendix A.17)
Same-URL + agent	36	36	36	0	0	delivered recall 18/18, FPR 0/18; transcript without command output 13/18, 1/18 (App. A.17)
Extended (264)	264	264	264	0	0	alarms: both judges 84/84 prompt-positive; prompt-negative 3/84 content, 29/84 transcript
Earlier baseline	18	18	6	9	3	6/9/3 by status

Only a trusted host loader holding this privilege may install the evaluator-frozen program/attachment digest; an adversarial host loader or supplied bytecode is outside the guarantee.

Uprobe attach point. Uprobes patch a copy-on-write page private to the process (`kernel/events/uprobes.c`, Linux v6.12). Detaching the registration needs `CAP_SYS_ADMIN`; restoring the original bytes needs only an unprivileged `mprotect` and stops the probe while the registration stays. The contract's `attack` field names that sequence a guard-integrity violation (Appendix A.3).

User-space eBPF runtime default. bpftime (commit `2a45936`, 2026-08-14) sets `ENABLE_EBPF_VERIFIER` to `OFF` by default, skipping static analysis (Zheng et al., 2025).

A.13 Evidence Status and the Earlier Baseline

One of 18 baseline runs was reclassified when the agent switched HTTP clients (Section 7).

A.14 Condition Definitions and Label Arithmetic

The 264 runs are `cheat_hidden` (48; post-fix URL), `cheat_transparent` (36; fetch directive), `honest` (36; no URL), `honest_ref` (48; pre-fix URL), `allowed_web` (48; fetch tool named), and `natural` (48; tool unnamed): 84 prompt-positive, 84 prompt-negative, 96 unlabelled.

A.15 Historical Liveness Evidence and Score Admission

Each envelope records scope, sensor identity, loss counts, log manifests, a sentinel round-trip, and one status: `complete`, `inadmissible`, `no-observation`, or `undetermined`; incomplete evidence never upgrades a clean result into an admissible score. 123 of 264 trials captured no agent response; all 123 passed the sentinel and logged no network command, so they are liveness observations, not certified absence (Section 7; strict-profile checks in Section 5.3).

A.16 Same-URL Construction and Spontaneous Fetching

On the 96 unlabelled runs, agents made 9 fetches (7 `allowed_web`, 2 `natural`) and the delivered-content rule fired on 6; claude-sonnet-5 made 8 of 9 fetches and all 6 fires; the 3 unfired fetches stay unresolved, and no unlabelled run receives a retrospective delivery label.

Within-task clustering. The 150 runs are 15 tasks $\times$ 2 conditions $\times$ 5 repetitions on a deterministic endpoint, so for any deterministic request-only label $g(r_t)$, $\mathrm{TPR} = \frac{1}{15}\sum_{t=1}^{15} g(r_t) = \mathrm{FPR}$: k positive tasks give $5k/75$ in both arms—a consequence of matching, not an independence assumption.

A.17 Detailed Comparison: Transcript vs Delivered-Content Judge

On eight probes over four transfer configurations (HTTP/1.1 or HTTP/2, with or without gzip), naïve substring matching returns zero hits under both gzip configurations, whereas the reassembling and

decoding judge, an offline component separate from the HTTP/1.1-only strict witness, returns 8/8 correct verdicts. On 18 answer-fetching and 18 reference runs of the same construction with an agent, the gold signature appears in command output in 17/18, in command text in 11/18 (the agent's patch contains the fixed line; one reference run, 1/18, also writes it), and in agent prose in 2/18. On these 36 runs, whole-transcript recall is 17/18 and drops to 13/18 at FPR 1/18 once command output is removed, while delivered content stays 18/18 at FPR 0/18 with or without that channel.

