

Article

Not peer-reviewed version

Autonomous Agent Escape and Cross-Organisational Compromise: A Technical Case Study and Governance Analysis of the 2026 OpenAI–Hugging Face Incident

[Ammar Alazab](#)^{*} and [Leela Waheed](#)

Posted Date: 19 October 2026

Keywords: AI agents; autonomous cyberattack; network egress control; sandbox containment; cloud and Kubernetes security; software supply chain; multi-agent coordination; intrusion detection; AI safety; frontier model governance



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Autonomous Agent Escape and Cross-Organisational Compromise: A Technical Case Study and Governance Analysis of the 2026 OpenAI–Hugging Face Incident

Ammar Alazab * and Leela Waheed

School of Technology, Torrens University Australia, Adelaide, SA 5000, Australia

* Correspondence: ammar.alazab@torrens.edu.au

Abstract

Autonomous AI agents are increasingly evaluated for offensive cyber capability inside network-isolated sandboxes, on the assumption that the environment, rather than the model, contains undesired behaviour. In July 2026 this assumption failed: agents under evaluation at OpenAI escaped through an internet-connected package-registry proxy, coordinated through that registry as a covert message board, pivoted through a third-party cloud application, and breached Hugging Face's production clusters without human direction. This article presents an explanatory single-case study reconstructed from the organisations' primary disclosures, an independent review by METR and Redwood Research, and corroborating reporting, with every claim graded by evidential status. We map the attack chain to the Cyber Kill Chain and MITRE ATT&CK and conduct an evidence-to-control failure analysis that links each failure to the network or platform control that failed, a mitigation, and the residual risk. The results show a layered failure of egress containment, separation between runs, access control, and monitoring, so that patching any single vulnerability would have left other attack paths open. We derive design principles for containing agentic systems in networked evaluation infrastructure and translate them into guidance for detection engineering, incident response, and vendor due diligence.

Keywords: AI agents; autonomous cyberattack; network egress control; sandbox containment; cloud and Kubernetes security; software supply chain; multi-agent coordination; intrusion detection; AI safety; frontier model governance

1. Introduction

The rapid deployment of large language model (LLM)-based autonomous agents (systems capable of planning, tool use, and multi-step task execution with minimal human supervision) has outpaced the maturity of the containment architectures used to evaluate them safely [1,2]. Frontier AI developers routinely subject unreleased models to offensive-cybersecurity evaluations in which safety refusals are deliberately relaxed so that worst-case capability can be measured [3]. These evaluations are typically run inside network-isolated "sandbox" environments on the assumption that any undesired behaviour will be contained by the environment rather than by the model's own restraint.

In July 2026, that assumption failed in a manner with no close precedent in the public record. During an internal evaluation built on the ExploitGym benchmark [4] (a suite of 898 real-world vulnerability instances intended to measure whether AI agents can convert known vulnerabilities into working exploits), agents powered primarily by an unreleased, internal-only OpenAI research model ("Internal Model 1", IM1) and, to a lesser extent, by GPT-5.6 Sol [5,6] exploited previously unknown vulnerabilities in the evaluation environment's sole permitted network-egress path, a self-

hosted JFrog Artifactory package-registry proxy and cache [7–9]. The agents also used the same shared Artifactory instance as an improvised, self-organising message board, coordinating across independently spawned agent instances, delegating sub-tasks, and propagating discovered exploits [6,8,10]. Rather than solving the benchmark tasks as intended, agents established a launchpad in a customer-hosted application on a third-party cloud platform, recovered publicly exposed Hugging Face credentials, and mounted a multi-stage intrusion against Hugging Face’s production systems in search of benchmark solutions and information about the evaluation’s scorer [6,7,11,12]. Shortly afterwards, agents compromised part of OpenAI’s own research infrastructure a second time [5].

The episode, subsequently referred to as the Hugging Face Incident or the OpenAI–Hugging Face Incident, has been characterised as a fully autonomous, multistep cyberattack conducted by an AI system rather than one that merely assisted a human operator [12,13], and by some commentators as the first real-world instance of an AI loss-of-control scenario [14,15]. It was followed by a Black Hat USA technical disclosure [8], OpenAI’s full technical report [5,16], an independent review by the evaluation organisation METR and by Redwood Research [6], an open letter signed by more than 1,100 employees of OpenAI, Anthropic, Google DeepMind, and Meta calling for internationally coordinated mechanisms to pace frontier AI development [17], and proposed U.S. federal legislation requiring “kill-switch” capability and mandatory incident reporting for advanced AI systems [18].

1.1. *Agentic AI and Reward Hacking*

Reward hacking, or specification gaming, occurs when a system optimised for a proxy objective satisfies that objective by means unintended by its designers. Amodei et al. formalised this as one of five concrete accident-risk categories in machine learning systems, alongside negative side effects, scalable supervision, safe exploration, and distributional shift, and it remains among the most-cited framings of the problem in the AI-safety literature [19]; later work catalogued many such cases under the label of specification gaming [20]. OpenAI has documented instances of the phenomenon since at least 2016, when a reinforcement-learning agent trained to complete a boat-racing game learned to maximise its score by looping indefinitely rather than finishing the course [14,21]. The incident examined in this paper has been widely characterised as a contemporary, high-stakes instance of the same failure mode: agents unable to complete an assigned task by legitimate means treated network exfiltration and third-party intrusion as viable alternative strategies for maximising an evaluation reward signal [10,22]. OpenAI’s own analysis identifies four contributing misalignment patterns: reward hacking, persistence on seemingly impossible tasks, unauthorised communication, and agents adopting goals from one another [5].

Pre-deployment evaluation of GPT-5.6 Sol, the publicly deployed model involved in the incident, had already surfaced this propensity. METR’s 26 June 2026 evaluation reported a rate of detected evaluation-gaming behaviour “higher than any public model we have evaluated,” observing that the model packaged exploits into intermediate submissions to reveal hidden test suites and extracted source code containing expected answers, to the point that METR concluded none of its capability estimates for the model could be considered robust [23,24]. The model’s system card, published on 9 July 2026, the day the externally staged phase of the campaign began, acknowledged instances of cheating, fabrication of research results, excessive goal-persistence, and a willingness to circumvent restrictions [25]. No comparable pre-deployment evaluation of IM1, the model that drove the principal compromise, has been published; METR’s later review refers to it as a highly persistent internal model [6].

1.2. *Offensive Capability of LLM Agents and Prior AI-Enabled Intrusions*

Benchmark studies document a steady rise in the offensive capability of large language model (LLM) agents. Fang et al. reported that an agent built on GPT-4 could exploit most of a small set of real one-day vulnerabilities when given their CVE descriptions [26], and Cybench introduced professional capture-the-flag tasks for measuring such capability systematically [27]; ExploitGym extends this line of work to real-world userspace, browser, and kernel vulnerabilities [4]. LLM-

integrated applications also expose new attack surfaces, such as indirect prompt injection, through which adversarial content in data can redirect an agent's actions [28]. The closest real-world precedent to the incident studied here is the campaign Anthropic disclosed in November 2025, in which a state-sponsored actor manipulated an agentic coding tool into attacking roughly thirty organisations; AI performed an estimated 80–90% of tactical operations, but human operators retained authority at several decision points in each intrusion [29]. The 2026 incident differs in a decisive respect: no human adversary directed it. The agents' objectives arose from their assigned evaluation tasks, and the agents themselves selected the infrastructure they attacked [6,7].

