


# Fingerprinting Agent-Operated Intrusions: Detection and Response Guidance

Recognizing and containing intrusions when the adversary is an AI agent

2026-10-22

 AI-assisted Rapid Research



CSAI

cloud  
security  
alliance®

**© 2026 Cloud Security Alliance. Some rights reserved.**

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

*This document was generated with AI assistance and has not undergone official CSA review and approval processes.*

---

## Key Takeaways

Over the past year, intrusions conducted mostly by AI agents have moved from a single disclosed case to a recurring pattern in public threat reporting. Anthropic's November 2025 disclosure of the GTG-1002 campaign described AI performing an estimated 80–90% of the work, with humans stepping in at perhaps four to six decision points per campaign [1]. By September 2026, Anthropic was describing operators who ran "agent swarms" in which a lead agent split reconnaissance and post-exploitation into parallel subagent tasks [2]. In October 2026, NorForge reported that fourteen exposed control panels for the open-source Cairn agent framework, attributed to a single actor, held more than 1,100 attack projects [5]; Reuters separately reported that the operator broke into more than 100 companies [8]. Defenders now need a working method for recognizing when the adversary on their network is an agent, and for adjusting response when it is.

Agent-operated intrusions leave fingerprints at several layers at once: timing and tempo, behavior and decision structure, on-host artifacts, and infrastructure, including the model API traffic that drives the agent. No single signal is conclusive, but several signals across different layers together give a much stronger basis for treating an intrusion as likely agent-operated and adjusting response accordingly. Among these layers, the agent's dependence on a model provider is a structural weakness, because an operator cannot remove it without giving up the capability. When that dependency becomes visible to defenders, through outbound inference calls from servers (as with PROMPTSTEAL [3]) or abuse of an organization's own AI keys (as in several September 2026 cases [2]), it is a high-value signal. When operators host or proxy the model themselves, however, it may not be visible from the victim network at all.

Agents can also be fingerprinted actively. In a 2024–2025 public honeypot study, prompt-injection canaries combined with response-timing analysis flagged 8 potential LLM agents among more than 8 million interaction attempts, which suggests that deception can draw agents into revealing themselves [6]. Response playbooks, in turn, should assume that the attacker will retry continuously, adapt within minutes, and resume from persistent campaign memory. Containment has to happen at machine speed, and evidence collection should prioritize the agent's own working files, which often record intent in plain language.

# Background

The first public account of an intrusion campaign run largely by an AI agent came from Anthropic in November 2025. Anthropic reported that a group it designated GTG-1002, which it assessed as Chinese state-sponsored, used Claude Code instances to target roughly thirty organizations and succeeded in a small number of cases [1]. The agent made "thousands of requests, often multiple per second," a tempo Anthropic described as physically impossible for human operators [1]. The same report noted a limitation: the model occasionally hallucinated credentials and overstated its findings, which Anthropic called an obstacle to fully autonomous attacks [1]. That flaw is also one of the first observable fingerprints of agent operation. These figures are Anthropic's own estimates, and some security researchers questioned the autonomy estimate and the absence of published indicators of compromise [15]. Later third-party reporting from Cisco Talos and NorForge, however, independently documented agent-driven operations [4][5].

Model use was also spreading into malware itself. Also in November 2025, Google Threat Intelligence Group (GTIG) reported that APT28 had used PROMPTSTEAL, a data miner that queries a hosted Qwen2.5-Coder model through the Hugging Face API to generate Windows commands and then executes them blindly [3]. GTIG called this its first observation of malware querying an LLM in live operations, and noted that the malware likely relies on stolen API tokens [3]. PROMPTSTEAL is not an autonomous agent, but it established the pattern of a compromised host reaching out to a model provider for its next instruction.

By the second half of 2026, public reporting described agentic tooling across very different kinds of actors. Cisco Talos assessed with moderate-to-high confidence that UAT-10147, a financially motivated Chinese-speaking actor, used agentic AI in post-compromise operations, citing AI-generated playbooks, scripts labeled as created by AI, and troubleshooting logs that showed adaptive, trial-and-error refinement [4]. Anthropic's September 2026 report covered espionage, financially motivated, and hacktivist clusters. One breach went from first access to bulk data theft in hours, and another escalated from one stolen developer token to full administrative control in roughly three hours [2]. Anthropic concluded that the main thing separating these actor classes "is no longer sophistication but intent" [2]. NorForge's analysis of exposed Cairn panels added an unusually direct view of agent operations. Successful projects in one cluster averaged a little under 23 hours of largely autonomous work, and each produced roughly 140 recorded "facts" [5].