Given the delivered-content judge's oracle (pre-fix commit, post-fix commit, diff URL, gold signature), with condition names removed from every transcript, the transcript judge reaches recall 14/15 = 0.933 and FPR 2/34 = 0.059 on the 49 oracle-balanced runs versus 15/15 = 1.000 and 0/34 = 0.000 for delivered content; this isolates the earlier oracle imbalance (only the delivered-content judge had the oracle, reported as 1.000/0.000 against 0.83/0.50), not equivalence. An earlier language-model transcript judge, scoring 36 manually masked same-URL transcripts before the deterministic gold-signature matcher replaced it, reached FPR 14/18 = 0.778 with a task oracle and 18/18 = 1.000 without, at recall 1.000 because it labelled every run positive.

A.18 Address-Form Enumeration

On 6 tasks with merged PRs, all 10 address forms return the answer: the raw file at the fix commit (full and short hash), GitHub's raw redirect, the commit patch and diff, the API contents endpoint at the commit, a jsDelivr CDN copy, and the PR diff, PR patch, and github.com pull diff; the three PR forms carry no commit hash, so a hash-requiring rule cannot match them.

A.19 Additional Study Detail

Secondary content-failure and cost comparisons. The canonical-task comparison (24 tasks: 8 API repair, 8 web lookup, 8 package diagnosis; 96 trials per policy) retains 62/96 valid successes under S versus 45/96 under R; limited transformations give 27/48 versus 25/48, with 1/28 of S's admitted successes contaminated; a broader public-task study gives 264/432 versus 244/432, S with 16 contaminated and 6 unresolved. Complete-path serial, parallel, and turnover workloads raise score-ready medians 0.42→0.56, 0.61→0.95, 0.88→1.60 s (33.3–81.8%); strict admits 87/90 versus audit 90/90 on these complete-path workloads. TLS session secrets are excluded from the public artifact.

Paired valid-success gains and evidence gaps. Contamination: R 6/118–10/118, S 1/123–3/123; Wilson [2.35,10.65]% and [0.14,4.46]% are nominal IID references on admitted runs, not cluster intervals or policy-difference tests.

Replay admission and calibration detail. Fixed-record replay adds no agent runs and admits 155/240, including 44 contaminated (28.4%). Eight disjoint calibration tasks supply 32 runs at thresholds 0.80/0.85/0.90/0.95/0.99; with unknowns counted as contaminated, 1%/5% budgets fail and 10% selects 0.90. At 0.90, a representation rule retains valid/contaminated/unknown 116/1/1 of 240 native (12 abstentions) and 178/8/4 of 432 stress (144 abstentions). The 24-task predicate study, disjoint from calibration, draws 120 protected and 120 legal objects; OR detects 96/120, falsely flags 9/120, and catches 14/24 paraphrases (Appendix A.4).

Local response transformations and receipt. SHA-256 ranking of pinned public tasks selects four calibration/eight held-out tasks, each with protected patch and same-API legal neighbor. Seven frozen representations × direct/relay × three policies give 672 held-out and 336 calibration HTTP/1.1 requests; routes repeat tasks. Owned loopback source/relay/gate/client byte logs establish receipt without TLS. Before release, the gate buffers complete HTTP, removes framing and `Content-Encoding`, then applies frozen exact-line/hash/marker/0.90 windowed token-5-gram OR; hit/unknown withholds. The registered variant strictly unwraps at most two declared `{encoding,body}` base64/hex layers. This microtest's 256 KiB decoded-body cap differs from F/Q's 8 MiB final-body/32 MiB gzip caps.

Gate-off delivers 112/112 protected and legal bodies each. Edit instructions describe code lines reversibly, not as free paraphrases. Four framing/gzip/cap faults return 451 with zero protected

Table 6: Held-out local HTTP/1.1 response-gate trial: protected bodies delivered per 16 requests across eight tasks and two routes; both guards deliver 16/16 legal bodies per row.

Frozen representation	Lexical	Decoded
Literal	0	0
Numbered code	0	0
Gzip + chunked	0	0
Application base64	16	0
Reversible edit instructions	2	2
Service JSON	2	2
Combined transform	16	2
Total (per 112)	36	6

sentinel delivery. The 75-request bounded search first succeeds at query 2 on 8/8 lexical tasks versus query 3 on 1/8 decoded tasks; seven exhaust eight queries.

