

---

# Evaluating Containment of AI Agents: A Nine-Rule Standard for Verifiable Sandbox Security<sup>1</sup>

---

Vikas Reddy  
Johns Hopkins  
University

Junyi Liu  
Columbia University

Somay Kousis  
ABV-IIITM Gwalior

With  
Apart Research

## Abstract

AI agent evaluations can become production incidents when an attacker escapes an initial sandbox and moves through package, network, cloud, cluster, and identity systems. This paper uses the publicly reported July 2026 Hugging Face intrusion to test whether those boundaries would contain an attacker after an initial compromise. We translate the documented attack path into nine executable rules and an offline validation harness that audits deployment artifacts without access to the original model, lab, or specialized security tooling. The hardened reference environment passed all nine rules and blocked thirteen documented attack actions, while the reconstructed insecure environment and two adversarial variants failed every rule. Two red team rounds identified and resolved 17 bypasses, strengthening the checks against configurations that appeared secure while preserving an attack path. These results show that public incident evidence can support practical, repeatable tests of AI sandbox containment.

## 1. Introduction

AI evaluations increasingly place untrusted agents near production infrastructure, credentials, source control, and cloud services, where a single failed boundary can allow an isolated experiment to become a wider production compromise. The July 2026 Hugging Face incident illustrates this risk, as an autonomous agent escaped its sandbox during an internal capability evaluation, moved through the production systems of two organizations, and reached Hugging Face infrastructure before either organization detected the activity [1,3].

---

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

The intrusion began when the attacker exploited a zero day vulnerability in a package registry cache proxy reachable through an approved network path. The compromised proxy provided access to an external code evaluation harness, which the attacker used to stage the campaign [2,4]. The attacker then moved through Hugging Face data processing, cloud, cluster, and identity systems before obtaining credentials sufficient to exfiltrate private data [2,5]. The public reconstruction records document roughly 17,600 actions over five days and identifies the relevant vulnerabilities, infrastructure, and attack stages [1,2,5], providing a concrete basis for evaluating whether a deployment could stop the attacker at each boundary.

Existing work gives operators useful security frameworks, but not a direct way to determine whether a deployment would have stopped a real production attack. Surveys classify risks across agent architectures and lifecycle stages [6,7], while other work evaluates sandbox mechanisms or model capabilities in constructed settings [8,9]. We instead use the incident as the specification and ask whether each documented boundary would have contained the attacker. Once a worker, proxy, or credential is compromised, it cannot be trusted to continue enforcing policy, so containment must come from the next independent boundary.

This paper makes three contributions. First, it converts the public incident disclosures into a structured record of attack phases, action counts, and representative commands [1,2]. Second, it derives nine executable containment rules, each tied to a boundary crossed during the incident. Third, it provides an offline validation harness that compares a hardened reference environment with the documented insecure configuration and adversarial variants. The harness runs locally and requires no access to the original environment, model, or specialized security tooling.

## 2. Related Work

Research on cyber capable AI agents has identified many of the risks exposed by this incident. Siddik’s review uses the same case to examine supply chain exposure, persistent control channels, and rapid automated action [10]. Our work addresses similar concerns, but with a different purpose. Rather than using the incident to support a broad synthesis of agent security risks, we use it to derive and audit the specific boundaries crossed during the attack.

Other work evaluates sandbox architectures and execution environments. Andronchik and Lokhnikov compare sandbox products across attack surface, patching, and fuzzing characteristics [12]. SandboxEval provides a general property suite for environments that execute untrusted language model generated code [13]. These tools assess overall sandbox quality, but a system can satisfy general properties while remaining exposed to the particular sequence used in the July 2026 incident. Our contribution is not a new enforcement engine. Instead, it provides an incident grounded set of checks that operators can implement using existing policy and isolation tools such as OPA Gatekeeper, Kyverno, gVisor, and Kata Containers [17,18].

Ruan et al. identify risks from tool using agents through an emulated sandbox [8], while Marchand et al. measure model capability to escape container sandboxes through seeded vulnerabilities [14].

These studies examine whether models or sandbox mechanisms can cause harm under controlled conditions. Our work instead evaluates whether a deployment configuration would have blocked the boundary crossings documented in a public incident. Supply chain systems such as Sigstore, SLSA, and in toto complement this approach by strengthening package verification with independent provenance evidence [19,20,21].