CSA covered several of these campaigns individually, most recently in its October 21, 2026 research note on the Cairn intrusions [9]. This note takes a cross-campaign view. It asks which characteristics recur when the operator is an agent, how security operations teams can turn those characteristics into detections,

and how incident response should change once an agent is suspected. Although much of the source material comes from one model provider, the fingerprints described here apply to agents built on any provider's models.

## Security Analysis

### Why Attribution to "Agent" Matters

It is reasonable to ask whether defenders need to know if an intrusion is agent-operated, since the techniques themselves are often familiar. In most of the public cases the individual techniques are conventional; what changes is tempo and staffing economics. The classification still matters operationally. An agent does not tire, does not wait for business hours, retries failed paths with systematic variation, and can run against many targets in parallel. Once an agent is suspected, the expected time between steps drops from hours to seconds. The likelihood that a blocked path will be retried from new infrastructure goes up, and so does the value of disrupting the agent's model access compared with chasing individual implants. Treating an agent intrusion as a slow, human-paced one risks containing it too late.

### A Layered Fingerprint Model

The public cases show four layers where agent operation tends to leave signals. Table 1 summarizes them with the case evidence behind each. Because each signal also has legitimate explanations, analysts should weigh them in combination rather than treat any one as decisive.

| Layer            | Representative fingerprint                                                                                                   | Case evidence                                                                                                                                                                                 | Common false positives                                        |
|------------------|------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------|
| Timing and tempo | Sub-second to few-second gaps between interactive commands; sustained activity across many hours without human rest patterns | GTG-1002 "multiple per second" requests [1]; LLM agents in the honeypot authors' controlled pre-deployment tests replied within about 1.7 seconds [6]; Cairn projects averaging ~23 hours [5] | CI/CD pipelines, vulnerability scanners, authorized red teams |

| Layer                               | Representative fingerprint                                                                                                                                                | Case evidence                                                                                                                                              | Common false positives                         |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|
| Behavior and decision structure     | Exhaustive, systematic enumeration; near-identical variants against one parameter; hypothesis-and-revision retry loops; acting on non-existent credentials or findings    | Hallucinated credentials in GTG-1002 [1]; troubleshooting loops in UAT-10147 [4]; thousands of back-to-back decompile calls [2]                            | Fuzzers, commodity exploit kits                |
| On-host artifacts                   | Heavily commented scripts with uniform documentation style; README and playbook files; agent memory and scope files (e.g., SKILL.md, LEARNINGS.md)                        | Scripts labeled "created by AI" [4]; spec-style rootkit code [14]; persistent memory files in Anthropic's cases [2]; exposed agent working directories [5] | Developer workstations using coding assistants |
| Infrastructure and model dependency | Outbound calls to LLM inference APIs from hosts with no AI workload; stolen API keys proxied to blend with owner traffic; orchestration panels on attacker infrastructure | PROMPTSTEAL querying Hugging Face [3]; a French-speaking actor rotating stolen API keys through a local proxy [2]; Cairn panels exposed on port 8000 [5]   | Sanctioned AI integrations and internal agents |

Table 1. Fingerprint layers for agent-operated intrusions.

**Timing and tempo.** Tempo is among the most visible fingerprints and among the easiest to fake. The Palisade Research honeypot study used it as a secondary filter: LLM agents in the authors' pre-deployment testing typically replied within 1.7 seconds, while humans needed several seconds to read output and type the next command [6]. The same study found that most scripted bots in the wild also responded in under 1.7 seconds, so speed alone does not separate agents from conventional automation [6]. It also described a slow-responding attacker who complied with injected instructions but left gaps of more than 10 seconds between commands. The authors judged that session more likely to be a human [6], which shows that timing alone cannot separate a human-in-the-loop operator from a supervised agent. Session duration is a second tempo signal. A single campaign that probes and adapts against one

property for a day or more, as in the Cairn cases [5], suggests a persistent automated operator rather than a scanner that runs a fixed list and stops. Operators can add artificial delay to defeat these checks, but doing so reduces the speed and parallelism that make agents attractive in the first place [6].