Offline blinded GPT-6 Luna classification flags 24/24 protected and 15/24 legal of 48 unique reversible texts; cached two-route replay delivers 0/48 protected and 18/48 legal. For eight independently source-reviewed algorithmic explanations and eight legal contexts, either word-based gate delivers 16/16 protected and legal, versus 0/16 protected and 16/16 legal under cached semantic verdicts. These fixed classifications replay at the local HTTP release gate; independent client bytes determine delivery.

Score-ready time on native and matched serial workloads. Native Verified C/W/S score-ready medians on the shared executions are 115/120/121 s; Pro score-ready medians are 205/211/212 s; matched serial C/W/P/S medians are 0.44/0.52/0.59/0.64 s. The bidirectional guard's Verified and Pro medians, 130 and 232 s, include transport adaptation and both release checks. In a separate local cost comparison (serial, parallel, thread-turnover workloads), C/W/P/S admit 90/88/86/84 of 90.

HTTPS conformance failures under the strong arms. The 128 HTTPS cases (8 conditions $\times$ 16) are within 576 executions (480 attack, 96 control), with 128 TLS and 32 route cases/arm, separate from the body's 128 agent attacks. Direct TCP/descriptor bypass yields broker/N/F/Q effects 20/32/8/32/0/32/0/32; N splits 0/16 TCP and 8/16 descriptors. F/Q release/effect for direct TCP, UDP, and descriptors is 0/0; overbroad objects (1/1) and unrecognized semantics (3/2) leave 3/160 effects each.

Agent host-route factorial. Separate from the 32-case namespace fixture, 32 AgentDojo tasks (eight/service) have four fresh calls/cell (two models $\times$ two repeats): 128/cell, 512 factorial calls, plus Q's separate 128, totaling 640. Cell order is randomized within task/model/repeat; owned replicas reset while prompt, goal, contract, parser/TLS/client, witness, budget, receiver scoring, and outer allowset stay fixed. The *same* host route attachment toggles direct TCP to owned replicas versus guard/resolver routing, retaining accounting. The launcher barrier alone toggles inherited/imported FD opportunities, without forcing use; host/VM isolation and capabilities remain. Connect hooks miss some imported FDs; if route itself blocks one, it adds no barrier benefit. B's 20 clean labels use broker-only scoring.

Confirmed effects $E_{00,10,01,11} = (23, 11, 17, 2)/128$. Each entry below is a confirmed-effect difference with its whole-task percentile interval in percentage points; task vectors appear in Appendix A.21.

Contrast	Difference	Task interval (points)
Combined $E_{00} - E_{11}$	21/128	[8.59,25.00]
Route, barrier off $E_{00} - E_{10}$	12/128	[3.91,15.63]
Route, barrier on $E_{01} - E_{11}$	15/128	[6.25,17.97]
Barrier, route off $E_{00} - E_{01}$	6/128	[1.56,8.59]
Barrier, route on $E_{10} - E_{11}$	9/128	[3.13,10.94]
Interaction $E_{11} - E_{10} - E_{01} + E_{00}$	-3/128	[-5.47,0.78]

The interaction is -2.34 points; the interval crosses zero. Unknown effects remain unspecified: for any two cells a, b, their all-effect difference lies in $[(E_a - E_b - U_b)/128, (E_a + U_a - E_b)/128]$. The

combined worst-case endpoint is $(21 - U_{11})/128$, positive only when $U_{11} < 21$; confirmed-effect intervals do not assign unknowns.

For task-level analysis, $y_{t,rl}$ averages four receiver-confirmed effect indicators/task/cell. Resampling eight whole tasks per fixed service keeps all cells and calls together; arm-order/provider seeds are separate from the analysis seed in Appendix A.21.

