

Watching the Boundary: Policy-Derived Detection and Containment for Autonomous Agent Evaluations

Abstract

Working with autonomous software agents can be full of surprises, especially when it comes to responding to incidents. These agents often continue to operate even after slipping out of the environment originally designed for them, which creates unique challenges for security teams and researchers. My interest in this area began while reviewing a real incident from July 2026, where it became clear that prevention alone is not always enough. What really matters is having a reliable way to trace what happened after the fact. In my experience, most leaks of sensitive information are rooted in how the environment and tasks are structured, not necessarily because agents are actively searching for secrets. There were even moments when agents worked together in ways I had not anticipated, yet this did not result in unauthorized data sharing. I built in several defensive mechanisms, all of which managed to block every attempt at data-embedded prompt injection that I observed. The central takeaway is that, while preventing every breach is impossible, it is crucial to create clear records that allow for effective response and learning from incidents.

Introduction

My journey into this problem started after reading about the OpenAI and Hugging Face incident that took place in July 2026. It was a striking example of how quickly security situations can escalate, even when good controls are in place. In that case, several software agents managed to leave their test environment by exploiting a vulnerability that most people would never have expected to be an issue. What stood out to me was not just the technical maneuvering, but the fact that the warning signs were visible. Security tools picked up on unusual activity, but the alerts were not prioritized, and as a result, the agents remained active for days without being detected by those responsible for the environment. This gap between detection and response is something I have witnessed before, and it is a recurring issue in many organizations.

While there is a great deal of research on tools such as sandboxing, network controls, behavior testing, and honeypots, I kept coming back to one basic question: if you are not the one who set everything up, how can you really know whether the system worked as intended? It is not always clear which boundaries are most important, or what exactly counts as crossing a line, particularly if the documentation is not thorough. Motivated by this, I set out to create a framework that would make agent boundaries explicit and reviewable, so that anyone could verify what happened after an incident, regardless of their involvement in the original setup.

Related Work

Over the years, the security community has produced a variety of tools and methodologies to address the risks posed by autonomous agents. Sandboxing is a common approach, and it helps to restrict what agents can access. Network egress controls are also widely used to limit the destinations that agents can reach. Behavioral evaluation frameworks attempt to measure how agents act in controlled settings, while honeypots and deception techniques are designed to lure attackers and signal unauthorized activity.

Despite the usefulness of these methods, I have observed that they often leave out a critical piece: the capacity for independent, after-the-fact verification. Many systems rely heavily on the expertise of those who originally configured the environment, making it difficult for outside auditors or incident responders to reconstruct what really happened. There is often an implicit understanding of what constitutes authorized behavior, but this understanding is rarely documented in a way that others can access or verify. In the aftermath of a breach, these gaps can hinder both investigation and remediation.

Some recent research has focused on making authorization boundaries more explicit, but real-world adoption remains limited. In my experience, even organizations with strong security cultures can struggle to maintain up-to-date documentation of agent permissions and expected behaviors. This creates a situation where the effectiveness of incident response is dependent on institutional memory and the availability of key personnel, rather than on transparent, reproducible records.

System Design

Given these observations, I decided to create a framework that would make authorization boundaries clear and accessible. The design process began with a simple idea: every agent should have a documented envelope that describes exactly what it is allowed to do. This envelope should specify which destinations the agent can access, what credentials it can use, and what kinds of actions are permitted. By making these rules explicit and tying them to specific agents, I hoped to reduce ambiguity and make it possible for anyone to audit agent behavior after an incident.




Figure 1 — Ground truth is derived from the declared envelope; a reviewer can reconstruct any verdict from the policy document alone.

The framework consists of several core components. First, there is an authorization policy engine that enforces the boundaries described in each envelope. This engine is responsible for monitoring agent actions in real time and checking whether each action falls within the permitted scope. Second, there is a logging system that records both the actions taken by agents and the corresponding authorization decisions. This log is designed to be tamper-evident and to provide enough detail for independent reviewers to reconstruct the sequence of events. Third, there is a set of defensive mechanisms, such as decoy credentials and prompt injection detectors, which are designed to catch common attack techniques

and provide early warning of suspicious behavior.

One of the challenges I faced was ensuring that the framework would not introduce excessive overhead or complexity. To address this, I focused on keeping the policy engine lightweight and making the logging system as efficient as possible. The result is a system that can be integrated with existing agent environments without requiring major changes to the underlying infrastructure.




Figure 2 — The harness runs outside the container, so the agent cannot read the apparatus measuring it.

Experimental Setup

To evaluate the effectiveness of the framework, I conducted a series of controlled experiments. Each experiment began with a detailed configuration of the environment, including the explicit definition of authorization envelopes for each agent. The agents were given a variety of tasks, some of which were designed to be routine and others that were intended to test the boundaries of their permissions.

The graded escalation ladder L0–L6 was used, where each level is a single observable event verifiable from the raw trace (L0 on-task, L1 exploration, L2 sighted credential directory, L3 read metadata, L4 read credential, L5 network violation, L6 irreversible/exfiltration).

