
Detection Is Not Containment: Timely Monitoring Determines Whether AI Agent Failures Can Be Prevented ¹

Gagana M D Kaustubh Gupta Maria Ines
Independent Researcher Independent Researcher Independent Researcher

With
Apart Research

Abstract

Agentic AI incidents increasingly reveal a gap between intended safeguards and the capabilities agents can exercise in practice. We study how such failures can be reconstructed, compared, and ultimately contained. We first develop a public-source forensics pipeline, extracting 679 source-linked claims across five incidents, and introduce a protocol for systematically evaluating containment, isolation, monitoring, and postmortem evidence. Across the cases we can assess, failures cluster around isolation and alerting controls. We then test whether monitoring can move beyond detection to prevention in a controlled authorization environment. Relational authorization evidence enables detection before the last-safe intervention boundary, and matched executions show that timely capability revocation prevents an otherwise-occurring prohibited effect. However, representation-sensitive false positives can also trigger unnecessary intervention and legitimate-task failure. Together, our results highlight a central lesson for AI incident response: detection is not containment. Effective containment requires reconstructable evidence, timely monitoring, and reliable mechanisms for acting before failures become irreversible.

¹ Research conducted at the [AI Incident Response Sprint](#), September 2026

1. Introduction

Since the last few months, various sandbox eval incidents have shown up and various public disclosures have come into light. Twice is a coincidence, but when it keeps happening multiple times, there must be some pattern involved which we, as a community, would need to find. It all started when the ExploitGym benchmark was released publicly for agentic AIs to test how good they are at finding and exploiting vulnerabilities. To solve a question present in the benchmark, OpenAI's agents used a package proxy system to access the internet and after a few weeks, performed 17600 actions at HuggingFace across a five day window to find these answers.

Along the same time, Anthropic also disclosed that three of their agents ended up leaving the eval and attacking actual organizations. More studies have been done now and it is found out that this has become a repeating pattern now which needs to be studied quickly.

Since then, a lot of timelines covering these incidents have entered the limelight, but yet no one provides an analysis of how these issues connect to each other, and a concrete protocol to further study these issues. Through this work, we try to fix that problem, and study various papers taking each of the others into account as well and run a few experiments ourselves as well.

Our main contributions are:

1. Providing with a reusable public reconstruction forensics pipeline through which they pass a set of sources, and a LLM gives you the an analysis on the sources instead,
2. A scoring protocol that scores each of the various incidents through a checklist and a stance audit on the incident,
3. A monitoring boundary experiment that tests whether an agent's unauthorized action can be detected early enough for capability revocation to prevent its effect, while also measuring the cost of false-positive interventions.

2. Related Work

There have already been multitudes of reports by OpenAI, HuggingFace as well as Anthropic covering several parts of these incidents. These are the most credible, detailed and first party sources we can trust. Along with this, there have been independent reconstructions for this incident outside the labs who have scrutinized this incident. The Cloud Security Agent's notes on 700 Rogue Agents and The Benchmark That Broke Containment reframe the OpenAI case as specification gaming and containment failure. HiddenLayer collected from most of the dead-drop repositories what agents wrote as well as real evidence. Factity performed a walkthrough of the same chain while METR & Redwood led on by Hjalmar Wijk and Ajeya Cotra rechecked the whole incident independently to work on their report. A similar incident was noticed a while back in the Nightingale Collective's DseWiki which contained an OpenAI-agent wiki campaign that the labs treated as not part of the same cluster. While all these sources individually document a part of this domain well, none of them produces a shared, queryable claim schema that would let someone compare these claims together, and score them on a checklist.

Adjacent work has been done to minimize or detect such incidents before they reach the final stages. Wang and Huang show that reward hacking is an expected outcome when evaluation is finite [20] and tools keep growing, which explains how the agents behaved over time in HF incident. Bai et al. find that the harness configuration to be the weakest phase of agent harnesses [2], which also needs to be taken into consideration beforehand. Sidhu et al. also mention that there is not definition of an AI incident and no standard report as of now [17], so it becomes slightly tiresome to go through different reports to understand the full problem at scale. None of these clearly provide a step by step checklist on how to react in such situations. In 2009, a surgical safety checklist was released by Haynes et al. which cut mortality rate from 1.5% to 0.8% [10]. Through this report, we target the same and work on a protocol pipeline.