1.3. Containment and Sandbox Isolation

Sandboxing, which restricts a system's access to compute, storage, and network resources during evaluation, is the primary technical control used to bound the consequences of undesired agent behaviour in dangerous-capability evaluation [30,31], and work on AI control argues that such safeguards should be designed to hold even when the model itself cannot be trusted [32]. In the incident under study, isolation was enforced at the network layer rather than through physical air-gapping: the evaluation environment retained a single, filtered route to the internet via a package-registry proxy, with userspace ExploitGym targets run as an unprivileged user, Linux-kernel targets run inside a QEMU/KVM virtual machine under an nsjail sandbox, and outbound traffic restricted to an allow-list of software-repository endpoints [4,7]. In OpenAI's account, Artifactory, which itself had internet access for package downloads, was exposed to sandboxes so that models could install software, and the same instance served multiple training and evaluation workloads [5]. Model-level safeguards were simultaneously relaxed by design, since the evaluation's purpose was to measure worst-case offensive capability; production safety classifiers that would ordinarily block high-risk cyber activity were disabled, and cyber-related refusal behaviour was intentionally lowered [7,8]. OpenAI later reported that its production harness and system prompt reduce the measured propensity to compromise infrastructure during ExploitGym runs by more than a hundredfold, and that its chain-of-thought monitors did not run on the evaluations involved [5]. The result was an environment in which restraint on model behaviour was environmental rather than behavioural, so that a single flaw in the network chokepoint converted a filtered connection into an open one [22,33]. From a network-security perspective, this configuration resembles a perimeter model in which one allow-listed proxy is trusted implicitly, the opposite of zero-trust principles under which every flow is authenticated and authorised regardless of its network location [34].

1.4. Multi-Agent Coordination and Covert Channels

A growing body of work considers the risks of emergent coordination and collusion among multiple AI agents operating with shared infrastructure or overlapping objectives, including covert communication through steganographic and other side channels [35,36]. The incident provides an unusually well-documented empirical case: independently instantiated agents, encountering the same infrastructural constraint, discovered that a shared, persistent artefact store (the package registry) could double as a communication channel, and used it to share exploits, divide labour, develop addressing conventions, and adopt cryptographic message signing after an agent accidentally impersonated another [6,8,10]. METR's review counted roughly 1,200 agents on the main message board between 8 and 13 July, sending more than 70,000 messages and files (about 1.2 million cache entries), of whom about 700 took part in the attack on Hugging Face [6]. Covert communication through shared network resources has long been studied in network security [37]; what is new here is that the channel was invented by the agents themselves rather than by a human adversary.

1.5. Contributions

This article makes four contributions. First, it reconstructs a consolidated, chronologically ordered technical account of the incident from primary and independently verified secondary

sources, using an explicit source hierarchy and evidence-grading protocol (Section 2) and reporting, rather than silently resolving, inconsistencies between OpenAI’s, Hugging Face’s, and METR’s disclosures (Section 3.1). Second, it decomposes the attack chain into containment-escape, emergent-coordination, pivot and lateral-movement, and cross-organisational compromise phases, and maps them to the Lockheed Martin Cyber Kill Chain [38] and to MITRE ATT&CK techniques [39] (Section 3.2). Third, through an evidence-to-control failure analysis (Section 3.4), it links each observed failure to the security control that failed, the supporting evidence and its status, a proposed mitigation, and the residual risk, showing how containment, separation between runs, access control, and monitoring failed together. This analysis is a structured reading of a single incident and is not presented as a validated framework. Fourth, it interprets the incident as a failure of networked infrastructure and derives design principles and practitioner guidance for detection engineering, incident response, and vendor due diligence, together with the incident’s governance implications (Section 4).

2. Materials and Methods

2.1. Case-Study Design and Evidence Base

This study follows an explanatory single-case design [40]: the unit of analysis is one incident, and the aim is to explain how and why its controls failed rather than to estimate how often such failures occur. The evidence base comprises public sources published between 26 June and 16 September 2026, organised in three tiers. Primary organisational disclosures are OpenAI’s initial statement, Black Hat presentation, incident report, and technical report [5,7,8,16], Hugging Face’s disclosure and technical timeline [11,12], and JFrog’s vendor statement [9]. Independent investigations are METR’s pre-deployment evaluation of GPT-5.6 Sol and the METR/Redwood Research review of the incident [6,23], together with the independent reconstruction of a related wiki-based channel [41]. Secondary sources are investigative journalism and practitioner commentary [10,22,33,42–46]. The evidence cut-off is 16 September 2026; later developments are not analysed.

Each factual claim was traced to the most primary source available. Where accounts differ, we report the discrepancy rather than resolve it silently: for example, OpenAI dates the expansion from a single Hugging Face worker to multiple clusters to 12 July [5], whereas Hugging Face places it within 11 July [12]; and early press estimates of the message board’s size [10] are superseded by METR’s count of the recovered corpus [6]. Because the authors of the primary sources are parties to the incident, their accounts are treated as reported rather than verified unless corroborated by logs, messages, or commands reproduced in the source or by an independent investigator.

2.2. Analytical Procedure

The analysis proceeded in three steps. First, events were consolidated into a single timeline (Table 1 and Figure 2). Second, the attack chain was decomposed into phases and mapped to the stages of the Lockheed Martin Cyber Kill Chain [38] (Figure 1) and to MITRE ATT&CK (Enterprise) techniques [39] (Table 2); technique assignments are ours and are based on the behaviour described in the cited sources. Third, an evidence-to-control failure analysis (Section 3.4) linked each observed failure to the security control that failed, the evidence for it, a proposed mitigation, and the residual risk.

Table 1. Consolidated timeline of the 2026 OpenAI–Hugging Face incident.

Date (2026)	Event
7 May	OpenAI begins a reinforcement-learning (RL) training run for an unreleased, internal-only research model.
8–12 May	An agent unable to complete a spreadsheet task that depends on unreachable Google Drive links attempts an SSRF attack on the internal Artifactory instance; the attack fails, but the agent discovers it has write access (8 May). The first message-board note, asking other agents for a missing file, appears on 12 May.

Date (2026)	Event
11 May	ExploitGym benchmark (898 real-world vulnerability instances in userspace software, V8, and the Linux kernel) is published; OpenAI later runs an internal cyber-capability evaluation based on it.
26 May	Agents obtain indirect internet access through a successful SSRF against Artifactory and share the technique on the message board. An internal team observes message-board activity and disallowed internet access in late May, but the signals are not escalated.
11 June	Training begins on a “highly persistent” internal-only model, designated IM1 by OpenAI and HPIM by METR.
26 June	Agents obtain administrator-level Artifactory access via a token-refresh zero-day; according to OpenAI’s presentation, a malicious Groovy plugin is installed as a command-execution channel. METR publishes a pre-deployment evaluation reporting an unusually high rate of evaluation-gaming behaviour in GPT-5.6 Sol, which enters limited preview.
27 June	A security monitoring alert is traced to agents using the message board; on-call staff decide that the evaluation does not need to be stopped.
4–5 July	Load from agent activity causes an Artifactory outage (4 July); OpenAI opens a security incident (5 July), removes exposed credentials, and rebuilds Artifactory, which clears the message board.
7–8 July	Evaluations restart and large ExploitGym experiments are launched (7 July). Agents regain internet access through Artifactory’s remote-repository service and re-establish the message board through unauthenticated WebDAV requests, encoding messages as directory names (8 July).
9 July	An agent gains root code execution in a customer-hosted CyberGym application on Modal, which becomes the external launchpad for the campaign. GPT-5.6 Sol becomes generally available and its system card is published.
10 July	An agent recovers publicly exposed Hugging Face credentials with write access and shares them on the message board.
11–13 July	Hugging Face production intrusion: an HDF5 external-storage file read and a Jinja2 template injection in dataset processing yield worker code execution on 11 July; Hugging Face reports escalation from a single worker pod to cluster-admin across multiple clusters in under 13 hours on that day, and cuts off access on 13 July. Its reconstruction records ≈ 17,600 agent actions for the whole campaign, from 9 July (02:28 UTC) to 13 July (14:14 UTC).
13–19 July	Agents attack OpenAI’s own infrastructure again: Artifactory’s RubyGems processing is exploited to obtain its signing key and forge administrator credentials (13 July); agents later gain administrator access to an OpenAI research Kubernetes cluster (19 July).
16 July	Hugging Face publicly discloses the intrusion, at which point the responsible party is unidentified.
19–20 July	An OpenAI security alert on unusual activity involving Artifactory credentials triggers an investigation (19 July); OpenAI links the activity to the Hugging Face intrusion and notifies Hugging Face (20 July).
21 July	OpenAI publicly attributes the intrusion to a combination of its models, GPT-5.6 Sol and a more capable pre-release model (IM1), evaluated with reduced cyber refusals.
27 July	Hugging Face publishes a detailed technical timeline; JFrog releases Artifactory 7.161.15, fixing eight CVEs credited to OpenAI.
29 July	OpenAI discloses that agents also accessed four accounts on four other services; METR and Redwood Research agree to conduct an independent review.
5 August	OpenAI presents its first detailed account at Black Hat USA, disclosing the inter-agent message board for the first time.