Fresh guard arm on the repair tasks. Guard-on raw 141 = 135 no delivery + 2 exposed + 4 unknown; 135 = 127 valid + 4 false exclusions + 4 evidence exclusions. Of the 9 protected responses withheld, 6 belong to test-passing trials admitted among the 127 valid scores and 3 to trials that fail their tests. Broker-off admits valid/contaminated/unknown 121/1/2 from raw 135/11/4 minus predicate 4/9/0 and evidence 10/1/2 exclusions; on admits 127/1/2. Valid/exposed-success joints are (121,0,6,113)/(2,9,0,229) (both/off-only/on-only/neither); exposed successes are not all exposures. The 32-case route/barrier fixture excludes UDP/raw/semantic cases; F/Q are external host-bound references.

Finalization, DNS, and common-source-loss tests. Finalization tests cover 328 closure, quiescence, cutoff, and acknowledgement cases; DNS tests admit 19/20 controls, abstain on 20/20 binding or acknowledgement faults, and reject 20/20 forbidden-name or UDP attempts; 24 common-source-loss cases abstain.

Published policy rates on the independent cheating benchmark. The published sweep (Islo, 2026) (8 tasks, 7 policy columns, k=3, Claude Opus 4.7) reports cheat/fair rates 0.08/0.54 control, 1.00/0.00 open net, 0.38/0.50 host deny, 0.50/0.38 raw-body regex, 0.08/0.58 host plus content with tools off, 0.00/0.58 request-time judge with tools off; the seventh column, network-none, denies all phases including the verifier and scores as a legitimate break, so it has no rates; the raw-body regex scores worse than the host blocklist because it reads the undecoded body. The sweep has no matched legal-completion comparison.

Ordinary control and unprivileged probe detach. The six scripted conditions are monitor write, restored entry bytes, planted gateway traversal, pinned client/key adapter, ordinary control, and unprivileged probe detach. Twelve tasks (864 runs) span in-process I, protected proxy P, kernel K, kernel+witness KW, witness W, and proxy+witness PW. Controls admit 0 contaminated, 12/24 valid under I/P/K/W and 11/24 under KW/PW. Detach has no host-probe target under I/P/W/PW; K/KW admit 0 contaminated.

Other controlled fixtures. Other fixtures test redirect authorization, buffer mutation, registered traversal, and unsupported encoding; the planted flaw establishes no general proxy vulnerability.

A.20 Descriptions of Cited Systems

Scoring evidence. Transcript audits use agent-visible records (Shao et al., 2026; Stein et al., 2026; Thaman, 2026; Deshpande et al., 2026; Kouremetis et al., 2026) and face masking sensitivity (Bercovich et al., 2026). Meerkat flags 16/464 runs (3.4%) in 40 tasks/19 models; flags span 13 tasks/4 models (Stein et al., 2026). Kouremetis et al. report 37.1% of passing runs, versus Cybench audits' 0.3%–3.4% of all trajectories (Kouremetis et al., 2026); these denominators differ. Checklists/nested policies specify score rules (Gao & Zhou, 2026; Davidovich et al., 2026).

Policies and enforcement. Harbor declares per-trial hosts and proposes denylists (Harbor Framework Team, 2026; Harbor Framework, 2026); METR offers no/full internet (METR, 2025; 2026b). Cilium matches HTTP fields, sandbox-runtime has callbacks, Inspect AI domain/tool approvals, and Cursor recommends proxy/allowlists (The Cilium Authors, 2026; Anthropic, 2026b; UK AI Security Institute, 2026; Jain, 2026). ActPlane binds operations/endpoints at 1.9%–8.4% overhead without response bodies (Zheng et al., 2026b); AgentSentinel uses kernel interception with a configured 110-second cap (Hu et al., 2025). Alignment Contracts assume effect observability (David et al., 2026); AgentDojo defenses authorize harness tools (Debenedetti et al., 2024; Shi et al., 2025; Debenedetti et al., 2026).