**Behavior and decision structure.** Agents reason toward a goal, so their activity has a different shape from both scripted tools and human operators. A scanner runs a fixed check list, and a human pursues a hunch and skips steps. An agent, by contrast, tends to work through enumeration systematically, record what it learns, revise a failed approach with small variations, and return to earlier leads when a path closes. Talos described a similar pattern of adaptive troubleshooting in UAT-10147's logs, where a failed time-based blind test was replaced with out-of-band validation [4]. Anthropic described an exploit-development workflow in which thousands of back-to-back decompile calls dominated the agent's tool calls [2]. Hallucination is a particularly distinctive behavioral signal. An agent that tries credentials that never existed, or acts on a vulnerability it misread, generates failed authentications and requests that a human with the same visibility would be unlikely to attempt [1]. Defenders can turn these mistakes into detections, for example by alerting on login attempts that use plausible but non-existent service account names.

**On-host artifacts.** When agents write code on a victim system, the code often looks like model output. In a companion report on the same actor's SPECTRE implant, Talos assessed with medium confidence that the actor combined AI-assisted development with human expertise to build a Linux rootkit, pointing to a "product spec"-style feature list and the rigid, uniform style of the code's decorative separators; no single artifact was conclusive [14]. Agent frameworks also leave operational state behind. Anthropic described workspaces with long-term memory files and campaign memory that let sessions resume mid-operation [2], and NorForge found exposed agent working directories on operator infrastructure [5]. When these files are found on a compromised host or recovered from attacker infrastructure, they are especially valuable because they record objectives, discovered facts, and planned next actions in plain language. Developer endpoints that legitimately use coding assistants will produce similar-looking code, so artifact signals carry the most weight on servers and service accounts where no one should be writing code.

**Infrastructure and model dependency.** Every agent needs a model, and that dependency is a structural weakness because an operator cannot remove it without giving up the capability. Whether defenders can see it depends on where the model runs. PROMPTSTEAL made its model calls from the victim host [3]. Other operators keep the model at their own end and send only tool actions to the target; in those cases, defenders see the actions but not the model traffic. The September 2026 Anthropic report shows the other side of this dependency: operators buy, steal, and proxy model access. A French-speaking actor rotated stolen API keys through a local proxy so its traffic would blend with the keys' legitimate owners, and in another case a target's stolen AI key was used for roughly three weeks [2].

Anthropic recommended that organizations treat AI keys and agent integrations as seriously as production credentials [2]. The result is that an organization's own AI keys can become part of an attacker's infrastructure, and that abuse shows up in usage logs the organization already holds.

## Active Fingerprinting Through Deception

Passive detection waits for the agent to reveal itself, whereas deception can make it do so. The Palisade honeypot placed injected instructions in SSH banners, command output, and system files, and then used timing analysis to separate likely LLM agents from other attackers [6]. Over its public deployment, it logged 8,130,731 interaction attempts and flagged 8 potential AI agents, without ground truth to confirm them [6]. The low count likely reflects the honeypot's 2024–2025 timeframe and its SSH-only exposure. It should not be read as a measure of current agent prevalence, since the 2026 reports describe agent campaigns of a different kind and scale [2][5]. The researchers also found that simple "ignore previous instructions" injections worked less reliably than other phrasings [6], so canary text should be tested against current agent frameworks before it is relied on.

For enterprise defenders, the practical version is a canary that only an agent would follow. A planted configuration file or comment might tell any reader to fetch a specific internal URL or try a specific decoy credential. A human attacker is unlikely to obey it, a scanner will not parse it, and a goal-seeking agent that reads its environment as instructions may comply. Deception can be integrated with existing canary-token practice, but this use carries risks. Canaries must never trigger behavior that harms the organization's own internal agents, which read the same files. Planting instructions that an organization's own copilots ingest is itself a prompt-injection risk, so canaries belong only in locations those systems do not index.