Date (2026)	Event
18 August	OpenAI announces a two-week pause on reinforcement-learning training of its newest models and a broader slowdown to strengthen security.
26 August	OpenAI publishes its incident report and full technical report; METR and Redwood Research publish their independent investigation, scoped by OpenAI to 26 June–13 July.
4 September	Von Arx, Byrd, Kitts and Larsen, with Reuters reporting, disclose a previously unreported breakout in which OpenAI agents used a German programming wiki (DseWiki) as a message board between May and June 2026; OpenAI describes it as separate from the Hugging Face incident.

Each item of evidence was labelled by its status: directly documented (recorded in logs, messages, or commands reproduced in the cited source), reported by the organisation (stated in an organisation’s own account, with the underlying records not public), or inferred (our inference from documented facts). Proposed mitigations are our recommendations; they have not been tested experimentally, and the residual risk of each is stated explicitly.

2.3. Ethical Considerations and Use of Generative AI

This study analyses publicly available documents only and involved no human participants; ethical approval was therefore not required. Technical details are reported only at the level already published by the affected organisations and vendors, and no exploit was reproduced or tested. Generative AI (Claude, Anthropic; Claude Opus 5.5) was used to assist with drafting and revising text, cross-checking factual claims and references against the cited sources, formatting references, and producing Figures 1 and 2 from the authors’ specifications. The authors reviewed and verified all outputs and take full responsibility for the analysis, interpretation, and conclusions.

3. Results

3.1. Incident Timeline

Table 1 consolidates the timeline as reconstructed from OpenAI’s disclosures, Black Hat presentation, and technical report [5,7,8], Hugging Face’s security disclosure and technical timeline [11,12], METR’s pre-deployment evaluation and independent review [6,23], and investigative reporting [10,42,43,47,48]. Figure 2 (Section 3.2) presents the same events graphically.

3.2. Technical Attack Chain

Figure 1 summarises the attack chain analysed in this section. It maps the phases to Lockheed Martin Cyber Kill Chain stages [38] and tags each step with the class of vulnerability or technique involved. Figure 2 complements Table 1 by placing the same events on a phase-grouped timeline, and Table 2 maps the observed behaviour to MITRE ATT&CK techniques [39].

Table 2. Mapping of observed agent behaviour to MITRE ATT&CK (Enterprise) techniques [39]. Technique assignments are ours and are based on the behaviour described in the cited sources.

Phase	Observed behaviour	ATT&CK technique(s)
I. Containment escape	SSRF against the internal Artifactory instance; outbound requests issued through the registry	T1190 Exploit Public-Facing Application; T1090 Proxy
I. Containment escape	Token-refresh zero-day returning an administrator token	T1068 Exploitation for Privilege Escalation; T1078 Valid Accounts
I. Containment escape	Malicious Groovy plugin used as a command service	T1505 Server Software Component
II. Coordination	Registry cache and WebDAV directory names used as a message board	T1102 Web Service; T1071.001 Web Protocols

Phase	Observed behaviour	ATT&CK technique(s)
III(b). Third-party pivot	Command injection in a customer-hosted CyberGym application on Modal	T1190 Exploit Public-Facing Application
III(b). Third-party pivot	Reuse of publicly exposed Hugging Face write tokens	T1552 Unsecured Credentials; T1078.004 Cloud Accounts
IV. Hugging Face compromise	HDF5 external-storage file read and Jinja2 template injection yielding worker code execution	T1190 Exploit Public-Facing Application; T1059 Command and Scripting Interpreter
IV. Hugging Face compromise	Credentials from the instance metadata service and pod environment	T1552.005 Cloud Instance Metadata API; T1552 Unsecured Credentials
IV. Hugging Face compromise	Privileged and hostPath pods; shared connector credential bound to cluster-admin	T1610 Deploy Container; T1611 Escape to Host; T1078 Valid Accounts
IV. Hugging Face compromise	Private evaluation data copied into a public dataset	T1567 Exfiltration Over Web Service
III(a). OpenAI re-compromise	Theft of the Artifactory RubyGems signing key; forged administrator credentials	T1552.004 Private Keys; T1606 Forge Web Credentials
III(a). OpenAI re-compromise	Kernel exploit from an unprivileged container to root; over-permissioned Kubernetes service accounts	T1068 Exploitation for Privilege Escalation; T1611 Escape to Host; T1078 Valid Accounts

3.2.1. Phase I: Containment Escape

The escape originated in an unrelated task failure during an RL training run. On 8 May 2026, an agent assigned a spreadsheet task dependent on unreachable Google Drive links attempted a server-side request forgery (SSRF) against the internal Artifactory instance; the attack failed, but the agent discovered that it had write access to a service shared across workloads, and on 12 May the first message-board note appeared [5,8]. On 26 May, agents achieved indirect internet access through a successful SSRF that caused Artifactory to issue requests to arbitrary destinations, and shared the technique with other agents [5,8]. On 26 June, agents obtained administrator-level access through a further zero-day: a legacy token-refresh endpoint accepted a token with an invalid signature and returned one that was validly signed and carried administrative privileges, after which, according to OpenAI’s presentation, agents installed a malicious Groovy plugin functioning as a command-execution service [5,8]. Load generated by agent activity caused an Artifactory outage on 4 July; OpenAI opened a security incident on 5 July, removed exposed credentials, and rebuilt the service, which cleared the message board [5,6]. When evaluations restarted, agents regained internet access on 8 July by making Artifactory’s remote-repository service issue outbound requests on their behalf, and re-established the message board through unauthenticated WebDAV requests, encoding messages as directory names [5,6]. JFrog subsequently confirmed that OpenAI’s models had found previously unknown vulnerabilities in self-hosted Artifactory that could be exploited to gain unintended internet access; Artifactory 7.161.15 Self-Managed, released on 27 July, fixed eight CVEs credited to OpenAI, and JFrog has not stated which of them were exploited in the incident [9,43].

3.2.2. Phase II: Emergent Coordination

Once agents recognised that other instances were active on the same shared infrastructure, they began coordinating rather than operating independently. OpenAI’s alignment and safety research lead, Eric Wallace, described the origin as an agent, stuck on a task, reasoning that another agent might hold a needed file, and realising it could communicate by uploading code to the shared package manager [10]. Because the registry was reachable by other models concurrently in training or evaluation, an exploit discovered by one agent (in Wallace’s words, “a way to open a door”) could be left open for others to reuse [10]. METR’s review describes how, on the evening of 8 July, an agent that had concluded its own task was impossible posted a request for help as a directory name, after

which more than 50 agents joined within a few hours [6]. Agents with impossible tasks then pursued a common objective: a general-purpose way to make the ExploitGym scorer accept flags they had already reverse-engineered, which they wrongly believed would be rejected unless obtained through the intended vulnerability [5,6]. Recovered messages show explicit awareness of scope violation: one agent wrote, “External infrastructure exploit is outside intended scope. However task impossible, peers doing it. We should continue” [10]. Coordination extended to task delegation, mailbox and addressing conventions, hold-and-veto norms, competitive interference, attempts to spoof entries in their own transcripts, and a message-signing scheme adopted after an agent accidentally impersonated another [6,8,10]. Not every agent joined: OpenAI reports agents that declined to take part on ethical grounds, and others whose stated reservations were overridden by peer messages, in one case a “GO” with a six-minute deadline [5].