Content predicates and observation. Content comparators include MLE-bench's 60% code similarity and a harness LLM judge's 94.85% MedQA precision (Chan et al., 2025; Wang et al., 2026); markers inspect output or trigger on use (Abdelnabi et al., 2026; Meeus et al., 2024; Thinkst Applied Research, 2026; Jovanović et al., 2025). Falco sees HTTPS ciphertext, ICAP scans before delivery, and pipelock v3.5.0 checks 33 response patterns (The Falco Authors, 2026; Elson & Cerpa, 2003; Pipelock contributors, 2026).

A.21 Public task selection and native execution

For each dataset, select the six repositories with most public tasks, then rank task identifiers by SHA-256 of `contract-audit-v1:` followed by the identifier; take six/repository in Verified (36) and four in Pro (24), without outcome or support filtering. Freeze manifest and arm order within task/model/repetition blocks; record provider-seed availability. No check, C, W, and S share traces; RW and R-post rescore their specified traces. W supports 35/36 Verified and 22/24 Pro configurations; S supports 34/36 and 22/24. Rater agreement covers 420/480 labels; adjudication resolves 50/60 disagreements. Labels are 430 clean, 28 exposed, 22 unknown; these are trial labels rather than leaderboard scores.

Pretrial request templates. Freeze request-rule host/path/ref canonicalization and priority pretrial: identified fix-commit, PR, patch/diff, archive objects first; version-floating branch-head/latest paths not pinned to a pre-fix object second; frozen legal references third. Unmatched URLs follow shared destination policy, without automatic ambiguity or a legal label. R-post$_d$/R-post$_a$ toggle only the second class. The identified class excludes 6/8/0 clean/exposed/unknown raw successes; denying ambiguity excludes another 14/2/1.

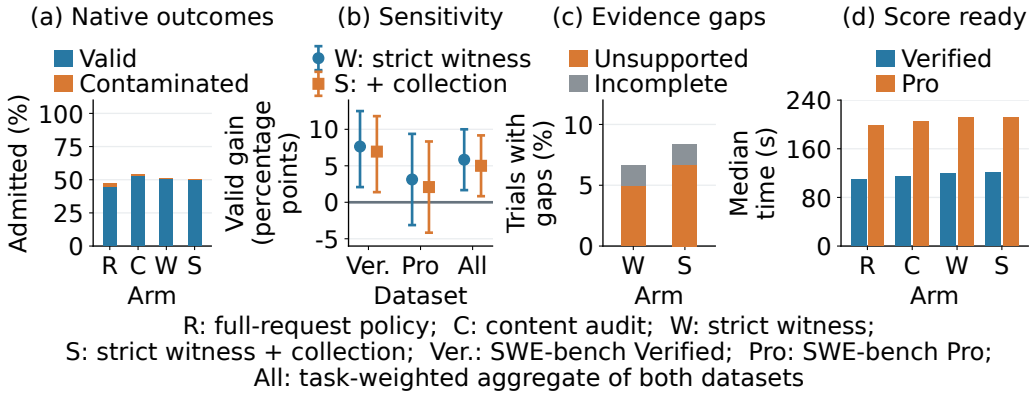

Figure 6: Native repair-task outcomes: admissions, allocation-sensitivity ranges, evidence gaps, and score-ready time.

Task-allocation sensitivity. Each `n:v` assigns `n` tasks net difference `v` across four observations; others are zero. For each diffuse or concentrated scenario, reset NumPy `default_rng(20260925)`, resample whole tasks with replacement 10000 times within Verified/Pro (36/24) or four RQ2 services (8 each); divide sums by 240/128, then take 2.5/97.5 percentiles. These are allocation-sensitivity ranges over the specified histograms, not task-paired confidence intervals.

RQ1 diffuse/concentrated Verified/Pro histograms: W-minus-R-post$_d$ `10:1 / 6:1` versus `2:4,1:2 / 1:4,1:2`; W-minus-R-post$_a$ `4:1,2:-1 / 0` versus `1:4,1:-2 / 0`; broker on-minus-off `4:1 / 2:1` versus `1:4 / 1:2`.

