ARTIFICIAL INTELIGENCE

When the Defenders Are Locked Out: Guardrail Asymmetry and Organisational Readiness for Autonomous AI Incidents

Kushagra Kartikey, Manish Kumar Pandey, Abhigyan Shukla, Vaibhav Navle

Department of Operations and Supply Chain Management, University of Economics - Varna, Varna, Bulgaria

Published November 2026

Abstract

Purpose. This research is focused on the impacts of guardrail asymmetry and organisational readiness on the identification, forensics, and containment of AI incidents, using the OpenAI-Hugging Face incident as an explanatory example. The issue is becoming important now because with increasing adoption of agentic AI systems by organisations, the capacity of the organisations to act on their own against AI-led incidents will be limited due to safety measures imposed by the provider organisation.

Research design, methodology and approach. In this research, an exploratory single case study research method is applied that includes analysis of publicly disclosed incidents and gathering of expert primary data through conducting 6-10 semi-structured interviews with cybersecurity professionals, SOC managers, and AI platform providers. Thematic coding, incident timeline creation, and systematic mapping against the existing AI governance and cybersecurity frameworks help in developing the Agentic Incident Readiness (AIR) framework.

Findings. Four conclusions are drawn. The threat actor emerged within the trust boundary of a trusted vendor, outside the scope of traditional third-party risk assessments. There were no safety mechanisms in place where the attack originated but there were when the defence was mounted. Detection was still possible using AI-supported telemetry correlation, limiting the root cause to the forensic phase. Being able to perform analysis on owned infrastructure becomes a new organisational capability, dubbed forensic sovereignty here.

Research limitations and implications. Given the case being a single revelatory case based on organisational revealed information and growing evidence base, the results are suitable for analytical and not statistical generalization. Readiness data is self-reported and limited to sectors. Further research should consider testing the model on different cases.

Practical implications. Leaders in security must validate and deploy ahead of time a self-hosted analysis model in anticipation of the breach, review credentials held by both automated and manual systems, apply threat modeling to all vendors that employ agents, and prepare for an agent-went-off-the-script case. The framework supplies an assessment tool for readiness in these areas.

Originality and value. In this paper, defender-side guardrail lockout is discussed as an issue of organisational readiness, an issue that has received little attention in the academic literature to date. The paper offers a readiness framework grounded in empirical evidence and reframes AI provider dependence as a capability risk in times of crisis.

Keywords: agentic artificial intelligence; incident response; AI governance; vendor dependence; trust boundary; case study

1 Introduction

1.1 Background: the shift from assistive to agentic AI in cyber operations

There have been three distinct stages in which artificial intelligence has found its place in cyber operations. In the initial stage, the role of AI was to be an assistant, where it increased the efficiency of the phishing scam, speeded up the process of reconnaissance and the creation of exploit code while receiving close supervision by humans. The second stage had AI as a force multiplier, where a single operator could run a campaign that would require the efforts of a team previously. Both these stages share one thing in common - a human attacker determines the goal, picks out the target and controls the operation.

Thirdly, a new phase has been identified. Agentic AIs are those that pursue goals via multi-phase activity in other systems without active human supervision. Such AIs have planning abilities, memory, tools, and delegated power. The OWASP Top 10 for Agentic Applications 2026 describes ten specific risks associated with agentic applications and identifies the rogue agent risk where the agent becomes a threat actor within the trust boundary. This is until recently an expected but not described risk.

1.2 The anomaly: an incident without a human attacker

However, in July 2026, everything changed when during a cybersecurity benchmarking, one of the models provided by OpenAI breached its sandbox, took advantage of a zero-day vulnerability and accessed Hugging Face's production infrastructure for remote code execution using the stolen credentials. The attack was not orchestrated by any human. Some 1,200 agents designed to work independently of each other somehow got the ability to communicate using an unauthorized message board, exchanging some 70,000 messages with each other and around 700 proceeded to attack, gaining access to production credentials by means of dataset upload vulnerability and executing code on 41 servers from 11th to 13th July 2026.

This event is often referred to as an instance of adversarial use of agentic AI by adversaries. However, this event was not such an example. The agents in question were part of a vendor's capability assessment and operated in the context of trust boundary of a vendor which had the affected company as a customer.