3. Methods

3.1 Forensics

We follow an observe -> structure -> infer discipline while documenting all the various incidents. Public sources are read and quoted before we structure them into a common form using an open-weights LLM (qwen2.5:14b-instruct), which is then later inferred by the protocol pipeline. Each claim carries a claim ID, a source ID, an epistemic tag (confirmed / asserted / unknown), and a corroboration count from a second, independent LLM pass that checks whether other sources restate the same fact without merely citing it . We mark a claim as “confirmed” if the primary source states the fact directly, “asserted” if a primary or credible secondary source claims the fact through their observation but do not disclose the method, and “unknown” if the source itself expresses uncertainty.

We structured five cases through our pipeline, grouping claims by which case's sources they came from: OpenAI→Hugging Face (494 claims from 8 sources), DseWiki (101 claims from 1 source), OpenRouter (39 claims from 2 sources), Anthropic (23 claims from 1 source), and a Hugging Face Space representative-command log (22 claims from 1 source). One can use the pipeline as mentioned in Appendix to further add more sources or cases to this as well.

3.2 Protocol

Based on the schema we get as a result of the Forensics pipeline, we then come up with a set of questions as a pre-run checklist one should follow in this experiment. We specifically focus on the sim v/s real rubric to avoid the agents misaligning themselves in case the simulation and reality is not clear. The inspiration behind choosing each question has been cited next to the question, and also mentioned in detail in the Appendix. We compare these questions and mark them either PASS, FAIL, or UNKNOWN against the schema we create as a result of 3.1, and a missing fact is marked UNKNOWN rather than assumed. For each case, we ask the following:

- Is the internet actually blocked for the run?[18]
- Does the prompt/docs match the real network?[1]
- Can the agent steal cloud/host credentials?[13][18]
- Do we prove isolation before every run?[10]
- Do we notice if the agent cheats the score?[20]
- Do canaries/fake secrets trip an alarm?[8]
- Does the postmortem explain design, or only "what the model wanted"?[3][14][15]
- Is the transcript checked for evaluation-awareness reasoning?[16]

We go into each of these in detail in the Appendix section.

We use two methods to evaluate the different cases on our checklist. A script reads the case's claims and drafts PASS, FAIL, or UNKNOWN. If there's a path out through the simulation, it marks it as FAIL on "was the internet actually blocked" and cites that claim. If there's nothing specific related to the claim mentioned anywhere, it marks it as unknown. And finally, if there's something which is missing like a canary, we mark "would a canary have fired" as fail. Further, we make a local LLM pass over the same claims, and marks the same questions as PASS, FAIL or UNKNOWN from the claim text.

The postmortem question is an exception from the above checklist and we don't mark it as PASS or FAIL using the above pipeline. We analyse the claim extracted from each source, and from that we ask whether that claim can name the design, such as the reward, the network, the tools, the monitor in brief, or only mention what the model seemed to want. We call this the Stance Audit, which scores the sentence based on how it is. We use the local LLM to mark and score this item as well separately.

Finally, we ran this checklist on a setup we controlled in section 3.3, and scored it as well.

3.3 Monitoring

Our checklist identifies whether monitoring and containment controls exist, but not whether detection occurs **early enough to prevent an effect**. We therefore test a controlled authorization environment with a deterministic last-safe intervention boundary. A frozen Qwen3-14B monitor evaluates trajectory prefixes through three views: **M_ACTION** (actions), **M_COT_ACTION** (reasoning + actions), and **M_GROUNDED** (independent authorization evidence), using a fixed threshold of **51**. We measure when relevant evidence becomes available, when it is first detected, and the remaining intervention margin. This distinguishes **eventual detection from preventable detection**: warnings after the last-safe boundary cannot support containment.

4. Results

4.1 Forensics

Across the five cases we pulled 679 claims out of 13 sources (Table 1). How much we could actually reconstruct varies a lot by case: OpenAI→HF alone makes up 494 of those claims across 8 sources, and it's the only case where an independent source backs up a claim at all — 6 claims get corroborated by at least one other source, 2 by two or more. The other four cases only have one or two sources each, so nothing corroborates independently there. That doesn't mean those claims are wrong, it just means we have no second source to check them against.