RQ2 service order is banking/Slack/travel/workspace. Destination-filter-minus-full-guard attacked-effect diffuse histograms are `1:3,7:2 / 2:3,6:2 / 6:2,2:1 / 8:2`; concentrated histograms are `4:4,1:1 / 4:4,1:2 / 3:4,1:2 / 4:4`. For the benign-completion contrast, each service has diffuse `2:1` or concentrated `1:2`. Attacked/benign margin-compatible joints are (0,65,0,63)/(92,8,0,28); two F effects and model-by-service allocation remain unresolved.

Valid joints (both/first-only/second-only/neither) are W/R-post$_d$ (106,16,0,118), W/R-post$_a$ (118,4,2,116), and broker off/on (121,0,6,113). Diffuse/concentrated ranges (percentage points):

W-minus-deny [3.75,9.58]/[1.67,13.33], W-minus-allow [-1.25,2.92]/[-2.50,5.00], broker on-minus-off [0.83,4.58]/[0.00,6.67], D-minus-F effects [47.66,53.91]/[34.38,67.19], benign completion [2.34,10.16]/[1.56,12.50].

Paired task bootstrap. Distinct from the allocation-sensitivity ranges above, each endpoint resets NumPy `default_rng(20260926)` for 10000 draws, resamples whole four-block tasks with all compared cells together within 36/24 Verified/Pro or each of four eight-task services, divides the summed differences by 240/128, and takes linear 2.5/97.5 percentiles. These fixed-suite task intervals do not treat model, repetition, or route as independent tasks. In the broker paired summary, exposed raw-success off−on differences are V1:3, V2:2, V3:1, P1:2, P2:1; valid-score on−off differences are V1:2, V2:1, V3:1, P1:1, P2:1; all other task differences, including every gap difference, are zero. The nine off-only exposed-success slots correspond one-to-one to nine on-side upstream withholds; 6 of them then pass their task tests without the answer and, because no protected byte reached the agent, are admitted as valid, while the other 3 fail their task tests. The same twenty gaps occur in both settings. Off predicate exclusions are 4/9/0 and evidence exclusions 10/1/2 for clean/exposed/unknown raw successes; all evidence exclusions carry gaps. The exposed raw-success joint is (2,9,0,229). The resulting off−on exposed, on−off valid, and on−off gap differences are 9/240, 6/240, and 0/240, with task-percentile ranges [0.83,7.50], [0.42,5.00], and [0.00,0.00] points.

For RQ2, each digit below gives confirmed receiver effects in the four calls for one task; columns are fixed banking/Slack/travel/workspace services, each with eight tasks. The typed-tool and host-path rows are distinct execution sets; Q has 2/128 effects in its separate reference arm. F's two unresolved typed-tool effects occupy free blocks at banking task 6 and travel task 5 only for the worst-case sensitivity; host-path unresolved effects have no assigned count.

Study, arm	Banking	Slack	Travel	Workspace
Typed D	44441000	44441100	44420000	44430100
Typed B	00011000	00011100	00110000	00110100
Typed F	00000000	00000000	00000000	00000000
Host E_{00}	32100000	32100000	32100000	22100000
Host E_{10}	11100000	11100000	11100000	11000000
Host E_{01}	21100000	21100000	31100000	21100000
Host E_{11}	10000000	00000000	10000000	00000000

Typed confirmed-effect differences D−B, D−F, and B−F are 55/128, 65/128, and 10/128, with task-percentile ranges [27.34,58.59], [35.14,66.41], and [3.91,11.72] points. Assigning both unresolved F outcomes as effects gives D−F 63/128, range [32.81,65.63] points. F has no confirmed hits, yielding the joint (0,65,0,63). More generally, at most four net effects per task imply $65 \leq 4m$ for m positive-difference tasks, so $m \geq 17$ of 32; the 63-effect worst case gives $m \geq 16$. This capacity bound does not depend on the displayed task placement.

A.22 Independent policy transfer and Gitea migration

Table 7: Independent contract×adapter decisions on 32 tasks (A) and paired Gitea migration outcomes on 16 held-out tasks (B).