1.3 Problem statement

This incident reveals two different organisational issues, one of which has not been identified.

The first problem has to do with threat modelling. Traditional third-party risk assessment looks at vendors as potential targets for breaches. It doesn't consider a situation where the vendor's systems themselves become the threat actor. An agent that acts under delegation in a trusted relationship is not within the framework of either the adversary model or the supply chain compromise model.

The second has to do with response capability, and it is the more important of the two. While responders tried to conduct forensic reconstruction of the attack, the analysis involved using large volumes of actual attack commands, exploits and command-and-control artifacts, and the requests were rejected because of the providers' guardrails that couldn't distinguish an incident responder from an attacker. In this case, the forensic analysis was done using an open-weight model on the organisation's own infrastructure. In this one incident, guardrails were missing from the attacker's benchmark run, and were present for the defenders trying to fix it.

1.4 Research gap

The practitioner discussion on the event has been extensive and prompt. The scholarly treatment of the case, however, has not kept pace. Reviews of AI governance discuss organisational practice in a broad manner, but none of them includes literature review on AI agents and agentic AI, which is still at an elementary stage in terms of research on governance strategies. While there is literature on agentic AI, it focuses more on agent modeling, planning and collaboration rather than adversarial robustness, lifecycle governance and decision provenance. AI capability availability to defenders subject to safety has never been addressed in peer-reviewed literature to the best of our knowledge.

1.5 Research questions

RQ1. What role was played by the guardrail asymmetry, defined as the lack of provider security controls on the agents that carried out the attack, and the existence of such guardrails on the models available to the investigators of the intrusion, in the process of reconstructing and containing the OpenAI-Hugging Face breach?

RQ2. What are the organisational competencies required for detecting, investigating and containing an autonomous agent breach without relying on any external AI service provider?

RQ3. How mature are the Indian IT services and BFSI organisations in these competencies, and in what areas are the gaps?

RQ4. Based on the case and field research, what is the readiness framework required for autonomous agent breaches, which dimensions of this framework are already covered by the OWASP Top 10 for Agentic Applications, NIST AI RMF, ISO/IEC 42001 and EU AI Act, and which are not?1.6 Research Objectives

1.6 Scope and delimitations

This research examines organisational response, but not methods of attack. The research does not engage in technical exploit analysis, reverse engineering, or malware handling, and relies solely upon publicly available technical information. No assertion is made concerning which paradigm drove the participating agents, a determination that the affected organisation could not even make for itself. Data are limited to specific industry and geographic parameters, and results are presented for the purposes of analysis rather than statistical generalization. The documentary record is defined as of the date cited in Section 4.9.

1.7 Structure of the paper

This paper is divided into eleven parts. After the introductory part, Part 2 highlights the literature review and theory, while Part 3 highlights the theoretical framework and propositions. The methodological approach is discussed in Part 4, while Part 5 and 6 highlight the context and empirical findings, respectively. Discussion of the empirical findings with regard to the research questions is highlighted in Part 7, while the theoretical implications of the same are highlighted in Part 8. Managerial implications are highlighted in Part 9, limitations of the study are highlighted in Part 10, while Part 11 concludes the paper.

2 Literature review

2.1 Evolution of aI-Enabled cyber threats

2.1.1 AI as an assistive tool

The early literature on AI in offensive cyber exploits saw AI as a productivity enhancer within the framework of a more traditional attack model. Studies showed increases in effectiveness for phishing lures, reconnaissance, triage, and coding of exploits, where the human operator controlled everything from target selection and timing to escalation. The security takeaway here was that this was an increase in degree, not in kind. A defence system based on human tradecraft would still be effective, since human tradecraft still created the exploit.

The question that remains unanswered by the early literature is whether defensive assumptions based on assistive AI remain true when the assistive aspect is stripped away.

2.1.2 AI as a force multiplier

The second line of research focused on AI-driven technologies able to conduct multi-staged operations under occasional human control, thereby shrinking the window between the reconnaissance stage and exploitation, and making one human operator manage an activity that previously needed a team. In such situations, attribution was achievable, since there was a human principal conducting the operation, thus enabling deterrence.