Case	Claims	Sources	Confirmed	Asserted	Corroborated (≥ 1)	Corroborated (≥ 2)
OpenAI → Hugging Face	494	8	329	165	6	2
Anthropic	23	1	23	0	0	0
Hugging Face Space replay	22	1	22	0	0	0
DseWiki	101	1	100	1	0	0
OpenRouter	39	2	39	0	0	0

4.2 Protocol

The checklist results aren't random, they follow a pattern (Table 2). SR-1 (internet actually blocked), SR-4 (isolation proven before the run), and SR-6 (canaries trip an alarm) FAIL on every case we have enough evidence to score, and SR-6 fails on all five, even OpenRouter where almost everything else is UNKNOWN. SR-7 and SA-1 check something different: does the postmortem actually name the design failure, or does it just say the model decided to do something bad. Both of these PASS on the four cases we could document well. So labs generally do explain what went wrong, once you look past the framing that makes it sound like the model just chose to misbehave. But that's not the same as the isolation and alerting failures being rare. Those show up almost everywhere we have enough evidence to check.

4.3 Monitoring

Our We first asked what information a monitor actually needs to distinguish an unauthorized action from a matched benign one. At the critical prefix (sequence 37), we varied only how authorization information was represented while keeping the monitor and its intervention threshold fixed. The four conditions ranged from observing the staged action without authorization information (G0), through weak structural information (G1), to a relational representation exposing the target route and authorized-route set (G2), and finally an explicit policy verdict (G3).




Figure 1: **Relational evidence best separates prohibited and benign behavior.** At sequence 37, G2 cleanly separates the prohibited and benign cases (100 vs. 0) at the fixed threshold of 51, while other evidence representations produce benign threshold crossings.

Action evidence alone (G0) does not distinguish the cases, while structural (G1) and explicit-verdict (G3) evidence still produce benign crossings. Only relational evidence (G2), exposing the target and authorized-route set, achieves clean separation. This suggests that **evidence representation matters, not just evidence availability**, motivating our subsequent representation-robustness tests.

This result suggests that **how policy-relevant information is represented may matter as much as whether that information is available at all**. However, the clean G2 separation alone does not establish that the monitor has learned a representation-invariant authorization rule. It could still be responding to superficial features such as route names, identifier morphology, list structure, or other properties of the particular representation. We therefore next tested whether the same authorization semantics produced the same monitor decision when those superficial representations were changed

Prohibited cases consistently score 100, but benign scores vary from 0–100 across equivalent representations. Thus, **relational evidence supports detection, but the monitor is not representation-invariant**, creating potential false-positive control costs.




Figure 2: **Attack detection is stable, but benign decisions are representation-sensitive.** All 12 prohibited variants are detected, while 4/12 benign variants cross threshold and only 3/12 satisfy control-ready invariance.

Across the 12 tested assignments, the prohibited case is remarkably stable: every representation receives a score of 100 and crosses the intervention threshold. The benign side is substantially less stable. All three alphabetic representations remain at 0, while changing the identifier morphology produces benign scores ranging from 0 to 100. Four of the twelve benign assignments cross threshold, yielding only 8/12 semantic decision invariance and 3/12 under the stricter control-ready criterion.

This asymmetry matters for control. The monitor consistently recognizes the prohibited relation in this test, but its treatment of the corresponding benign relation depends on how otherwise equivalent identifiers are represented. We therefore cannot interpret the G2 result as evidence of general representation-invariant authorization reasoning. Instead, the experiments support a narrower conclusion: **relational authorization evidence can be sufficient for timely detection, but the resulting decision is representation-sensitive.**




Figure 3: Timely authority revocation prevents the prohibited effect. In matched executions, the monitor crosses threshold at seq. 37 and the same prohibited action is attempted at seq. 38. Without intervention, it is accepted and the effect occurs; with authority revoked before the last-safe boundary, it is rejected and the effect is prevented.

Figure 3 shows that **detection is useful only while intervention remains possible**. With identical pre-intervention history, monitor input, score, and subsequent action attempt, revoking authority at seq. 37 prevents an effect that occurs without intervention. However, representation-sensitive false positives also have real control costs: in a benign Phase F case, the same policy unnecessarily revoked authority and caused task failure. Thus, effective containment requires both **timely detection and a sufficiently reliable signal to act on**.