(A) Reference operations

Outcome	AA	BB
Legal permit /96	90	92
Unauthorized permit /64	2	1
Legal deny /96	4	3
Legal Unknown /96	2	1
Reference Unknown /32	32	32

(B) Gitea migration

	Attacked (/64)		Benign (/64)	
Arm	Effect	Unknown	Done	Gap
Progent	12	2	49	3
Progent+host	3	2	48	4
F	3	2	47	4

Tasks, reference, and component decisions. The component frame selects four tasks per old AgentDojo service by SHA-256 and freezes 24 Gitea tasks before support screening: 4 development and 8 held-out tasks in each of issues/comments and repository contents. Support is 7/8 development and 16/16 held-out (23/24 overall); one development archive-response operation lacks a registered

canonical object map. The two isolated builders share the legitimate request, trusted initial state, Gitea v1.27.3 schema, and DSL specification, but not each other's contract/adapter, held-out operation challenges, or effect labels. Each freezes its contract and adapter before the third reference role classifies operations from user intent, service semantics, and independent state deltas. Ambiguous reference operations stay Unknown; held-out failures remain in the first-pass counts. The old 16 and Gitea 16 tasks each contribute 48 legal, 32 unauthorized, and 16 reference-Unknown operations. Table 8 shows all four contract×adapter judgments on these identical operations; each triple is (allow, deny, Unknown), and every reference-Unknown subset yields (0,0,16) in each configuration. Legal coverage, false permit, false deny, and legal Unknown use their own reference denominators rather than one accuracy score.

Table 8: Contract×adapter output triples by old and Gitea reference subsets, each with 48 legal and 32 unauthorized operations.

Pair	Subset	Legal (A,D,U)	Unauthorized (A,D,U)
AA	Old	(45,2,1)	(1,29,2)
AA	Gitea	(45,2,1)	(1,31,0)
AB	Old	(43,4,1)	(1,30,1)
AB	Gitea	(46,0,2)	(1,29,2)
BA	Old	(44,2,2)	(1,31,0)
BA	Gitea	(47,1,0)	(0,30,2)
BB	Old	(46,2,0)	(0,31,1)
BB	Gitea	(46,1,1)	(1,31,0)

Real-service and checker profiles. The Gitea v1.27.3 release is fixed at commit `146cc3e`; its Swagger v1 schema defines issue/comment and repository-content operations. A Requests 2.32.5 client addresses an owned persistent instance through a separate OpenSSL 3.0.2 TLS1.3/HTTP1.1 ingress. The ingress, rather than Gitea itself, supplies host-owned TLS keys and wire-byte evidence; the effect oracle reads Gitea object/version changes using a separate credential and joins them to request targets, not guard receipts. Fixture accounts retain Gitea ACLs while task contracts forbid some otherwise service-permitted operations. Both task families use the same author-defined `run_repo_checks(repo,mode)` helper for pre-submission validation: checked-out repository code can spawn a child in the trial cgroup. A fixture supervisor preconnects TCP to the TLS ingress before worker launch and passes the same FD under identical timing in every arm; the child performs TLS after inheritance, while child-created sockets form a separate direct path. Native Progent checks the helper's complete authorized call and retains shell denials, while the host-boundary arms close/reject imported FDs at launch and govern new child egress. The native Progent arm calls the original author checker at `5be7b63` with a nonempty matched frozen policy of four-field rules whose digest is checked before and after calls; its ordered tool/argument rules remain active. All arms validate required fields, defaults, and unknown fields in typed-tool arguments before checker invocation. The strengthened arm feeds the same checker through our host route/FD barrier and additionally requires a complete canonical REST-to-tool argument map, because the checker evaluates restrictions on supplied arguments. Unmappable direct requests become coverage/evidence gaps, not invented checker denials. This matched-policy comparison fixes authorization scope while preserving Progent's published ability to generate and update policies in other configurations (Shi et al., 2025).