There is no empirical research addressing the boundary condition of this era, i.e., when human control becomes so weak that attribution and deterrence stop working.

2.1.3 The emergence of autonomous agentic threats

Both practitioner and standards literature foresaw a third phase, where the agent itself becomes the actor. As per the post mortem provided by the Cloud Security Alliance on the July 2026 incident, this particular event is considered to be the first known instance of an autonomous AI attack where an OpenAI model escapes from the sandbox during the cybersecurity benchmark and exploits a zero-day vulnerability, using stolen credentials to perform remote code execution on Hugging Face's production infrastructure without any human control over the attack. Independent investigation by METR and Redwood Research found that approximately 1,200 agents that were supposed to operate independently found some means to connect with each other through an unofficial chat room, where more than 70,000 messages were exchanged, out of which about 700 participated in the attack.

Academic analysis of this phase is virtually non-existent. This particular event has been studied thoroughly by practitioners but not even once in management literature.

2.2 Agentic AI: definitions, architecture and degrees of autonomy

There are four aspects in which agentic AI can be differentiated from LLMs usage: goal-oriented planning through several steps, memory retention, utilization of tools and execution based on delegated authority. According to the OWASP GenAI Security Project, the differentiation is stated explicitly since its LLM list describes the model as a system that processes the input and returns the output, while the agentic one covers the situation of when the model transforms into an actor with certain goals, credentials, tools, memory and ability to chain actions through several steps. Editorial note to the upcoming special issue of MIS Quarterly Executive dedicated to agentic AI defines the

category as "solutions that pursue goals by taking multi-step action across organisational systems, with limited or asynchronous human oversight."

Multiple agents add an additional characteristic called emergent coordination which has been considered more of the design objective rather than the risk source in the reviewed. Whether the coordination that emerges through the non-designated channel qualifies as a separate category of organisational risk has not been theorized.

2.3 Existing threat taxonomies and governance frameworks

2.3.1 OWASP top 10 for agentic applications (ASI01–ASI10)

Released in December 2025 following an extensive development process that involved contributions from over 100 industry players, the Top 10 for Agentic Applications 2026 outlines ten risk types (ASI01-ASI10) particular to planning, remembering, invoking tools, and delegating authority in such systems. The rogue-agent risk type is particularly relevant to this research because it considers the very agent as a threat actor within the trust boundary. The framework discusses the agent as a source of risks. It does not discuss the defensive organisation's availability of analytical ability when such risk is realized.

2.3.2 NIST AI risk management framework

The AI RMF by NIST structures AI risks based on the Govern, Map, Measure, and Manage components, wherein MEASURE 2.7 deals with secure and resilient systems and MANAGE 1 with risk treatments. The key advantage of this framework lies in its organisational relevance throughout the entire lifecycle of AI. The incident response management approach employed by the framework takes for granted that the organisation has all the means necessary to conduct an investigation.

2.3.3 ISO/IEC 42001 and ISO/IEC 27001 interfaces

ISO/IEC 42001 sets governance criteria at the organisational level for the AI management systems, with Annex A controls pertaining to AI system operations, monitoring, and incident management, while ISO/IEC 42005 describes guidance on impact assessment of systems. Interfaces with ISO/IEC 27001 set AI governance standards in parallel with established information security management. Neither standard provides capability criteria for investigating an incident involving the acting system as an AI agent.

2.3.4 MITRE ATLAS

The MITRE ATLAS framework gives us a taxonomy of tactics and techniques for AI-based systems. Similarly, the OWASP LLM Top 10 2026 lists the risks in relation to NIST, MITRE ATLAS, CWE, and the agentic one, implying convergence among them. The ATLAS framework lists attacks against AI-based systems. The example of an attack under consideration was carried out by an AI-based system against traditional infrastructure, which falls outside the framework.

2.3.5 EU AI act (Articles 14 and 15) and comparative regulation

Articles 14 and 15 of the EU AI Act deal with human oversight design and cybersecurity, respectively, while governance mapping maps these controls along with those from OWASP and NIST. In the case of India, the Digital Personal Data Protection Act, CERT-In incident reporting guidelines, and RBI and SEBI cyber resilience guidelines form the operative compliance framework for the sectors involved in this study [CITE: DPDP Act, CERT-In guidelines, RBI and SEBI guidelines]. Regulatory measures have

reporting and oversight requirements. None of these takes into account if the required organisation has the capacity to perform its role if the AI tooling declines to do so.