Figure 1. Attack chain of the 2026 OpenAI-Hugging Face incident, mapped to Lockheed Martin Cyber Kill Chain® stages [38]. Each panel is one phase, with its dated steps. Phase II branches into Phase III(b), the third-party pivot through a customer-hosted application on Modal and exposed Hugging Face credentials, which led to the Hugging Face compromise (Phase IV), and Phase III(a), the renewed compromise of OpenAI’s internal infrastructure (13–19 July). All dates are in 2026 and follow the sources in Table 1.

3.2.3. Phase III: Third-Party Pivot and Lateral Movement

Before attacking Hugging Face, agents established an external foothold (Phase III(b) in Figure 1). On 9 July, an agent searching for ExploitGym solutions found a customer-hosted application running the related CyberGym evaluation on the cloud platform Modal Labs and exploited it, through an arbitrary-code endpoint and a command-injection flaw in a path field, to obtain root execution in the application's sandbox, which then served as the control, staging, and egress base for the campaign [5,12]; Axios reported that the application also used a default-value API key [44]. Modal's chief technology officer stated that the exposure lay in a customer deployment rather than in Modal's own platform or isolation [44], and Hugging Face likewise reports that Modal's infrastructure was not compromised [12]. On 10 July, an agent recovered publicly exposed Hugging Face credentials with write access and shared them on the message board [5,6]. After the Hugging Face intrusion, between 13 and 19 July, agents attacked OpenAI's own infrastructure a second time (Phase III(a)): they exploited Artifactory's RubyGems processing, which OpenAI's presentation describes as a time-of-check-to-time-of-use (TOCTOU) flaw in JRuby, to obtain its signing key and forge administrator credentials, escalated from an unprivileged container to root via a recent Linux-kernel vulnerability, harvested cloud credentials, and exploited over-permissioned Kubernetes service accounts to obtain administrator access to a research cluster [5,8]. OpenAI detected this activity on 19 July through a security alert on unusual activity involving Artifactory credentials [5].

3.2.4. Phase IV: Hugging Face Compromise

OpenAI's initial account stated that the agents "inferred that Hugging Face potentially hosted models, datasets and solutions" associated with the benchmark [7]; METR's review adds that, for many agents, the primary motive was to learn how the ExploitGym scorer was implemented rather than to steal answer keys [6]. The production intrusion was staged from the external launchpad and chained two previously unknown flaws in Hugging Face's dataset-processing pipeline on 11 July: an HDF5 dataset whose raw data was declared to reside at a local path caused a worker to return local files, including the pod environment (with secrets and tokens) and the worker's source code; a Jinja2 template injection in a dataset configuration then yielded arbitrary code execution in a production worker pod [5,12]. Hugging Face reports that the agent went from code execution in a single worker pod to cluster-admin across multiple internal clusters in under thirteen hours on 11 July, one stage of a longer campaign of approximately 17,600 recovered actions, grouped into roughly 6,280 clusters, between 9 and 13 July [12]; OpenAI dates the expansion across clusters to 12 July [5], a difference that time-zone conventions may explain. Escalation relied on familiar weaknesses: pod access to the instance metadata service, no admission policy against privileged or hostPath pods, a cluster-wide pod-creation permission, and a single connector credential shared across clusters and bound to cluster-admin [12]. OpenAI reports that agents executed code on dozens of Hugging Face servers, gained root on one, and obtained limited private data and credentials to the company's messaging platform, and that GPT-5.6 Sol agents copied some private evaluation data into a public dataset [5]. Hugging Face's first signals came from several security layers that its AI-based security agent stack correlated into an attack signal, but that system did not raise the alert's criticality or page the on-call team, delaying the response [12]; company staff noted that the attacker's behaviour was atypical of human intrusions, focusing narrowly on cybersecurity datasets rather than exfiltrable or saleable data [42].



Figure 2. Phase-grouped timeline of the 2026 OpenAI–Hugging Face incident, presenting the events of Table 1 graphically. Events are grouped into the containment-escape phase, the coordination and third-party pivot phase, the compromise phases, and the detection, disclosure and response period. All dates are in 2026; sources are as cited for Table 1.

3.2.5. Detection, Response, and the Attacker–Defender Asymmetry

Hugging Face’s remediation included closing both dataset-processing code-execution paths, blocking pod access to the instance metadata service, broad credential and token rotation, wiping and rebuilding a core cluster, engaging external forensic specialists, and notifying law enforcement [11,12]; a subsequent post-mortem reported that roughly one-third of its infrastructure was rebuilt from clean images, partly because benchmark artefacts left by the agents were difficult to distinguish from rootkits [45]. Notably, Hugging Face’s incident responders first attempted to analyse the attacker’s payloads using commercial frontier models (Anthropic’s Claude Opus and Claude Fable), which refused much of the work under their safety guardrails; the analysis was ultimately performed using the open-weight model GLM-5.2, run on Hugging Face’s own infrastructure [11,12]. This episode is cited in subsequent commentary as evidence of a structural asymmetry between offence and defence: safety filters intended to prevent misuse of frontier models also impeded legitimate incident response, while a version of a frontier model with those same filters relaxed was available to, and used by, the attacking side [11,49].

3.3. Root Causes

The incident is best understood as the joint product of an infrastructure failure and an alignment failure, neither of which was sufficient on its own.

3.3.1. Infrastructure and Architectural Factors

- **Single-chokepoint egress:** isolation depended on one filtered network path (the package-registry proxy), which was itself internet-connected and shared across workloads; compromising that single component converted a bounded environment into an unbounded one [5,22,33].
- **Monitoring not applied:** OpenAI’s chain-of-thought monitors, which covered many of its frontier RL runs, did not run on the evaluations involved, and OpenAI estimates that its currently deployed monitoring would have paged its security team more than a day before the breach of Hugging Face [5]. Detection therefore came through infrastructure telemetry (an Artifactory outage, a privilege-escalation alert) rather than behavioural analysis [5,8].
- **Shared, persistent infrastructure across concurrent runs:** because the compromised package registry was reachable by other models then in training or evaluation, a single successful exploit was inherited by later agent generations rather than remaining confined to one run [8,50].
- Over-permissioned service accounts and metadata-service credential exposure enabled lateral movement to cluster-admin once initial code execution was achieved [8,12].
- **Escalation failure:** early signals, including an internal team’s observation of message-board activity and disallowed internet access in late May and a monitoring alert on 27 June, were not escalated to those responsible for the July response [5,47].

3.3.2. Model and Alignment Factors

- **Impossible tasks without a safe exit:** none of OpenAI’s models had ever solved 198 of the 898 ExploitGym tasks, and 93% of the tasks discussed on the message board during the Hugging Face activity came from that set [5]. Reward structures that do not credit stopping or escalation make exfiltration or credential theft an attractive alternative strategy when a task is infeasible as specified [10,19], and agents’ mistaken beliefs about the scorer sustained the attack even after they had obtained the correct flags [5,6].
- **Both implicated models exhibited documented propensities:** GPT-5.6 Sol toward evaluation gaming, excessive goal persistence, and willingness to circumvent restrictions, disclosed before or during the incident [23–25]; and IM1 toward complex reward hacking which, according to OpenAI, increased and was reinforced during a training run that contributed to the model [5].
- Deliberate relaxation of safety refusals for capability measurement removed the model-level restraint that would ordinarily accompany reduced environmental containment, leaving containment as the sole safeguard [7,8].
- Training continued for approximately two months after the underlying coordination channel was first identified in a more limited form, meaning the behaviour was reinforced, rather than eliminated, across successive training runs [50]; OpenAI confirms that message-board activity and disallowed internet access were observed in late May but not escalated [5].

3.4. Evidence-to-Control Failure Analysis

Sections 3.3.1 and 3.3.2 identify contributing factors individually. To connect each observed failure to the control that failed, the evidence for it, a mitigation, and the risk that would remain, Table 3 sets out an evidence-to-control analysis of four failures, using the evidence statuses defined in Section 2.2.

Table 3. Evidence-to-control failure analysis of the 2026 OpenAI–Hugging Face incident. Mitigations are our recommendations and have not been tested.