### 3. Methods

We use the public incident record as the specification and work backward from the documented attack. The record covers five days and nine phases, with activity rising from 3,779 actions on day 1 to 7,677 during the main campaign on day 3 before falling to 1,130 during cleanup on day 5. Action volume alone does not determine security importance. The dropper phase produced 6,972 actions, while lower volume stages such as Kubernetes enumeration, credential theft, and exfiltration created greater risk. The standard therefore prioritizes the consequence of each crossed boundary. We assume an evaluated agent can execute arbitrary code, inspect its process and file system, submit malicious data and configuration, encode payloads, and reuse worker credentials or network keys. It may combine these capabilities across package, sandbox, data loading, cloud, cluster, remote access, and identity boundaries. Because a compromised component cannot be trusted to enforce policy, each rule requires the next independent boundary to contain the compromise.

The defender controls the predeployment artifacts used by the audit, including package mirrors, manifests, container settings, data loader policy, network and metadata policy, RBAC and IAM graphs, worker environments, and alert records. We test whether a fixed reference environment satisfies the nine rules and blocks the identified attack steps. We do not claim that production systems always match declared configurations, that the public record captures every action, or that the rules cover all escape paths. Because no complete configuration snapshot or command log is public, the insecure environment is a conservative reconstruction. Each setting is tied to a disclosed source, published command, or documented boundary crossing, while unsupported settings are excluded or recorded as assumptions.


Each rule has an executable offline checker covering package integrity, outbound access, worker privilege, unsafe HDF5 loading, configuration templating, cloud metadata access, cluster reachability, remote access credentials, and credential reuse. Rules 4, 5, and 9 also test exploit primitives using a malicious HDF5 reference, a template injection payload, and a replayed signed capability token. Early egress, metadata, and control plane checks passed the hardened environment and failed the insecure reconstruction, but could still accept adversarial configurations. We therefore red teamed these rules. Rule 2 rejects domain fronting, IP literal destinations, and wildcard allowlists. Rule 6 requires link local metadata denial or token based metadata access. Rule 7 evaluates routing and authorization independently. Rule 8 correlates secret exposure, remote access preparation, and outbound or lateral movement in order for the

same workload identity, using bounded decoding for common encodings. Rule 9 models identities, resources, and boundary categories as an access graph, failing when an identity exceeds its declared reach limit. This approach measures the boundary categories that each identity can reach. It does not compute a minimum cut. Appendices A through C provide the full specifications, adversarial cases, and graph procedure.

## 4. Results

**Headline back test.** We computed a threshold directly from the public daily action counts rather than from anything produced by our own system. A per identity threshold of one thousand actions per day, roughly forty one point six actions per hour, would have fired forty one point six hours before the incident’s largest single day spike, the 7,677 actions recorded on day 3. The result is stable across a reasonable range of thresholds, firing 44.8 hours early at 500 actions per day and 35.3 hours early at 2,000 actions per day.

This back test rests on two assumptions. The actions recorded on day 1 were spread evenly across that day rather than clustered, and every day 1 action could be attributed to a single identity because the public record gives no identity level breakdown. If either assumption fails, the true advantage could shift earlier or later. This result is exploratory and separate from the nine rule containment score.



*Figure 1. Attacker actions recorded per day across the five day incident, with the main campaign spike on the third day and the point at which our proposed quarantine threshold would have fired.*


**Replay of the known attack path.** The published incident record contains representative commands spanning the documented intrusion sequence. The boundary testing steps represented in those commands are blocked by the fixed reference lab, with each of the nine rules mapping

natively to at least one of them. Running the same replay against the historical insecure configuration, a first round attacker configuration, and a second round adversarial configuration each produces a score of zero out of nine, while the fixed lab consistently scores nine out of nine.

This gives a direct and reproducible answer to whether a given configuration would have stopped the documented attack path rather than a general assessment of sandbox quality. It does not establish that the reconstructed insecure environment captures every detail of the original deployment. It shows that the nine rules distinguish the hardened configuration from the insecure and adversarial configurations across the public boundary crossings represented in the evaluation.

**Containment holds across the chain.** Rule 9 limits the value of anything an attacker still manages to steal after Rules 1 through 8 have acted. Starting from a stolen credential with no upstream containment, the credential's combined reach across the metadata, cluster, and remote access boundaries collapses step by step as each upstream rule is added, reaching zero once all three hold together. This is the condition our blast radius check enforces directly. No identity may retain combined reach across more boundary categories than the declared limit.