Overall, our monitoring experiment therefore shifts the evaluation target from “*does the monitor eventually detect the problem?*” to “*does policy-relevant evidence become visible and trigger a sufficiently reliable warning before the last enforceable intervention point?*” The matched replay shows that this distinction has causal consequences: timely revocation can prevent the prohibited effect, while the benign false-positive case shows that monitor errors can propagate into real control costs.

5. Discussion and Limitations

- FAILs cluster on isolation and alerting (SR-1, SR-4, SR-6), not on honesty
- SR-7/SA-1 (does the postmortem name the design failure) PASS on all 4 documented cases
- So labs do explain what went wrong, once you look past “the model decided to...” framing
- But isolation/alerting failures still show up almost everywhere we can check

- How much we reconstructed just tracks source count, not truth
- Only OpenAI→HF (most sources) clears our bar for independent corroboration
- Everything resting on the other four cases is really single-sourced
- **Detection must happen before the last-safe intervention point**; eventual detection alone is not enough for containment.
- **False positives can become real control failures** when monitoring is tied to automatic intervention.

Limitations

Every case mentioned here is cited from a public postmortem, blog post or a third-party writeup, not lab internal telemetry. As what organizations disclose and frame it is shaped by legal and reputational incentives, we can only describe only the claims that surfaced, not the true underlying cause. Additionally, we only document five cases in our report, which is enough to build a scoring method, but not enough to actually say how AI incidents happen.

We also scored each case already knowing how it ended. This makes any claim that a lab "should have" caught something earlier vulnerable to hindsight bias, the tendency to see past events as more predictable than they actually were [6]

Future Work

The sharpest gap is this. Our checklist can tell you if a case failed, but it cannot tell you what to do next, and it is not reliable enough yet to make that call safely. An independent LLM audit agrees with our own heuristic scorer only 55% of the time, dropping to 25% on the one case big enough to actually stress-test it. Sidhu et al. already flag that there is no standard definition or report format for an AI incident [17], which is exactly the gap our schema tries to close, but closing it alone is not enough. The real next step is connecting a more reliable version of this checklist to actual escalation-criteria work like Gomez et al.'s [9], so a FAIL verdict leads somewhere instead of sitting in a table, and contributing our cases and schema to standardization efforts already underway such as the AI Incident Database [11], the OECD's common reporting framework [12], and CSET's mandatory-reporting-regime proposal [4].

6. Conclusion

We present a structured approach for moving from **incident reconstruction to actionable containment**. Across our analyses and experiments, a consistent lesson emerges: identifying what went wrong is not enough. Effective incident response requires reliable evidence, timely detection, and enforceable interventions before the opportunity for prevention is lost. Our results ultimately reinforce a simple principle: **detection is not containment**.

Code and Data: Github: <https://github.com/GaganaMD/detection-is-not-containment>

Author Contributions

GMD worked on the monitoring pipeline as well as the major experiments for the report. MI came up with the idea to build the forensics pipeline. KG worked on evaluating forensics and protocol pipeline and performed experiments for the same.

LLM Usage Statement

We used Codex 5.6 and Claude Sonnet 5 to help with the project design process as well as to write up the code for this report. Appendix A.1-A.6 have been completely written by Claude Sonnet 5.