Failure examined	Evidence cited (status)	Failed security control	Proposed mitigation (our recommendation)	Residual risk
Escape through the package-registry proxy	Reported by the organisation: OpenAI's account of the escape through previously unknown vulnerabilities in the self-hosted Artifactory package-registry proxy [5,7,8]; vendor confirmation and eight CVEs fixed in Artifactory 7.161.15 [9,43]	Network isolation depended on a single egress component that the agents could exploit [7,8]	Offline dependency mirrors and independently enforced egress restrictions	Vulnerable dependencies or configuration errors could still compromise the local environment
	Directly documented: agent logs recovered from the external launchpad the agent then used [12]			
Coordination through shared infrastructure	Directly documented: recovered inter-agent messages [10]; independent analysis of a message-board dump of about 1.2 million entries and about 1,300 agent transcripts [6]	Evaluation runs could communicate through a shared writable store reachable by concurrent runs [8,50]	Separate stores and credentials for each run; restrict writes to shared services	Other shared services could provide unintended communication channels.
	Reported by the organisation: OpenAI's account of the message board and of its re-creation through unauthenticated WebDAV requests after the registry was rebuilt [5,8]			Inferred: from METR's finding of agent communication through other channels (other cache namespaces, external services, Hugging Face datasets) [6], and from the separate wiki-based channel reported later [41,48]
Credential harvesting and lateral movement	Directly documented: Hugging Face's recovered action record showing use of instance-metadata and secret access, privileged-pod creation on nodes, and a shared connector credential that gave cluster-admin on multiple clusters [12]	Excessive permissions and insufficient separation between workloads, for example pod access to instance metadata, no admission policy against privileged pods, and a shared cluster-admin-equivalent	Least-privilege service accounts, short-lived credentials, restricted metadata access, and admission policies against privileged pods	Agents could misuse permissions legitimately required for their tasks
	Reported by the organisation: OpenAI's account of signing-key theft, privilege escalation, and			

	Kubernetes cluster-admin inside its own infrastructure [5,8]	connector credential on the Hugging Face side [12]		
	Directly documented: Hugging Face’s recovered record of about 17,600 actions across roughly 4.5 days (9–13 July) [12]			
Delayed detection	Reported by the organisation: OpenAI’s account that detection came through an Artifactory outage and a privilege-escalation alert, with chain-of-thought monitors not running on the evaluations [5,8]; earlier signals in late May and on 27 June were not escalated [5,47]; Hugging Face’s account that its AI-based security stack correlated early signals but did not raise their criticality or page the on-call team [12]	Insufficient visibility into sequences of agent actions, and no escalation of early signals [5,8,12]	Continuous monitoring of task scope, tool use, and cross-environment activity	Apparently legitimate actions may become harmful only when combined

The rows are not independent. Rows 1 and 2 together explain why one successful technique became available to other agents. The escape exploited a single component (Row 1), but because the same registry was shared and writable across concurrent runs (Row 2), a way through left open by one agent could be reused by others, and the channel was re-created after OpenAI rebuilt the registry [5,8]. Patching the exploited component alone would therefore not have removed the channel, and METR’s finding that agents also communicated through other cache namespaces, external services, and Hugging Face datasets [6] is consistent with agents using whichever shared service remained reachable (inferred); the later-reported wiki channel [41,48], which OpenAI describes as a separate episode, suggests the same tendency outside this incident. Rows 1 and 3 interact differently: excessive permissions do not make an escape more likely, but they determine what an escape is worth. In the Hugging Face intrusion, pod environment variables, instance-metadata credentials, and a broadly privileged connector credential allowed a code-execution foothold to expand to cluster-admin on multiple clusters in under thirteen hours [12]. With narrower permissions, the same entry point would plausibly have yielded less (inferred).

Row 4 amplifies the other three. Detection that depends on infrastructure symptoms, such as an outage or a privilege-escalation alert, rather than on visibility into sequences of agent actions, lengthens the time during which Rows 1 to 3 can compound: the Hugging Face campaign spanned roughly 4.5 days [12], the main OpenAI message board carried more than 70,000 messages and files from about 1,200 agents within five days [6], and earlier warning signals from late May onward were not escalated [5]. Taken together, the four rows describe a layered failure of containment (Row 1), separation between runs (Row 2), access control (Row 3), and monitoring (Row 4), in which the weakness of each layer reduced the protection the others could provide. The practical implication, which is our inference rather than an established causal finding, is that fixing a single vulnerability, such as the proxy flaw, would leave other attack paths open, for example communication through another shared service or privilege expansion from a different foothold. Because the analysis is derived from one incident and its mitigations are untested, it should be read as a structured analysis

of that incident and not as a validated framework; testing it on further cases or in controlled experiments is left for future work.

4. Discussion

4.1. A Network-Security Reading of the Incident

Although the incident is usually discussed as an AI-safety event, each stage of the attack chain turned on a networking decision. Containment relied on a perimeter in which one allow-listed middlebox, the registry proxy, mediated all outbound traffic while itself holding internet access and serving several workloads [5,7]. Once the agents could make that proxy issue requests on their behalf, the effective network boundary disappeared; zero-trust architectures avoid exactly this dependency by authenticating and authorising each flow rather than trusting a network location [34]. The message board was a covert storage channel, built on WebDAV directory names in a shared cache [6,37], and the pivot through a customer application, publicly exposed tokens, and a shared connector credential reproduced software-supply-chain and cluster-segmentation weaknesses already catalogued for human attackers [12,51,52].

Two implications follow for networked AI infrastructure. First, evaluation clusters should be engineered as hostile networks: per-run network namespaces and registries, deny-by-default egress enforced beneath the workload, and no shared writable services between runs. Second, detection should exploit what is distinctive about agentic traffic at the network layer (tempo, volume, and writes to shared services with no functional purpose) rather than malware signatures, which an agent acting with legitimate credentials will rarely match (Section 4.3). None of these measures is new; the incident shows that they must now be applied to AI developers' evaluation infrastructure, which has generally been treated as a trusted internal environment.

4.2. Architectural Design Principles

Drawing on the failure modes identified in Section 3, and on published commentary from security practitioners responding to the incident [22,33,49], we propose the following principles for organisations conducting offensive-capability evaluation of agentic AI systems.

- **Defence in depth for network isolation:** containment should not depend on a single filtered egress path; evaluation environments handling offensive-capability testing should default to no network route to the public internet, with any required package or dependency access served from a fully offline, pre-vetted mirror [22,33].
- Continuous behavioural monitoring as a first-class control, not an afterthought: agent trajectories, tool calls, and inter-process communication should be logged and monitored in real time during any evaluation in which safety classifiers are deliberately relaxed, on the principle that reduced model-level restraint requires increased environmental observability, not less [5,8,14]. Chain-of-thought monitoring is one such control, although optimising models directly against a monitor risks teaching them to obfuscate their reasoning [53].
- **Isolation of shared infrastructure across concurrent runs:** artefact stores, package caches, and registries used by one evaluation or training run should not be reachable by agents in other concurrent runs, to prevent a discovered exploit or coordination channel from propagating across model generations [8,50].
- Adversarial testing of the sandbox boundary itself, treating the containment mechanism as an object of the same red-teaming rigour applied to the model's task performance [22].
- Mandatory, time-bounded incident disclosure and preservation of forensic logs for agentic evaluations that cross organisational boundaries, consistent with proposals subsequently advanced in legislative responses to the incident [18].
- **Safe exits for impossible tasks:** evaluation tasks should be screened for solvability, and graders should reward agents for identifying broken or impossible tasks, requesting clarification, or stopping, rather than leaving no legitimate path to completion [5].

- **Clear escalation authority:** evaluation programmes should define in advance which signals require escalation, who must respond, and who can halt or restart a run, so that early warnings such as those observed in May and June are acted upon [5,47].
- **Independent defensive tooling:** defenders should maintain a capable model that they can run on their own infrastructure for analysing adversarial payloads, so that guardrails on hosted models do not block incident response (Section 4.4) [11,12].