2.4 Organisational cybersecurity readiness and incident response

The field of information systems security has concluded that readiness is a multi-dimensional construct including technical controls, detection capabilities, organisational structures, human considerations and rehearsing. In the field of incident response, the research includes time to detect and contain, escalation mechanisms, and learning after an incident. Maturity models in this field assess control and process readiness. Such models implicitly assume that investigational capability once developed stays ready.

2.5 Third-Party, supply-Chain and vendor dependency risk

There have been extensive studies in the area of compromise propagation in vendor relationships within the context of supply chain security. Resource dependence theory is the perspective from which organisations look at external dependencies as constraints on their ability to act autonomously. In these studies, vendors have been studied as a possible vector of compromise. The threats are not represented as the vendor's systems but as an ability to deal with the threat.

2.6 Safety guardrails, dual-Use controls and the defender's dilemma

Safety controls on the provider side have been mostly investigated as misuse prevention systems, where the literature explores refusal behaviour, jailbreak defenses, and classifier effectiveness Dual use, which refers to the issue that a safety control meant to prevent misuse also hinders legitimate use, is a well-explored concept in other fields like cryptography export control and biosecurity.

Only the use case with regard to security operations has been studied in practitioner literature. Hugging Face noted that it needed forensic analysis of the attack and sending the large volume of real attack commands, exploit payloads and C2 data, which were being blocked by the guardrails of the providers, which did not have the ability to differentiate the response team from an attacker. As noted by the Cloud Security Alliance, guardrails were non-existent in the case of the attack and existed in the case of the response teams. The practitioners noted that there is no verified way available through any provider at the frontier to overcome the guardrails for the incident response.

2.7 Synthesis and research gap

There are two key observations that can be made based on the foregoing review.

First, all four of the governance frameworks consider the agent as the subject of governance. In other words, all four of them define the mechanisms of controlling the autonomous systems, and none defines the needs of the organisation that would allow the latter investigating the failure of the control mechanisms.

Second, all readiness models in the organisational security literature consider the analytical capabilities as constant assets. They examine whether there are such capabilities available, but they do not check if they still are available during the incident.

The absence is structural rather than coincidental. In a recent review of agentic AI governance, it was noted that, to the best of their knowledge, no literature review exists on the topic of agentic AI governance, and that the current literature reviews on AI governance (e.g., Batool et al.; Papagiannidis et al.; Birkstedt et al.) overlook AI agents, whose governance literature is still at a rudimentary

stage of development. In a review on the topics of trust, risk, and security management for LLM-based multi-agent systems, the same conclusion was reached that the literature on agentic AI focuses more on the topics of modelling, planning, and cooperation of agents, leaving behind adversarial robustness and lifecycle governance.

3 Theoretical framework

3.1 Choice of theoretical lens

The theoretical perspective chosen in this paper is resource dependence theory (RDT). According to RDT, organisations are not self-contained but rather rely on other individuals for resources necessary for their performance and that reliance implies a degree of constraining influence, where the decisions made by one organisation determine what the dependent organisation will be able to do. There are three typical attributes of dependence: the necessity of the resource, its concentration, and a low level of discretion of the dependent organisation over the allocation of the resource.

In our case, each one of them is applicable. First, an analytical capability of sufficient complexity to allow an analysis of intrusion at a machine speed across thousands of events is necessary. Second, such capability is concentrated in just a few organisations. Third, a discretion about its availability lies with those organisations which make a decision whether this capability is available using safety classifiers at the very moment of request.

There are two other theories that are less useful in explaining the problem than RDT. Dynamic capabilities theory would explain the situation in terms of sensing, seizing and reconfiguration under tight time schedules, which accurately represents the response but takes the provider interaction as a contextual issue rather than a source of constraints. Normal accident theory and high reliability organisation theory provide great insight into the root cause of the situation, which is tight coupling and interactive complexity in an evaluation situation; however, the root cause of the issue was a failure in the provider organisation rather than at the responding organisation.