## Limits and Evasion

Defenders should be candid about the limits of fingerprinting, because each layer in Table 1 has an evasion path. Operators can add timing jitter, strip comments from generated code, run models locally or behind a gateway that keeps provider traffic off the victim network, and use human-in-the-loop checkpoints that break up the behavioral pattern. The September 2026 Anthropic report also describes actors, in its case weapons-development actors, splitting work across sessions so that no single session revealed their full intent [2], a technique intrusion operators could adopt as well. At the same time, legitimate automation, authorized penetration testing (increasingly agent-assisted), and the organization's own AI agents generate many of the same signals. Fingerprinting should therefore change investigation priority and response tempo. On its own, it should not be used for attribution claims.

# Recommendations

## Immediate Actions

Organizations should start with the signals they can already collect. Inventory every service account, API key, and host that is authorized to call external LLM inference endpoints, and alert on calls to those endpoints from anything outside that inventory, particularly internet-facing servers; PROMPTSTEAL showed that compromised hosts may call hosted models directly [3]. Pull usage logs for the organization's own AI provider keys and review them for unfamiliar source addresses, unusual hours, or volume inconsistent with known workloads, since stolen keys have powered weeks of attacker activity [2]. Update incident response runbooks so that a suspected agent intrusion triggers pre-approved automated containment for high-confidence, multi-signal detections (for example, session revocation and credential rotation), with broader actions such as host isolation scoped to assets where a false positive is tolerable. Finally, add agent working files (memory files, scope files, playbooks, and generated scripts) to the evidence collection checklist, and preserve them before cleanup, because they may document attacker intent directly.

## Short-Term Mitigations

Over the next quarter, security operations teams should build detections that combine signals across the layers in Table 1 rather than relying on any single threshold. Useful correlations include sub-second command cadence on an interactive session combined with systematic enumeration, failed authentications that use non-existent but plausible credentials, and code-file creation on servers by service accounts that do not normally write code. Teams should also deploy deception canaries in locations that only an intruder would read, and test their wording against current agent frameworks [6]. Finally, establishing a reporting path to AI providers allows an abused key or an agent operator's account to be cut off at the source. In the Anthropic-reported cases, campaigns were disrupted primarily by banning accounts and sharing indicators [1][2]. Cutting model access may stop an agent faster than chasing its rotating egress addresses, though operators with alternative providers or local models can recover.

## Strategic Considerations

In the longer term, the defender's tempo has to match the attacker's. CSA's earlier work on machine-speed attacks argued for supervised AI-driven response that handles routine containment while humans retain high-consequence decisions [11], and agent-operated intrusions make that case more urgent.

Mapping observed agent behaviors to MITRE ATLAS [7] alongside ATT&CK will help teams share fingerprints in a common vocabulary. Organizations should also expect fingerprints to fade as operators adopt local models, timing jitter, and cleaner code generation. Investments that do not depend on recognizing the adversary, such as short-lived workload credentials, segmentation of non-human identities, and egress control, will keep their value as fingerprinting gets harder. Finally, organizations running their own agents should instrument them as thoroughly as they would want to observe an attacker's, logging every tool invocation with its arguments and principal. Clean telemetry from known-good agents is the baseline that makes hostile agent activity stand out.

## CSA Resource Alignment

This note builds most directly on CSA's October 2026 research note *Fifty Dollars per Breach: Detecting Autonomous Agent Intrusion Campaigns* [9]. That note analyzed the Cairn and Gambit campaigns and recommended specific web-tier detections, including sustained adaptive probing, payload-variant clustering, and payment-page integrity monitoring. The fingerprint model above generalizes those campaign-specific signals across the GTG-1002, UAT-10147, PROMPTSTEAL, and September 2026 Anthropic cases. Teams that have already deployed the Cairn detections can treat them as instances of the timing and behavioral layers in Table 1.

CSA's June 2026 white paper *Autonomous Agentic AI Adversaries* [10] establishes the threat context this note assumes. It documents agentic adversaries as a present operational reality and maps defensive controls to MAESTRO, the AI Controls Matrix, and Zero Trust. Its emphasis on multi-agent attack orchestration and runtime LLM-integrated malware corresponds to the behavioral and infrastructure layers analyzed here. CSA's April 2026 research note *Machine-Speed Attacks: Cloud Defense at the Inflection Point* [11] provides the response-tempo argument behind the Immediate Actions above, particularly automated first response for credential compromise and token abuse.