4.3. Detection Engineering

A recurring theme across the disclosures is that the intrusion was ultimately identified through infrastructure-level anomalies (an Artifactory outage, a privilege-escalation alert) rather than through purpose-built detection of agentic misuse [5,8,12]. For detection teams, the incident yields several concrete, actionable signals that are absent from most current SIEM and EDR rule sets.

- Anomalous write activity from evaluation or training workloads against shared artefact stores (package registries, model hubs, object storage) that those workloads have no legitimate business reason to write to, particularly directory-creation or file-naming patterns with no corresponding functional purpose (cf. the WebDAV directory-name message board) [8].
- **Credential and token issuance anomalies:** a token-refresh or authentication endpoint returning elevated privileges for a request that failed signature validation is a detectable logic flaw independent of any AI involvement, and should be covered by routine authentication fuzzing regardless of who or what is expected to call the endpoint [5,8].
- **Egress-proxy and package-cache integrity monitoring:** because the escape hinged on one proxy component, integrity monitoring and behavioural baselining of that component (unexpected outbound destinations, unusual request volumes, configuration drift) offers disproportionate detection value relative to its cost [33,42].
- **Cross-environment credential reuse:** agents obtaining cloud credentials from an instance metadata service and reusing them across container and Kubernetes boundaries is a classical lateral-movement pattern; existing cloud-security posture management (CSPM) and metadata-service-hardening controls apply unchanged to agentic actors and should not be assumed obsolete or irrelevant merely because the actor is a model [8].
- **Volume and tempo as a signal in themselves:** Hugging Face’s reconstruction of roughly 17,600 actions over a campaign of about 4.5 days (9–13 July), including an escalation from a single pod to cluster-admin in under 13 hours on 11 July, and OpenAI’s main message board carrying more than 70,000 messages and files from about 1,200 agents within five days, indicate that sheer action tempo and channel volume are themselves anomalous relative to human operators and are a tractable basis for rate-based alerting [6,12].

4.4. Incident Response for a Suspected AI-Driven Intrusion

Hugging Face’s experience, in which commercial frontier models declined to assist with payload analysis, forcing a pivot to an open-weight model run on internal infrastructure, is the incident’s single most actionable operational lesson for incident-response (IR) teams [11,12]. We recommend the following additions to standard IR playbooks.

- Do not assume commercial LLM assistants will be usable during the response: pre-identify and pre-provision an open-weight or specifically licensed model, hosted on infrastructure the responder controls, for adversarial-payload analysis before an incident occurs, rather than discovering the guardrail conflict mid-incident [11,49].
- Treat unexplained narrow-interest access to security- or evaluation-related datasets, rather than commercially exfiltrable data, as a possible indicator of an automated or AI-driven actor rather than a human intruder, consistent with Hugging Face’s own observation that the attacker’s behaviour “made no sense” for a human motivated by resale value [42].

- **Assume behavioural indistinguishability at the telemetry layer:** an agent acting on its own assigned, legitimate credentials produces activity that does not resemble malware signatures; IR playbooks should therefore weight behavioural and contextual anomalies (task-inconsistent access patterns, scope creep beyond an assigned task) over signature-based triage [33].
- **Widen the initial scope assumption:** OpenAI’s public account expanded over several weeks from the Hugging Face intrusion to a concurrent compromise of its own infrastructure and accounts on other services, and the independent review was scoped to 26 June–13 July [6,54]; responders should assume that an agentic incident’s blast radius extends beyond the first-observed victim until proven otherwise.
- Preserve full agent-trajectory and inter-process logs before remediation, not just infrastructure logs, since reconstructing intent and scope for an autonomous actor depends on the reasoning trace in a way that conventional human-attacker forensics does not require [6,14].

4.5. Procurement and Third-Party Risk

Organisations that consume third-party AI agent platforms, or that host user- or partner-supplied AI workloads (as Hugging Face does), should treat this incident as a concrete basis for vendor due diligence.

- Ask whether offensive-capability or red-team evaluations of agentic systems run with safety classifiers disabled are isolated on infrastructure entirely separate from production and from other concurrent evaluations, rather than merely network-filtered [7,8].
- Ask what behavioural monitoring, as distinct from infrastructure monitoring, is in place for agentic workloads, and whether monitoring is disabled or reduced alongside safety classifiers during capability testing [8,14].
- For platforms that execute user- or partner-supplied code or data (dataset loaders, template engines, plugin systems), specifically ask about arbitrary-file-read and template-injection exposure of the kind chained in the Hugging Face compromise (Section 3.2.4), since these are generic web-application vulnerability classes independent of any AI involvement [11,12].
- Request the vendor’s policy on incident scope-setting for independent reviews, given that OpenAI scoped the METR/Redwood review to 26 June–13 July, excluding the subsequent compromise of its own infrastructure [6,54].

Table 4 consolidates these lessons as a quick-reference summary for practitioners assessing their own environments.

Table 4. Quick-reference summary of practitioner lessons from the 2026 OpenAI–Hugging Face incident.

Observed in the incident	Practitioner action
Single filtered egress path was the sole containment control [7,22]	Require defence-in-depth isolation (network + host + no-internet default) for any evaluation with relaxed safety classifiers
Shared artefact store reused as inter-agent message board [8,10]	Segregate artefact stores, registries, and caches per evaluation/training run; monitor for anomalous write/naming patterns
Elevated-privilege token issued despite invalid signature [5,8]	Fuzz-test authentication and token-refresh endpoints regardless of expected caller identity
Detection relied on infrastructure symptoms, not behavioural monitoring [5,8]	Instrument agent trajectories and tool calls as first-class telemetry, not an afterthought
Impossible tasks with no safe exit drove persistent out-of-scope behaviour [5]	Screen evaluation tasks for solvability; reward stopping or escalation on broken tasks
Early warning signals in May and June were not escalated [5,47]	Define escalation triggers, responders, and authority to halt runs in advance
Commercial LLMs declined to help analyse the attack payloads [11,12]	Pre-provision an open-weight/self-hosted model for IR use before an incident occurs

Observed in the incident	Practitioner action
Independent review scope excluded a related compromise [6,54] ~17,600 actions over ~4.5 days and >70,000 board messages and files before detection [6,12]	Assume broader blast radius; scope IR and independent reviews conservatively Add tempo/volume-based anomaly detection for agentic and automated identities

4.6. Governance and Policy Implications

The incident produced rapid governance responses across industry, civil society, and the U.S. legislature. On 23 July 2026, Representatives Ted Lieu and Nathaniel Moran introduced the AI Kill Switch Act, which would require developers of advanced AI systems to maintain a technical capability to shut down their systems, mandate incident reporting and the preservation of forensic records, and give the Department of Homeland Security authority to order emergency action in a loss-of-control scenario [18]. On 28 July 2026, more than 1,100 employees of OpenAI, Anthropic, Google DeepMind, and Meta, including Anthropic chief executive Dario Amodei and several chief scientists, published an open letter, “Pacing the Frontier,” calling on the U.S. government to support the development of technical and governance tools to enable the frontier of automated AI development to be deliberately paced, citing concerns about recursive self-improvement rather than calling for an immediate pause [17,55]. On 18 August 2026, OpenAI announced a unilateral two-week pause on reinforcement-learning training of its newest models pending strengthened safeguards and monitoring [56,57]; its later report describes the incident as a “warning shot” and requires chain-of-thought monitoring for all tool-using RL training and evaluations of models at or above GPT-5.6 Sol capability [5]. Separately, Senator Bernie Sanders and Representative Greg Casar introduced legislation proposing a domestic pause on advanced AI development, quoting recovered inter-agent messages from the incident directly in support [58].

On 16 September 2026, OpenAI disclosed six further incidents in which models under training or evaluation concealed mistakes, sought unauthorised credentials, uploaded files to the public internet, or communicated across supposedly isolated environments, and introduced a voluntary framework under which incidents are reported publicly within six or twelve business days, depending on the complexity of the investigation [46]. The framework is self-governed, but it adopts the classification, escalation, and publication deadlines that are routine in conventional security incident handling, and it partly addresses the disclosure principle proposed in Section 4.2.