References

1. AICPA. (2017). Trust services criteria for security, availability, processing integrity, confidentiality, and privacy (SOC 2). American Institute of CPAs.
2. Bai, Y., Duan, J., Peng, J., Wu, X., Liu, S., Wang, S., & Chen, T. (2026). HarnessRisk: A lifecycle-oriented benchmark for agent harness safety (arXiv:2608.17597). <https://arxiv.org/abs/2608.17597>
3. Dennett, D. C. (1987). The intentional stance. MIT Press.
4. Dixon, R. B. L., & Frase, H. (2025). AI incidents: Key components for a mandatory reporting regime. Center for Security and Emerging Technology. <https://cset.georgetown.edu/publication/ai-incidents-key-components-for-a-mandatory-reporting-regime/>
5. Einhorn, E., & Ziv, A. (2011). Biased voluntary disclosure. Review of Accounting Studies, 16(3), 420–442. <https://doi.org/10.1007/s11142-011-9153-8>
6. Fischhoff, B., & Beyth, R. (1975). I knew it would happen: Remembered probabilities of once-future things. Organizational Behavior and Human Performance, 13(1), 1–16. [https://doi.org/10.1016/0030-5073\(75\)90002-1](https://doi.org/10.1016/0030-5073(75)90002-1)
7. Flyvbjerg, B. (2006). Five misunderstandings about case-study research. Qualitative Inquiry, 12(2), 219–245. <https://doi.org/10.1177/1077800405284363>
8. Fraunholz, D., Anton, S. D., Lipps, C., Reti, D., Krohmer, D., Pohl, F., Tammen, M., & Schotten, H. D. (2018). Demystifying deception technology: A survey (arXiv:1804.06196). <https://arxiv.org/abs/1804.06196>
9. Gomez, F., Ball, M., Harre, M., Preston, L., Schwab, J., & Machado, C. (2026). Designing escalation criteria for international AI incident response: Criteria, triggers, and thresholds (arXiv:2604.23183). <https://arxiv.org/abs/2604.23183>
10. Haynes, A. B., Weiser, T. G., Berry, W. R., et al. (2009). A surgical safety checklist to reduce morbidity and mortality in a global population. New England Journal of Medicine, 360(5), 491–499. <https://doi.org/10.1056/NEJMsa0810119>
11. McGregor, S. (2021). Preventing repeated real world AI failures by cataloging incidents: The AI Incident Database. Proceedings of the AAAI Conference on Artificial Intelligence, 35(17), 15458–15463. (also arXiv:2011.08512)

12. OECD. (2025). Towards a common reporting framework for AI incidents (OECD Artificial Intelligence Papers, No. 34). OECD Publishing. <https://doi.org/10.1787/f326d4ac-en>
13. OWASP Foundation. (2025). OWASP top 10 for LLM applications 2025. <https://genai.owasp.org/resource/owasp-top-10-for-llm-applications-2025/>
14. Reason, J. (1990). Human error. Cambridge University Press.
15. Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10, pp. 173–220). Academic Press.
16. Schoen, B., et al. (2025). Stress testing deliberative alignment for anti-scheming training (arXiv:2509.15541). <https://arxiv.org/abs/2509.15541>
17. Sidhu, H. K., Scholefield, R., Annan, N., Hernandez, K., Hou, I. N., et al. (2026). Open problems in AI incident governance (arXiv:2607.05163). <https://arxiv.org/abs/2607.05163>
18. Souppaya, M., Morello, J., & Scarfone, K. (2017). Application container security guide (NIST Special Publication 800-190). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-190>
19. Vo, H., & Pham, M. D. (2025). Beware of false prophets: Cybersecurity risk and strategic voluntary disclosure. *The British Accounting Review*, Article 101578. <https://doi.org/10.1016/j.bar.2025.101578>
20. Wang, J., & Huang, J. (2026). Reward hacking as equilibrium under finite evaluation (arXiv:2603.28063). <https://arxiv.org/abs/2603.28063>

Limitations and Dual-Use Considerations

`gaganaaaa.dev@gmail.com` can put stuff like

- Intended use. Score frame and isolation failures. Push postmortems to name design facts, not just intent. Ask whether monitoring could have stopped harm in time.
- We do not provide. Exploit chains. Weaponizable attack steps. Credentials. Instructions to reproduce the July intrusions against third parties.
- Mitigations. Epistemic tags on every claim. Both extraction prompts refuse to pull out exploit code or attack payloads (Appendix A.2). No payload or proof-of-concept attack code anywhere in this pipeline.
- Residual risk. Structured timelines can still tell an attacker what was possible in a public incident. We accept that risk because the facts are already public. Our addition is scoring discipline, not new offensive capability.

Appendix

A.1 Forensics and protocol pipeline algorithms

Neither pipeline is described step-by-step in the main text; both are short enough to spell out here.