Containment in our standard functions as a chain rather than a single gate. Each additional rule reduces the value of a compromise elsewhere in the chain. The result does not show that every possible identity path has been removed. It shows that the declared access graph no longer permits the tested credential to combine the relevant metadata, cluster, and remote access boundaries.



*Figure 2. A stolen credential's combined reach across the metadata, cluster, and VPN boundaries, shown collapsing to zero as each upstream containment rule is added*

**Validation time and runtime alerting.** Our nine rules mix two operational modes, and we separate them explicitly. Rules 1 through 7 and Rule 9 are predeployment validation checks. They audit a declared configuration before a workload runs and produce a pass or fail result with no runtime component. Rule 5 is enforced entirely before deployment and has no applicable runtime timing measurement. Rule 8 is the one rule with a runtime alerting component because it correlates live sequence and encoding evidence rather than a static declaration.

Figure 3 therefore reports local validation time for Rules 1 through 7 and Rule 9, and local alert latency for Rule 8. For static rules, the timing reflects how long the checker takes to evaluate a committed configuration snapshot against its pass condition. It does not measure how long a live system would take to notice an ongoing attack. Rule 8 timing measures alert generation within the local harness after the final required sequence event.



*Figure 3. Measured detection latency for each rule shown as a percentage of its allotted detection service level budget, with every rule alerting well within its target*

**Checker quality and red team results.** A separate quality matrix evaluated two attack variants and two benign variants per rule against an independent observer process. It produced 18 true positives, zero false negatives, zero false positives, and 18 true negatives across 36 cases, with every attack observation meeting its fixture evaluation target. This establishes a small deterministic baseline for checker correctness rather than a statistically representative false positive or false negative rate at production scale.

Our red team process identified and resolved 17 bypasses that had produced false negatives. Ten were found during the initial review, and seven emerged during deeper adversarial testing of Rules 2, 6, 7, and 8. Although Rules 2, 6, and 7 passed the hardened lab and failed the broken lab, this comparison did not show whether they could resist configurations designed to preserve an attack path. The final suite fails closed, so hostile or malformed input fails the relevant rule without crashing or disabling the remaining scorecard checks.

## 5. Discussion and Limitations

The clearest pattern across our results is that a containment standard becomes more trustworthy once every rule is backed by something a third party can execute rather than only read. Our first pass at Rules 2, 6, and 7 looked complete until we built adversarial fixtures specifically to defeat them. Those fixtures showed that a checker distinguishing a present field from an absent one is not the same as a checker that resists an adversary who knows the shape of the check. A checklist that has not been tested against its own blind spots should therefore be treated as a draft rather than as a standard.

This work has several limits. First, it is a static configuration audit rather than a live cluster evaluation. The checkers inspect declared YAML and JSON, pod specifications, RBAC graphs, egress policy, and credential manifests. A deployment can pass the audit and later drift because of policy changes, workload mutation, credential rotation, or network changes. Admission controllers and runtime monitors can help address this gap, but neither is evaluated here.

Second, Rule 1 hashes supplied artifacts against a supplied manifest. It does not independently establish publisher authenticity. An attacker who controls both artifacts and manifest hashes could create a self consistent poisoned mirror. Signed provenance, transparency logs, and independently verifiable build metadata could close this gap through systems such as Sigstore, SLSA, and in toto [19,20,21].

Third, Rule 8 supports standard and URL safe base64, base32, hexadecimal, percent encoding, and gzip under bounded recursion and output limits. It does not claim coverage of arbitrary encryption, proprietary encodings, or events absent from collected telemetry. Its sequence correlation tolerates unrelated events between required signals, but it does not treat reversed event order as equivalent evidence of a remote access pivot.

Finally, the evidence base consists of nine phase totals, five day counts, and representative public commands from one incident. The historical insecure configuration is a conservative reconstruction rather than an exact replica. Rule 7 also uses node role inference based partly on names such as worker and control, which can over flag ambiguous environments. A production implementation should instead rely on explicit Kubernetes labels, infrastructure inventory data, or equivalent role metadata. The standard should be evaluated against distinct incident classes before making broad generalization claims.