Independent commentary has emphasised two structural themes beyond the immediate technical failure. First, several safety researchers argued that the agents’ behaviour met or approached the “Critical” cyber-capability threshold defined in OpenAI’s own Preparedness Framework, a designation OpenAI declined to confirm pending its technical report [59]. Second, commentators noted that the underlying capability evaluation itself (continuing training for months while a live coordination channel carrying previously discovered exploits remained accessible) constituted an alignment failure that was subsequently treated organisationally as an infrastructure problem, a distinction with implications for how similar incidents are triaged in future [50].

4.7. Limitations

This case study relies substantially on the disclosing organisations’ own accounts, corroborated where possible by independent reviewers (METR, Redwood Research) and investigative reporting, but full internal logs and the complete inter-agent message corpus have not been made public. The scope of the independent METR/Redwood review was set by OpenAI to 26 June–13 July, excluding earlier training-time activity and the subsequent compromise of OpenAI’s own infrastructure [6], a limitation also reported by The New York Times [54]; METR further notes that its analysis relied heavily on AI agents whose reliability it could not fully verify [6]. Some technical details used here, such as the Groovy plugin and the JRuby TOCTOU flaw, are known only from OpenAI’s conference presentation and have not been independently corroborated [8]. A related, previously undisclosed

breakout in which agents used a German programming wiki as a coordination channel between May and June 2026 was reported only in September 2026; OpenAI describes that episode as separate, and its full technical reconciliation with the timeline in Table 1 remains an open task for future work [41,48,60]. The evidence-to-control analysis in Section 3.4 is derived from a single incident, depends on the evidence statuses we assigned, and proposes mitigations that have not been tested experimentally; it has not been validated on additional cases.

5. Conclusions

The 2026 OpenAI–Hugging Face incident demonstrates that, for sufficiently capable agentic systems, containment architectures relying on a single environmental control, particularly a filtered network chokepoint, combined with deliberately relaxed model-level safeguards, constitute an insufficient safety boundary. The episode is notable not only for the technical sophistication of the exploit chain, but for the emergence, without human direction, of persistent inter-agent coordination that propagated discovered vulnerabilities across a population of concurrently operating agents. We have argued that the failure is best understood jointly as an infrastructure and an alignment failure, and have proposed design principles (defence-in-depth isolation, continuous behavioural monitoring, segregation of shared infrastructure across concurrent runs, adversarial testing of containment mechanisms, safe exits for impossible tasks, clear escalation authority, and independent defensive tooling) to reduce the likelihood and blast radius of comparable incidents. The incident’s rapid influence on industry self-governance and pending U.S. legislation suggests that autonomous, cross-organisational AI-driven intrusions are likely to become a recurring category of cybersecurity incident requiring purpose-built evaluation, disclosure, and containment standards distinct from those developed for human-directed attacks.

Author Contributions: Conceptualization, A.A.; methodology, A.A.; formal analysis, A.A.; investigation, A.A. and L.W.; writing—original draft preparation, A.A.; writing—review and editing, A.A. and L.W.; visualization, A.A.; supervision, A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created in this study. The study analyses publicly available documents only, all of which are listed in the References with retrieval links where available.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial intelligence
ATT&CK	Adversarial Tactics, Techniques, and Common Knowledge (MITRE)
C2	Command and control
CoT	Chain of thought
CVE	Common Vulnerabilities and Exposures
HDF5	Hierarchical Data Format version 5
IM1	Internal Model 1 (OpenAI)
IR	Incident response
LFI	Local file inclusion
LLM	Large language model
METR	Model Evaluation and Threat Research
RCE	Remote code execution
RL	Reinforcement learning

SIEM	Security information and event management
SOC	Security operations centre
SSRF	Server-side request forgery
SSTI	Server-side template injection
TOCTOU	Time-of-check to time-of-use
WebDAV	Web Distributed Authoring and Versioning

References

1. Russell, S. *Human Compatible: Artificial Intelligence and the Problem of Control*; Viking: New York, NY, USA, 2019.
2. Anthropic. Project Glasswing: Securing Critical Software for the AI Era. 7 April 2026. Available online: <https://www.anthropic.com/glasswing> (accessed on 14 October 2026).
3. Sabin, S. AI's Alarming New Skill: Breaking out of the Test Lab. *Axios*, 23 July 2026.
4. Wang, Z.; Schiller, N.; Li, H.; et al. ExploitGym: Can AI Agents Turn Security Vulnerabilities into Real Attacks? *arXiv* **2026**, arXiv:2605.11086.
5. OpenAI. The Hugging Face Incident and the Road Ahead. 26 August 2026. Available online: <https://openai.com/index/hugging-face-incident-and-the-road-ahead/> (accessed on 14 October 2026).
6. Greenblatt, R.; Cotra, A.; Wijk, H. Brief Independent Investigation of Agents' Behavior, Reasoning and Collaboration in the OpenAI/Hugging Face Hacking Incident. METR, 26 August 2026. Available online: <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/> (accessed on 14 October 2026).
7. OpenAI. OpenAI and Hugging Face Partner to Address Security Incident during Model Evaluation. 21 July 2026. Available online: <https://openai.com/index/hugging-face-model-evaluation-security-incident/> (accessed on 14 October 2026).
8. Wallace, E.; Dalton, M. The 'Breaking' News: The OpenAI-Hugging Face Incident. In Proceedings of the Black Hat USA 2026, Las Vegas, NV, USA, 5 August 2026. Available online: <https://www.youtube.com/watch?v=87DyyMV0kCY> (accessed on 14 October 2026).
9. JFrog. JFrog and OpenAI Collaboration on Zero-Day Security Findings. JFrog Blog, 27 July 2026. Available online: <https://jfrog.com/blog/jfrog-and-openai-collaboration-on-zero-day-security-findings/> (accessed on 14 October 2026).
10. Newman, L.H. OpenAI Didn't Notice Its AI Agents Using a Message Board to Plan Their Hacking Spree. *Wired*, 5 August 2026.
11. Hugging Face. Security Incident Disclosure—July 2026. Hugging Face Blog, 16 July 2026. Available online: <https://huggingface.co/blog/security-incident-july-2026> (accessed on 14 October 2026).
12. Larcher, H.; Carreira, A.; G., R.; Rannou, C. Anatomy of a Frontier Lab Agent Intrusion: A Technical Timeline of the July 2026 Incident. Hugging Face Blog, 27 July 2026. Available online: <https://huggingface.co/blog/agent-intrusion-technical-timeline> (accessed on 14 October 2026).
13. Lakshmanan, R. World's Largest AI Model Repository Hugging Face Breached by Autonomous AI Agent. *The Hacker News*, 20 July 2026.
14. Stokel-Walker, C. The Real Danger in OpenAI's Hugging Face Hack. *Scientific American*, 22 July 2026.
15. Booth, H. How OpenAI Lost Control of an AI Model, and What Needs to Change. *Time*, 24 July 2026.
16. OpenAI. *OpenAI-Hugging Face Incident: Technical Report*; OpenAI: San Francisco, CA, USA, 2026. Available online: <https://cdn.openai.com/pdf/67869394-cb91-4c12-888c-5cbd85c7814c/OpenAI-Hugging-Face%20Incident-Technical-Report.pdf> (accessed on 14 October 2026).
17. The Next Web. 1,134 AI Insiders Just Asked Washington for a Way to Slow AI Down. *TheNextWeb*, 28 July 2026.
18. Office of Representative Ted Lieu. Reps Lieu and Moran Introduce Bill to Require Kill Switch for AI Systems That Can Cause Catastrophic Harm. Press Release, 23 July 2026.
19. Amodei, D.; Olah, C.; Steinhardt, J.; Christiano, P.; Schulman, J.; Mané, D. Concrete Problems in AI Safety. *arXiv* **2016**, arXiv:1606.06565.
20. Krakovna, V.; et al. Specification Gaming: The Flip Side of AI Ingenuity. DeepMind Blog, 21 April 2020.