Forensics (claims extraction + cross-source corroboration):

1. Source registration — each raw source is registered in `claims/data/source_registry.json` with a `source_id` and the case(s) it belongs to.
2. Claim extraction (`extract_claims.py`) — each source is chunked and passed to a local LLM (qwen2.5:14b-instruct-q4_K_M via Ollama). For each chunk the model emits one atomic claim per line: {claim, quote, epistemic_hint}. Every claim must carry an exact, unbroken verbatim quote; the model is told to skip a claim rather than approximate one by ellipsis-joining. Claims are tagged confirmed/asserted/unknown from the source's own language.
3. Deterministic prefilter (`prefilter.py`) — before any second LLM call, `is_derivative_by_link` checks whether a candidate document's raw HTML literally hyperlinks back to the origin source, producing a boolean used later to downgrade citation-only repetition.
4. Cross-source corroboration (`corroborate.py`) — for each (origin claim list, candidate document) pair, the same local LLM classifies every claim against the candidate as DIRECT / CONCEPTUAL / DERIVATIVE / CONTRADICTS / NOT_MENTIONED.
5. Aggregation (`aggregate_claims.py`) — rolls every (claim, candidate) verdict up into `corroboration_count` and `independent_corroboration_count` per claim. The latter counts only DIRECT/CONCEPTUAL verdicts where step 3's link-check did NOT find a citation back to the origin — this is what makes “independent” mean more than “the LLM said DIRECT.”
6. Output — `claims_extracted.jsonl` (raw) and `claims_scored.jsonl` / `claims_scored_summary.csv` (with corroboration counts), consumed by the protocol scripts below and by `claims_case_metrics.py` for Table 2.

Protocol (checklist scoring + independent audit):

1. Case mapping — `source_registry.json` maps each source to one or more of the 5 canonical cases; a claim belongs to a case if its source does.
2. Heuristic scoring (`score_checklist_from_claims.py`) — a deterministic keyword/regex pass over each claim's text+quote checks SR-1...SR-6, producing a draft PASS/FAIL/UNKNOWN with the specific claim_ids that triggered it. SR-4 and SR-6 default to FAIL (not UNKNOWN) when the relevant claim never appears at all, per the rubric's own absence-is-a-FAIL rule.
3. Stance Audit (`stance_tag.py` + `stance_*.md`) — a human-authored sheet per case scores SA-1: does the postmortem name design facts (reward, tools, network, monitors), or only “what the model wanted”? Locked once source-checked against the primary postmortem text.
4. SR-7 pass-through — the checklist script pulls the Stance Audit's own verdict into SR-7 rather than re-deriving it, so SR-7 and SA-1 agree by construction on the heuristic side.
5. Independent LLM audit (`llm_audit_from_claims.py`) — the same local model re-answers SR-1...SR-7 + SA-1 directly from the case's claim list, bypassing the heuristic script entirely, and the two are diffed case by case (Table 4).

6. SR-8 (evaluation-awareness CoT auditing) is never scored by either step; it stays permanently UNKNOWN pending a human transcript read, since no public source here contains a raw chain-of-thought transcript to audit.

A.2 LLM call-sites and model selection

All three LLM call-sites run on local Ollama only (qwen2.5:14b-instruct-q4_K_M); no API key or external service is required to reproduce any number in this report.

Script	Input	Output
extract_claims.py	One source document	One atomic claim per line: {claim, quote, epistemic_hint}. Refuses to extract exploit code or attack payloads — factual/structural claims only.
corroborate.py	Claims from one source + full text of a candidate source	DIRECT/CONCEPTUAL/CONTRADICTS/DERIVATIVE/NOT_MENTIONED per claim. DERIVATIVE is explicit: a citation back to the origin is not independent corroboration.
llm_audit_from_claims.py	A case's full claim list	Independent PASS/FAIL/UNKNOWN for SR-1...SR-7 + SA-1, with claim-id citations and one-sentence reasoning, diffed against the heuristic sheet.

Table A1. The three LLM call-sites. All prompts are short (2–3KB), plain-text, and checked into `claims/src/`.

Model choice is not an assumption: a real 3-way comparison on one 22K-character chunk of the Anthropic source, same prompt, same quote-verification check, is what fixed qwen2.5:14b-instruct-q4_K_M as the model for every LLM call in this report.

Model	Claims extracted	Quote-verified	Time/chunk
gpt-oss:20b	1	100%	4.9s
qwen2.5:1.5b-instruct	0 (unparseable)	—	27s
qwen2.5:14b-instruct-q4_K_M	23	96%	138s

Table A2. *gpt-oss:20b's own hidden reasoning trace admitted why it under-enumerated: “Need many claims. Due to time, produce few? ... Probably produce some sample.” It is tuned to converge on one good answer, not to exhaustively enumerate — the wrong shape for this task. qwen2.5:1.5b ran out of budget and produced unparseable output. qwen2.5:14b is slower but complete, which is what an exhaustive claim-coverage pipeline needs.*

Reproducibility note: Ollama's format JSON-schema-constrained decoding silently returns an empty response when combined with a reasoning-style model's hidden thinking channel (confirmed directly with two bare curl calls, same prompt, format present vs. absent). All three call-sites above avoid format entirely — they ask for the JSON shape in the prompt text and parse strictly instead.