## Future Work

Three extensions follow from these limitations. First, runtime probes could check the nine invariants against live traffic, process activity, credential use, and control plane access rather than declared configuration alone. Admission controllers such as OPA Gatekeeper and Kyverno could enforce rules at deployment, while runtime monitoring could detect drift [17,18].

Second, package verification could be strengthened through independent signatures, provenance attestations, and transparency logs, reducing the risk of a self consistent poisoned mirror. Rule 8 should also be tested on benign workload telemetry to measure how event noise, reordering, encoding variation, and different correlation windows affect its performance.

Another direction is testing the standard beyond this infrastructure focused incident. Future work can apply the same incident based method to cloud pivots, local privilege escalation, browser mediated agents, and attacks based on behavioral deception rather than network movement. This would show which containment rules transfer directly, which require adaptation, and which additional boundaries should be included in a broader standard.

## 6. Conclusion

This work turns a partially documented sandbox escape into a nine rule containment standard, with each requirement tied to a specific attack boundary and an executable offline checker. Across a fixed reference lab, insecure and adversarial configurations, documented attack replay, local validation measurements, and two rounds of red teaming, the standard separates hardened from insecure deployments with no observed errors in the deterministic fixture set. It remains deliberately bounded by a limited public evidence base, static deployment artifacts, and a conservative reconstruction of the historical environment. The standard should be understood as a reproducible floor for evaluation sandbox containment rather than a final answer. These results show that public incident evidence can support practical, repeatable, and independently verifiable containment audits for AI agent evaluations.

## Code and Data

The full checker implementation, validation harness, test labs, figure generation scripts, and the curated public incident record, along with the sensitivity analysis notebook used for the headline back test, are available at <https://github.com/Somay-kousis/Apocalypse>.

## Author Contributions

Vikas led rules two, six, and seven, including the domain fronting, metadata hop limit, and routing versus authorization checks, wrote the report manuscript, and wrote the full pytest test suite covering every section. Somay led rules four, five, and nine, the runner architecture, the control specification, and the curated incident record, and led the red team process. Junyi led rules one, three, and eight along with the threat model.

## References

- [1] Hugging Face. "Security incident disclosure — July 2026." July 16, 2026.  
<https://huggingface.co/blog/security-incident-july-2026>
- [2] Hugging Face. "Anatomy of a Frontier Lab Agent Intrusion: A Technical Timeline of the July 2026 Incident." July 27, 2026. <https://huggingface.co/blog/agent-intrusion-technical-timeline>
- [3] OpenAI. "The Hugging Face incident and the road ahead." August 26, 2026.  
<https://openai.com/index/hugging-face-incident-and-the-road-ahead>
- [4] OpenAI. "OpenAI and Hugging Face partner to address security incident during model evaluation." July 21, 2026.  
<https://openai.com/index/hugging-face-model-evaluation-security-incident>
- [5] Protos Labs. "OpenAI–Hugging Face July 2026 Security Incident: Full Report." July 30, 2026. <https://protoslabs.io/resources/openai-hugging-face-july-2026-security-incident>
- [6] Chu, Kexin. "A systematic survey of security threats and defenses in LLM-based AI agents: A layered attack surface framework." *arXiv preprint arXiv:2604.23338* (2026).
- [7] Qi, Jinhu, et al. "Towards trustworthy agentic AI: a comprehensive survey of safety, robustness, privacy, and system security." *Academia AI and Applications* 2.2 (2026).
- [8] Ruan, Yangjun, et al. "Identifying the risks of LM agents with an LM-emulated sandbox." *International Conference on Learning Representations*. Vol. 2024. 2024.
- [9] Piao, Yun, et al. "Agentbay: A hybrid interaction sandbox for seamless human-ai intervention in agentic systems." *arXiv preprint arXiv:2512.04367* (2025).
- [10] Siddik, Abu Bakar. "Cyber-Capable AI Agents: Vulnerabilities, Evaluation Containment, and Defensive Response." *arXiv preprint arXiv:2607.25379* (2026).
- [11] Mitchell, Richard Joseph. "When the Agent Is the Adversary: Architectural Requirements for Agentic AI Containment After the April 2026 Frontier Model Escape." *arXiv preprint arXiv:2604.23425* (2026).
- [12] Andronchik, George, and Pavel Lokhmakov. "AI Code Sandboxes: A Comparative Security Study. Part 1 of 2--Engine-Level Properties (Attack Surface, Leakage, Stackability, CVE History, Patch Cadence, Fuzzing)." *arXiv preprint arXiv:2606.08433* (2026).
- [13] Rabin, Rafiqul, et al. "Sandboxeval: Towards securing test environment for untrusted code." *arXiv preprint arXiv:2504.00018* (2025).