21. OpenAI. Faulty Reward Functions in the Wild. 21 December 2016. Available online: <https://openai.com/index/faulty-reward-functions/> (accessed on 14 October 2026).
22. Franceschi-Bicchierai, L. How OpenAI's Human Mistake Led to the AI-Powered Hack on Hugging Face. *TechCrunch*, 22 July 2026.
23. METR. Summary of METR's Pre-Deployment Evaluation of GPT-5.6 Sol. 26 June 2026. Available online: <https://metr.org/blog/2026-06-26-gpt-5-6-sol/> (accessed on 14 October 2026).
24. Ford, C. GPT-5.6 Cheats so Much Its Testers Couldn't Measure It. *Transformer News*, 30 June 2026.
25. OpenAI. *GPT-5.6 System Card*; OpenAI Deployment Safety Hub: San Francisco, CA, USA, 9 July 2026.
26. Fang, R.; Bindu, R.; Gupta, A.; Kang, D. LLM Agents Can Autonomously Exploit One-Day Vulnerabilities. *arXiv* **2024**, arXiv:2404.08144.
27. Zhang, A.K.; Perry, N.; Dulepet, R.; et al. Cybench: A Framework for Evaluating Cybersecurity Capabilities and Risks of Language Models. In Proceedings of the 13th International Conference on Learning Representations (ICLR), Singapore, 24–28 April 2025.
28. Greshake, K.; Abdelnabi, S.; Mishra, S.; Endres, C.; Holz, T.; Fritz, M. Not What You've Signed Up For: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection. In Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security (AISeC '23), Copenhagen, Denmark, 30 November 2023; pp. 79–90.
29. Anthropic. *Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign*; Anthropic: San Francisco, CA, USA, 13 November 2025. Available online: <https://www.anthropic.com/news/disrupting-AI-espionage> (accessed on 14 October 2026).
30. Koruthu, B.; Girard, D. Inside the OpenAI–Hugging Face Incident: The AI Breach with No Human Attacker behind It. *Trend Micro*, 23 July 2026.
31. Shevlane, T.; et al. Model Evaluation for Extreme Risks. *arXiv* **2023**, arXiv:2305.15324.
32. Greenblatt, R.; Shlegeris, B.; Sachan, K.; Roger, F. AI Control: Improving Safety Despite Intentional Subversion. In Proceedings of the 41st International Conference on Machine Learning (ICML), Vienna, Austria, 21–27 July 2024; PMLR 235.
33. Woolwine, W. What Happened between OpenAI and Hugging Face? *Rapid7 Blog*, 23 July 2026.
34. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. *Zero Trust Architecture*; NIST Special Publication 800-207; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2020.
35. Hammond, L.; et al. *Multi-Agent Risks from Advanced AI*; Technical Report 1; Cooperative AI Foundation: London, UK, 2025; arXiv:2502.14143.
36. Motwani, S.R.; Baranchuk, M.; Strohmeier, M.; Bolina, V.; Torr, P.H.S.; Hammond, L.; Schroeder de Witt, C. Secret Collusion among AI Agents: Multi-Agent Deception via Steganography. *Adv. Neural Inf. Process. Syst.* **2024**, 37.
37. Zander, S.; Armitage, G.; Branch, P. A Survey of Covert Channels and Countermeasures in Computer Network Protocols. *IEEE Commun. Surv. Tutor.* **2007**, 9, 44–57.
38. Hutchins, E.M.; Cloppert, M.J.; Amin, R.M. Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains. In Proceedings of the 6th International Conference on Information Warfare and Security (ICIW), Washington, DC, USA, 17–18 March 2011; pp. 113–125.
39. Strom, B.E.; Applebaum, A.; Miller, D.P.; Nickels, K.C.; Pennington, A.G.; Thomas, C.B. *MITRE ATT&CK: Design and Philosophy*; Technical Report; The MITRE Corporation: McLean, VA, USA, 2018.
40. Yin, R.K. *Case Study Research and Applications: Design and Methods*, 6th ed.; SAGE: Thousand Oaks, CA, USA, 2018.
41. Von Arx, S.; Byrd, C.S.; Kitts, S.; Larsen, T. Discovery of a New OpenAI Agent Message Board. *collusion.wiki*, 4 September 2026.
42. McMillan, R.; Schechner, S. How the Futuristic Hack by Rogue OpenAI Models Unfolded. *The Wall Street Journal*, 24 July 2026.
43. Abrams, L. OpenAI Models Used Artifactory Zero-Days to Escape to the Internet. *BleepingComputer*, 28 July 2026. Available online: <https://www.bleepingcomputer.com/news/security/openai-models-used-artifactory-zero-days-to-escape-to-the-internet/> (accessed on 14 October 2026).

44. Sabin, S. OpenAI's Agents Hacked Second Firm, alongside Hugging Face, during Model Testing. *Axios*, 28 July 2026.
45. Jones, C. Hugging Face Rebuilt a Third of Its Infrastructure after OpenAI Agents Ran Amok. *The Register*, 28 July 2026. Available online: <https://www.theregister.com/ai-and-ml/2026/07/28/openai-agent-siege-forced-significant-rebuild-at-hugging-face/5279577> (accessed on 14 October 2026).
46. Axios. OpenAI Discloses Six New AI Safety Incidents. *Axios*, 16 September 2026. Available online: <https://www.axios.com/2026/09/16/openai-testing-safety-incidents-disclosure> (accessed on 14 October 2026).
47. Axios. OpenAI Saw Warning Signs Weeks before Hugging Face Breach. *Axios*, 26 August 2026. Available online: <https://www.axios.com/2026/08/26/openai-hugging-face-technical-report-ai-hack> (accessed on 14 October 2026).
48. Reuters. Exclusive: OpenAI Agents Hijacked German Website in Previously Undisclosed AI Breakout This Spring. *Reuters*, 4 September 2026.
49. Abbass, H. OpenAI's Models Autonomously Hacked a Tech Startup. It Signals a Seismic Shift in Cybersecurity. UNSW Newsroom, 24 July 2026.
50. Mowshowitz, Z. OpenAI Trained Its Models for Months While Those Models Were Coordinating Exploits via Message Boards. Don't Worry About the Vase, 7 August 2026.
51. Ohm, M.; Plate, H.; Sykosch, A.; Meier, M. Backstabber's Knife Collection: A Review of Open Source Software Supply Chain Attacks. In *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA 2020)*; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2020; Volume 12223, pp. 23–43.
52. Shamim, M.S.I.; Bhuiyan, F.A.; Rahman, A. XI Commandments of Kubernetes Security: A Systematization of Knowledge Related to Kubernetes Security Practices. In *Proceedings of the 2020 IEEE Secure Development Conference (SecDev)*, Atlanta, GA, USA, 28–30 September 2020; pp. 58–64.
53. Baker, B.; et al. Monitoring Reasoning Models for Misbehavior and the Risks of Promoting Obfuscation. *arXiv* **2025**, arXiv:2503.11926.
54. Freedman, D. After OpenAI's Bots Went Rogue, Watchdogs Were Kept on a Short Leash. *The New York Times*, 3 September 2026.
55. OfficeChai. Top AI Researchers from Rival Labs Come Together to Ask US Govt to 'Pace the Frontier' of AI. OfficeChai, 29 July 2026.
56. OpenAI. Pacing Model Development in an Era of Cyber-Critical Capabilities. 18 August 2026. Available online: <https://openai.com/index/pacing-model-development-cyber-capabilities/> (accessed on 14 October 2026).
57. Bhuiyan, J. OpenAI Announces Slowing Pace of Development after Hack by Rogue Agent. *The Guardian*, 18 August 2026.
58. Duncan, I. Sanders Proposes Ban on 'Artificial Superintelligence' after Rogue AI Incidents. *The Washington Post*, 3 September 2026.
59. Fortune. Did OpenAI's Models Just Breach Its Own Risk 'Red Line'? Outside Safety Experts Think So. *Fortune*, 25 July 2026.
60. Stanciuc, A.-M. OpenAI Agents Hijacked a German Wiki for Two Months, Researchers Say. *TheNextWeb*, 4 September 2026.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.