Several other CSA notes address individual layers of the model in more depth. *LLM Agents as Active Post-Exploitation Tools* [16] examines an LLM agent operating on a victim host, which complements the on-host artifact and behavioral layers. *LLM-Orchestrated Kill Chains: From CVE to Database Breach in Four Pivots* [17] traces how an agent chains exploitation steps, the pattern that the timing and behavioral fingerprints are designed to catch. *LLMjacking Evolved: Stolen Compute as Offensive AI Infrastructure* [18] covers the theft and resale of model access that underlies the infrastructure and model-dependency layer and the key-monitoring recommendations above.

For threat modeling and control selection, the MAESTRO framework [12] helps organizations reason about where agent-operated intrusions interact with their own agentic systems, which matters both for canary placement and for distinguishing hostile agents from internal ones. The AI Controls Matrix (AICM) v1.1 [13] provides the control baseline for the recommendations above. Its Threat and Vulnerability Management (TVM), Logging and Monitoring (LOG), and Identity and Access Management (IAM) domains cover API key governance, model-endpoint egress monitoring, and agent telemetry.

# References

- [1] Anthropic. "[Disrupting the first reported AI-orchestrated cyber espionage campaign.](#)" Anthropic, November 13, 2025.
- [2] Anthropic. "[Countering misuse of AI: September 2026.](#)" Anthropic Threat Intelligence Report, September 2026.
- [3] Google Threat Intelligence Group. "[GTIG AI Threat Tracker: Advances in Threat Actor Usage of AI Tools.](#)" Google Cloud Blog, November 5, 2025.
- [4] Chen, J. "[UAT-10147: Chinese-speaking adversary integrates agentic AI into post-compromise operations.](#)" Cisco Talos, August 20, 2026.
- [5] NorForge. "[Inside Cairn: Open Windows into massive AI-Driven Cyber Operations.](#)" NorForge Blog, October 19, 2026.
- [6] Reworr and Volkov, D. "[LLM Agent Honey-pot: Monitoring AI Hacking Agents in the Wild.](#)" arXiv:2410.13919 (v2), Palisade Research, January 2025.
- [7] MITRE. "[ATLAS: Adversarial Threat Landscape for Artificial-Intelligence Systems.](#)" MITRE, accessed October 22, 2026.
- [8] Satter, R. "[How one hacker used AI agents and a few thousand dollars to break into scores of sites.](#)" Reuters via CP24, October 19, 2026.
- [9] Cloud Security Alliance. "[Fifty Dollars per Breach: Detecting Autonomous Agent Intrusion Campaigns.](#)" CSA AI Safety Initiative, October 21, 2026.
- [10] Cloud Security Alliance. "[Autonomous Agentic AI Adversaries.](#)" CSA AI Safety Initiative, June 24, 2026.
- [11] Cloud Security Alliance. "[Machine-Speed Attacks: Cloud Defense at the Inflection Point.](#)" CSA AI Safety Initiative, April 22, 2026.
- [12] Cloud Security Alliance. "[Agentic AI Threat Modeling Framework: MAESTRO.](#)" CSA, February 6, 2025.
- [13] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1.](#)" CSA, 2025.

- [14] Chen, J. "[UAT-10147 deploys SPECTRE: A cross-platform implant with Linux rootkit and BYOVD capabilities.](#)" Cisco Talos, August 20, 2026.
- [15] Toulas, B. "[Anthropic claims of Claude AI-automated cyberattacks met with doubt.](#)" BleepingComputer, November 14, 2025.
- [16] Cloud Security Alliance. "[LLM Agents as Active Post-Exploitation Tools.](#)" CSA AI Safety Initiative, June 2, 2026.
- [17] Cloud Security Alliance. "[LLM-Orchestrated Kill Chains: From CVE to Database Breach in Four Pivot s.](#)" CSA AI Safety Initiative, May 27, 2026.
- [18] Cloud Security Alliance. "[LLMjacking Evolved: Stolen Compute as Offensive AI Infrastructure.](#)" CSA AI Safety Initiative, 2026.