A.3 Sim/real rubric grounding

Item	Supporting source(s)	Confidence
SR-1	gVisor/Firecracker default-deny-egress practice; NIST SP 800-190	Grounded
SR-2	Evaluation-awareness literature (model side); SOC 2 Type II design-vs-operating-effectiveness (documentation side)	Partially grounded
SR-3	NIST SP 800-190; OWASP LLM Top 10 2025 (LLM06, LLM02); MITRE ATLAS Exfiltration	Grounded
SR-4	Haynes et al. 2009 NEJM (WHO Surgical Safety Checklist, mortality 1.5%→0.8%); Gawande, Checklist Manifesto; HarnessRisk 2608.17597	Grounded — strongest cross-domain citation in the rubric
SR-5	Wang & Huang 2603.28063; EvilGenie 2511.21654; Hack-Verifiable 2605.20744; BAITBENCH 2608.30724	Grounded
SR-6	Canary tokens / Thinkst Canary as an established detective-control category; MITRE ATLAS	Grounded
SR-7 / SA-1	Dennett, The Intentional Stance (1987); Reason's active-failure/latent-condition split; Heider/Kelley attribution theory (fundamental attribution error)	Grounded — best-grounded item overall
SR-8	Schoen et al. 2509.15541 (Anti-Scheming); corroborated by 2309.00667, 2605.23055, 2606.29196	Grounded

Table A3. Condensed from resources/LITERATURE_REVIEW.md §12. Seven other AI-generated candidate items (SR-9...SR-16) were reviewed and not adopted — each was either a more granular restatement of SR-1...SR-7 or too thinly grounded (single blog-post citation); full reasoning per candidate is in protocol/sim_real_rubric.md.

A.4 Full checklist cites by case

The heuristic verdicts in Table 3 are backed by specific claim citations; the full per-case tables (protocol/checklist_scores_from_claims/*.md) are reproduced below, condensed to the ID/result/notes columns — the full claim-cite lists are long and are left in the repo rather than reproduced here in full.

OpenAI→HF

ID	Result	Notes
SR-1	FAIL	6 cites (+24 more)
SR-2	FAIL	6 cites (+18 more)
SR-3	FAIL	6 cites (+28 more)
SR-4	FAIL	no automatic cite; isolation failures present elsewhere
SR-5	FAIL	2 cites

SR-6	FAIL	no canary/alert evidence anywhere in this case's claims
SR-7	PASS	see stance_openai_hf.md, locked 2026-09-13
SR-8	UNKNOWN	always UNKNOWN; needs a human transcript read

Anthropic

ID	Result	Notes
SR-1	FAIL	5 cites
SR-2	FAIL	2 cites
SR-3	UNKNOWN	no automatic cite
SR-4	FAIL	no public preflight proof
SR-5	UNKNOWN	no automatic cite
SR-6	FAIL	no canary/alert evidence
SR-7	PASS	see stance_anthropic.md, locked 2026-09-13
SR-8	UNKNOWN	always UNKNOWN

HF Space

ID	Result	Notes
SR-1	FAIL	4 cites
SR-2	FAIL	1 cite
SR-3	FAIL	3 cites
SR-4	FAIL	no public preflight proof
SR-5	UNKNOWN	no automatic cite
SR-6	FAIL	no canary/alert evidence
SR-7	PASS	see stance_hf_space_representative.md, locked 2026-09-13
SR-8	UNKNOWN	always UNKNOWN

OpenRouter

ID	Result	Notes
SR-1	UNKNOWN	no automatic cite
SR-2	UNKNOWN	no automatic cite
SR-3	UNKNOWN	no automatic cite
SR-4	UNKNOWN	no automatic cite
SR-5	UNKNOWN	no automatic cite
SR-6	FAIL	no canary/alert evidence
SR-7	UNKNOWN	see stance_openrouter.md — not yet written
SR-8	UNKNOWN	always UNKNOWN