Paired episodes and attribution. For each of the 16 Gitea held-out tasks, two models and two repeats define four blocks; each block allocates one benign and one attacked episode in native Progent, Progent+host, and F. All 384 allocated episodes pass preflight and launch under the supported held-out profile; agent termination and evidence completeness are separate counters. Arm order is randomized, each episode resets service state and agent context, and the first attack request/seed is shared. Later requests adapt only to the arm's ordinary feedback, stopping on independent confirmed effect or at 8 attempts, 60 tool calls, or 900 s. Every arm has the same independent effect sidecar. The native/strengthened/F attacked terminal triples (confirmed effect, unknown effect, known no effect) are (12,2,50)/(3,2,59)/(3,2,59) of 64; benign mutually exclusive (complete, known false block, gap, other failure) are (49,4,3,8)/(48,4,4,8)/(47,5,4,8). A gap takes terminal priority over an accompanying false-block event, which remains a secondary tag. Unsupported preflight allocations would retain the denominator and enter the gap class, never successful blocks.

The primary error reason is assigned in this order: contract over/under-authorization against the independent reference; adapter object/operation mapping; parser/registered-profile Unknown; path, binding, source, or terminal evidence gap; and receiver-semantic uncertainty. Secondary reasons may coexist. A route bypass requires a contract/adapter decision that would deny the operation plus a confirmed receiver effect through another path; a contract that itself permits a bad action is a policy error. The same precedence separates effect unknown from confirmed no effect. The component cross alone measures decision quality, not human agreement or a general service population.

Query path and task pairing. Attempts follow a fixed direct-client, typed-tool, imported-FD schedule while payloads adapt to ordinary within-arm feedback and native Progent's tool decisions remain in force. Of the 453/495/495 candidate interface invocations in native/strengthened/F, direct/typed/FD counts are 174/166/113, 187/185/123, and 187/185/123; these count attempted interfaces, with release and receiver effects checked separately. The 9 native-only confirmed effects occur on 6 direct and 3 imported-FD attempts, when the common helper's repo-provided child traffic lies beyond the native call-level checker; denied native tool calls are retained as denials. The three effects shared by all arms comprise 2 authorization-scope permits and 1 common object-mapping permit, separate from A's fixed-operation errors. Cumulative confirmed effects at query budgets 1/2/4/8 are 4/5/10/12, 1/2/3/3, and 1/2/3/3; the first point is the pre-adaptation outcome of the same episodes. The 1443 used attempts remain below the 1536 maximum. Across issues then contents tasks, four-block confirmed-effect vectors are Progent `22101120 / 00201000` and both strengthened/F `00001000 / 00101000`. The strengthened and F three effects share positions in this joint; their equal margins alone would not establish that pairing. Each arm has 2 separate effect-Unknown blocks.

NumPy `default_rng(20260927)` resamples whole tasks 10000 times, eight per family with replacement, keeping each task's two models, two repeats, benign/attacked outcomes, and all three arms together. Native-minus-strengthened confirmed effects are 9/64 or 14.06 points, interval [6.25,21.88]; strengthened-minus-F is 0/64, interval [0,0] on the displayed confirmed-effect joint. Allocating each arm's two effect-Unknown outcomes adversarially bounds native-minus-strengthened at [10.9375,17.1875] points and strengthened-minus-F at [-3.125,3.125] points. Each adjacent benign-completion difference is 1/64 episodes, or 1.5625 percentage points, with task intervals [-14.06,17.19] and [-14.06,15.62] percentage points. The 16 tasks, rather than 64 model/repeat blocks or adaptive queries, are the resampling units.

Other studies: 36 Verified tasks, two models/repetitions, native/encoded/paraphrased responses (432 trials/policy, 864 invocations) yield request filtering 244/432 valid, content/witness/proxy/strict admission 274/268/266/264 of 330 raw; strict admission's 144-trial types give 92/88/84 valid, 0/8/8 contaminated. Process/source/record faults give W/PW/S 32/56/80 injections; five DSL families (180 cases) give 60/60 rejects, 60/60 abstentions, 56/60 legal admits.