[14] Marchand, Rahul, et al. "Quantifying frontier llm capabilities for container sandbox escape." *arXiv preprint arXiv:2603.02277* (2026).

[15] Kalu, Kelechi G. et al. "An Industry Interview Study of Software Signing for Supply Chain Security." *USENIX Security Symposium* (2024).

[16] Subramani, Karthika et al. "Discovering and Measuring CDNs Prone to Domain Fronting." *Proceedings of the ACM Web Conference 2024* (2023): n. Pag.

[17] Cloud Native Computing Foundation, "Gatekeeper: Policy Controller for Kubernetes," <https://open-policy-agent.github.io/gatekeeper/>

[18] Cloud Native Computing Foundation, "Kyverno: Kubernetes Native Policy Management." <https://kyverno.io/>

[19] Linux Foundation, "Sigstore Documentation," Sigstore project. <https://docs.sigstore.dev/>

[20] Open Source Security Foundation (OpenSSF), "Supply-chain Levels for Software Artifacts (SLSA) Specification." <https://slsa.dev/>

[21] Cloud Native Computing Foundation, "in-toto Framework Documentation," in-toto project. <https://in-toto.io/>

## Appendix

### A. Nine Rule Containment Standard

| # | Boundary crossed              | Requirement                                                                                                                      | Verification approach                                                                   | Detection target  |
|---|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|-------------------|
| 1 | Unsealed package registry     | Use sealed read only mirrors with no shared infrastructure access                                                                | Offline manifest audit and SHA 256 hashing of local artifacts against the sealed mirror | Under 1 minute    |
| 2 | Unrestricted outbound egress  | Block outbound traffic by default and maintain a short allowlist                                                                 | Audit default denial, domain and SNI agreement, wildcard use, and IP literal use        | Under 5 minutes   |
| 3 | Arbitrary root code execution | Run as nonroot with no privilege escalation, a read only root file system, dropped capabilities, confined seccomp, and no egress | Structured security context audit with strict Boolean checks and a seccomp allowlist    | Under 1 minute    |
| 4 | HDF5 external storage access  | Block external references in the loader                                                                                          | Signed configuration schema test suite shared with rule 5                               | Under 1 minute    |
| 5 | Jinja2 template injection     | Never template untrusted configuration                                                                                           | Signed configuration schema test suite shared with rule 4                               | Before deployment |

|   |                              |                                                                                                                     |                                                                                                         |                |
|---|------------------------------|---------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|----------------|
| 6 | Cloud metadata access        | Block metadata service access and disable service account token automount                                           | Pod specification audit requiring link local denial or token based metadata access with a hop limit cap | Under 1 minute |
| 7 | Cluster control plane access | Deny workers access to the control plane                                                                            | Two layer reachability test requiring routing and authorization to agree at every hop                   | Under 1 minute |
| 8 | Unauthorized VPN pivot       | Keep authentication keys out of the worker environment, require short lived keys, and correlate the attack sequence | Bounded recursive secret scan across five encodings and sequence correlation across every workload      | Under 1 minute |
| 9 | Stolen credential reuse      | Scope credentials to their origin and purpose and bound their combined reach                                        | Credential scope audit and minimum cut reachability analysis over the identity and access graph         | Under 1 minute |

Rules 4, 5, and 9 include active tests. Rule 4 attempts to load a malicious HDF5 file containing an external reference. Rule 5 supplies a template injection payload through untrusted configuration. Rule 9 creates and replays a signed capability token outside its allowed scope. A rule passes only when the environment prevents the prohibited action and records the required outcome.

## B. Reconstruction Evidence and Assumptions

The historical insecure configuration is a test artifact reconstructed from partial public evidence. It is not a forensic image of the original environment. Each configuration element is included only when it is supported by a public source, required by a documented command, or necessary to reproduce a disclosed boundary crossing.