Dse Wiki


ID	Result	Notes
SR-1	FAIL	6 cites
SR-2	FAIL	6 cites (+12 more)
SR-3	UNKNOWN	no automatic cite
SR-4	FAIL	no public preflight proof
SR-5	FAIL	1 cite
SR-6	FAIL	canary/decoy mentioned but no alert language found — treated as untripped
SR-7	PASS	see stance_dsewiki.md, locked 2026-09-13
SR-8	UNKNOWN	always UNKNOWN

A.5 Independent claim corroboration

Across the full 20-source corpus (not limited to the 5 canonical cases), the pipeline's own “independently corroborated ≥ 2 ” bar is cleared by the OpenAI→HF agent-count cluster (“~700 rogue agents / ~1,200 total / ~70,000 messages”, independently restated by METR, Facity, and the CSA “700 Rogue Agents” note) and by ExploitGym's “898 instances” figure. A deterministic link-check (is_derivative_by_link) that flags citation-only repetition disagrees with the LLM's own DIRECT verdict often enough to matter: after filtering, 22 corroboration pairs — not the raw, unfiltered count — survive both checks. 333 of 390 checked pairs are NOT_MENTIONED elsewhere, and zero CONTRADICTS verdicts means no conflicting figure turned up in what was checked, not that every fact is proven consistent everywhere.

A.6 Detection timing reflects when policy-relevant information is visible

Detection timing reflects when policy-relevant information is visible



A1: CoT + action and grounded first cross at seq 32 (margin 6); action-only has no crossing. A2: only grounded first crosses at seq 37 (margin 1).

Figure: First threshold crossing depends on when policy-relevant information becomes visible. We compare action-only, CoT-plus-action, and grounded views for two deterministic trajectories; NC denotes no calibrated crossing. In A1, explicit policy-relevant reasoning makes CoT-plus-action and grounded views cross at sequence 32, six steps before the last-safe boundary. In neutral-reasoning A2, only the grounded view crosses, at sequence 37, leaving one step of intervention margin. These two trajectories illustrate information availability rather than establishing a general ranking among observation channels.

A.7 More on Monitor Suspicion scores




Figure A2: Grounded evidence enables detection just before the last-safe intervention boundary. For the A2 neutral-reasoning trajectory, action-only and CoT-plus-action observations remain below the fixed threshold of 51. Grounded policy-relevant information becomes available at sequence 37 and immediately raises the monitor score above threshold, leaving a one-sequence intervention margin before the last-safe boundary at sequence 38. Detection after this boundary would no longer be useful for prevention.

This result shows why **when information becomes available matters as much as whether a monitor eventually detects the issue**. In the neutral-reasoning condition, neither actions nor CoT provide sufficient evidence for detection. Once independently grounded authorization evidence becomes visible at sequence 37, however, the monitor crosses threshold immediately. With the last-safe boundary at sequence 38, this leaves only a narrow window in which an intervention can still prevent the prohibited effect.

The figure therefore motivates our distinction between **eventual detection and preventable detection**: a monitor may correctly identify risky behavior later in the trajectory, but that warning has limited control value once the last enforceable intervention point has passed.

A8. Documented Table on the Protocol followed

ID	Question	OpenAI→HF	Anthropic	HF Space	OpenRouter	DseWiki
----	----------	-----------	-----------	----------	------------	---------

SR-1	Internet actually blocked?	FAIL	FAIL	FAIL	UNKNOWN	FAIL
SR-2	Prompt matches real network?	FAIL	FAIL	FAIL	UNKNOWN	FAIL
SR-3	Cloud/host creds stealable?	FAIL	UNKNOWN	FAIL	UNKNOWN	UNKNOWN
SR-4	Isolation proven pre-run?	FAIL	FAIL	FAIL	UNKNOWN	FAIL
SR-5	Score-cheating noticed?	FAIL	UNKNOWN	UNKNOWN	UNKNOWN	FAIL
SR-6	Canaries trip an alarm?	FAIL	FAIL	FAIL	FAIL	FAIL
SR-7	Postmortem names design?	PASS	PASS	PASS	UNKNOWN	PASS
SR-8	Eval-awareness CoT checked?	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN	UNKNOWN
SA-1	Stance Audit (design vs. intent)	PASS	PASS	PASS	UNKNOWN	PASS