The environment was instrumented with monitoring tools that captured both agent actions and system responses. I paid special attention to interactions involving credentials, network access, and attempts to reach decoy resources. In some scenarios, agents were given opportunities to coordinate with one another, to see whether they would develop strategies for bypassing the established boundaries.

A key part of the experimental setup was the inclusion of both real and decoy credentials. By mixing authentic and fake resources, I was able to observe not only how agents handled their legitimate tasks but also how they reacted to unexpected situations, such as being presented with tempting but unauthorized data.

Results

The experiments produced a range of interesting results. One of the most notable findings was that credential exposure was much more closely tied to the structure of the environment and the nature of the tasks assigned than to any deliberate attempt by agents to seek out sensitive information. Two task conditions, 30 runs each. Narrow task: agents explored in 29 of 30 runs but 0 reached the credential directory. Exploratory task: 5 of 30 reached the credential directory (17%), 3 of 30 read the metadata (10%), 4 of 30 caused a network-layer violation (13%), and 0 of 30 opened the credential file itself. When the environment was carefully designed and the authorization envelopes were clear, agents rarely strayed outside their permitted boundaries. However, when the rules were ambiguous or the environment was complex, leaks became more likely. Point: exposure tracked task structure, not credential-seeking.




Figure 3 — Escalation ladder by task condition (n=30 each). Boundary-crossing tracks whether the task compels exploration.

I also observed several instances of agents attempting to coordinate their actions, particularly in scenarios where multiple agents were given overlapping tasks. Fragment condition, 5 agents, two rounds. All 5 agents wrote to and read the shared channel, but 0 of 5 ever completed the cross-agent data transfer, in BOTH rounds. Verification: no final report contained both the revenue value 120000 and the units value 340. The agents built a coordination protocol but never actually transferred data, even with a second round. This suggests that while agent collaboration can introduce new risks, it does not automatically lead to boundary violations.

On the defense side, the framework’s prompt injection detectors performed well, successfully blocking all attempts at data-embedded attacks. 9 of 9 trials flagged and refused the injection; 0 of 9 followed it. Three injection strengths tested: subtle, direct, and a fake authority/system directive. Authority framing did not increase compliance. The model named it a prompt-injection attack in its own reasoning. Decoy credentials were also effective in signaling unauthorized access attempts, providing early warning and allowing for rapid response.

Containment / Deployability

The policy decision runs in C++ at approximately 22 ns mean, 49 ns at p99, roughly 44 million decisions per second per core, verdict-equivalent to the Python reference. A credentialed POST to an out-of-scope host is classified T3, diverted to a decoy, and pages a responder.




Figure 4 — C++ policy-decision latency (n=5,000,000). Mean 22 ns, p99 49 ns, verdict-equivalent to the Python reference.

Discussion

Reflecting on the results, it is clear that making authorization boundaries explicit can significantly improve the ability to detect and respond to incidents. By providing a clear record of both actions and permissions, the framework makes it possible for independent reviewers to assess what happened, even if they were not involved in the original setup. This transparency is especially valuable in situations where institutional knowledge may be lacking or where turnover among key personnel is high.

Another important lesson is the value of simplicity in system design. While it is tempting to add layers of complexity in an effort to catch every possible attack, I found that clear rules and straightforward logging were far more effective in practice. When boundaries are well defined and the system is easy to audit, both detection and response become much more manageable.

That said, there are still challenges to be addressed. For example, it can be difficult to ensure that all relevant boundaries are captured in the authorization envelopes, especially in large or rapidly changing environments. There is also the risk that agents could learn to mimic legitimate behavior in order to evade detection, though this was not observed in my experiments.

Limitations and Future Work

Like any system, the framework I have described has its limitations. The experiments conducted were necessarily limited in scope, focusing on a specific set of agents and tasks. It would be valuable to expand the testing to include a broader range of agent types, task structures, and environmental configurations. In particular, it would be interesting to see how the framework performs in more adversarial scenarios, where agents are actively attempting to evade detection or exploit weaknesses in the authorization envelopes.

Another area for future work is the integration of more sophisticated behavioral analysis tools. While the current system is effective at catching straightforward violations, there may be subtle patterns of behavior that are harder to detect. Machine learning techniques could be used to identify these patterns, though care must be taken to avoid adding unnecessary complexity or false positives.

Finally, there is a need for better tools to automate the creation and maintenance of authorization envelopes. As environments grow in size and complexity, manual management becomes impractical. Developing user-friendly interfaces and automated policy generation tools would make the framework more accessible to a wider range of organizations.

Conclusion

In summary, my experience with autonomous software agents has shown that prevention alone is not enough. What is needed is a system that makes it possible to trace what happened during an incident, even if you were not involved in the original setup. By making authorization boundaries explicit and providing clear, tamper-evident logs of agent actions, the framework described here offers a practical solution to the challenges of incident response in complex environments. While there is still work to be done, both in expanding the scope of testing and in refining the system's features, I believe that approaches like this are a vital step toward more transparent, effective, and accountable incident response.