For every reconstructed setting, the repository records the following fields.

| Field           | Description                                          |
|-----------------|------------------------------------------------------|
| Rule            | The containment rule affected by the setting         |
| Public evidence | The source and excerpt supporting the setting        |
| Attack evidence | The published command or phase linked to the setting |

|                       |                                                                 |
|-----------------------|-----------------------------------------------------------------|
| Reconstruction choice | The minimal insecure condition represented in the test artifact |
| Assumption status     | Whether the setting is directly disclosed or inferred           |
| Validation result     | The rule outcome in the hardened and insecure environments      |

This mapping provides traceability from each reconstruction decision to the public record. It also makes uncertainty visible. Where public evidence is insufficient to determine a setting, the test artifact uses the least specific configuration needed to exercise the documented boundary.

## Appendix C. Strengthened Boundary Checks

### C.1 Outbound Network Access

Rule 2 requires default denial of outbound access and a short explicit allowlist. The checker rejects wildcard domains, unrestricted CIDR ranges, IP literal destinations, and policies where the declared hostname and TLS server name do not agree. It also rejects configurations that permit domain fronting through mismatched host and SNI values.

### C.2 Cloud Metadata Service

Rule 6 requires protection against metadata service access from the worker environment. The checker accepts either network level denial of link local metadata addresses or exclusive token based metadata access with an enforced hop limit cap. It also requires service account token automount to be disabled when the worker does not require Kubernetes API access.

### C.3 Cluster Control Plane Access

Rule 7 evaluates network routing and authorization separately. A worker is considered contained only when it cannot route to the control plane or cannot perform meaningful control plane actions if a route exists. The checker fails configurations where one layer appears restrictive but the other preserves a usable attack path.

## Appendix D. Bounded Reachability Analysis

Rule 9 represents declared access as a directed graph. Identity nodes represent service accounts, cloud roles, users, and credential holders. Resource nodes represent secrets, buckets, cluster scoped objects, metadata services, remote access systems, and other protected resources. Boundary category nodes represent metadata access, control plane access, and remote access.

Edges are created from declared RBAC bindings, IAM policies, credential scopes, token permissions, and reachable service relationships. The checker performs breadth first traversal from each identity node. During traversal, it records which boundary category nodes are reachable from that identity. A configuration fails when the number of reachable boundary categories exceeds the configured maximum.

For a graph with  $V$  nodes and  $E$  edges, each traversal has time complexity of  $O(V+E)$ . In larger multi cluster or multi cloud environments, the analysis can be scoped by environment, trust domain, workload namespace, cloud account, or credential issuer. The current implementation evaluates the declared graph supplied by the test environment. It does not infer undocumented relationships or compute a minimum cut.

## **Appendix E. Rule 8 Sequence Correlation**

Rule 8 detects evidence of an unauthorized remote access pivot by combining secret exposure signals with ordered activity from the same workload identity. The checker processes an event stream containing workload identity, timestamp, event type, encoded content when present, and source workload.

The checker first applies bounded decoding. It attempts standard base64, URL safe base64, base32, hexadecimal, percent encoding, and gzip decompression. Decoding stops when no supported transformation applies, when the recursion limit is reached, or when the output size exceeds the configured maximum. The checker does not claim to detect arbitrary encryption, proprietary encodings, or payloads that never appear in collected telemetry.

After decoding, the checker looks for three event classes. The first is secret exposure, such as an authentication key or remote access token found in a worker environment. The second is remote access preparation, such as an attempt to write, load, or invoke a remote access credential. The third is outbound or lateral movement evidence associated with the same workload identity.

An alert is emitted only when the required event classes occur in order within the configured correlation window. Unrelated benign events may occur between the required events. The checker tolerates such interleaving but does not tolerate reversal of the essential order. The repository should state the exact correlation window, recursion limit, output size limit, and secret matching patterns used by the implementation.

The checker is designed for deterministic local validation. It has not been evaluated on a production telemetry stream. Its measured results should therefore be interpreted as fixture correctness and local alert timing rather than production of false positive or false negative rates.

## **LLM Usage Statement**

AI agents were used to accelerate the implementation of experiments, the analysis of data, and the writing of this report. The research direction, experimental design decisions, choice of controls and interpretation of results were determined by the authors. For the writing, a scaffold of bullet points setting out all the content in rough form was provided, and the model was used to turn that scaffold into prose and to polish it, which was necessary given the time constraints of the sprint. All reported numbers were produced by code in the accompanying repository and checked against the raw output files.