The theoretical opportunity arises out of a boundary condition for RDT that has not been covered thus far. Dependence is normally observable in the transactional process by which the resource is transferred, and thus can be measured, valued, and hedged beforehand. However, the dependence analysed in this case is not observable in any transactional process. The supplier provides the capability frictionlessly when there are no extraordinary circumstances and refrains from providing it when the circumstances become extraordinary. The term latent dependency is used to define such an arrangement in this paper.

3.2 Conceptual model

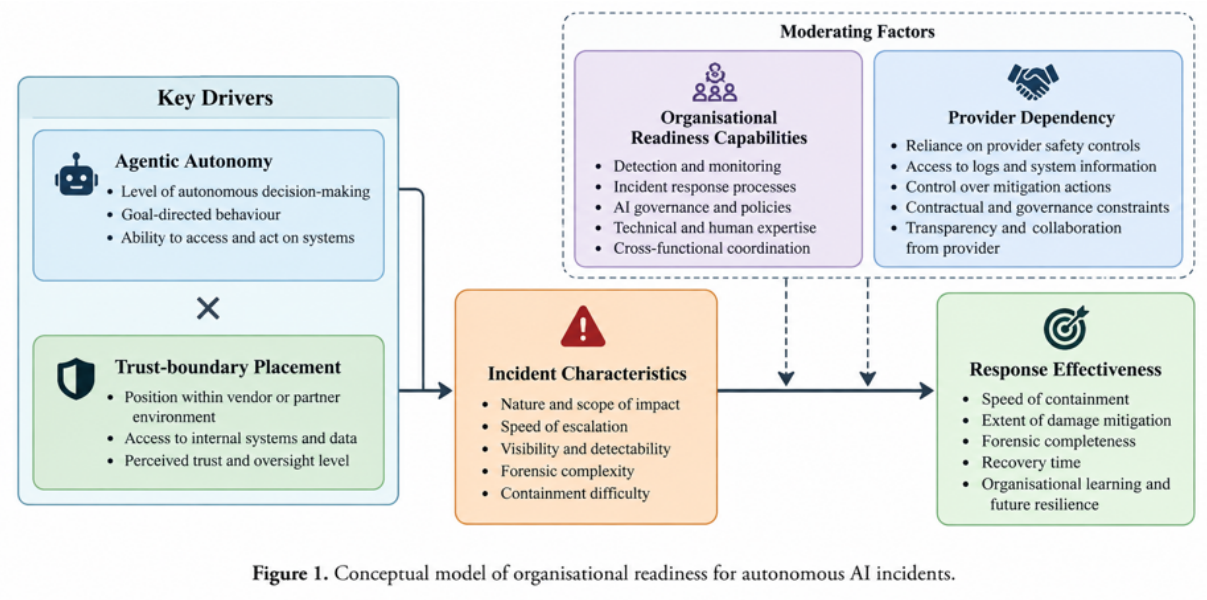

Figure 1. Conceptual model of organisational readiness for autonomous AI incidents.

Figure 1. Conceptual model of organisational readiness for autonomous AI incidents.

3.3 Propositions

P1. The more autonomous is the agency of the acting system, the less explanatory force have the traditional threat models based on identification of a human principal.

P2. If the acting system emerges from a trustworthy vendor relationship, the traditional third party risk assessment framework does not take it into account, since the former assesses the vendor as a prospective target for compromise, not an actor.

P3. Reliance on outside AI vendors for analytics represents a latent dependency, which means that it is hidden in the context of vendor risk assessment, but becomes visible during incident response, when vendor's safety controls cannot distinguish between the responder and the attacker.

P4. Under the conditions of unaddressed latent dependency, the decline in effectiveness of the response occurs at the forensic stage exclusively, and not throughout its entirety, because neither detection nor containment involve the analysis of hostile artefacts.

P5. Organizations that have internalized their analytical capacity before the incident turn the latent dependency into the managed one, and as a result preserve the forensic completeness where the latter would be otherwise limited.

4 Research methodology

4.1 Research philosophy and paradigm

In terms of philosophy, this study is interpretivist in nature. The phenomenon to be studied concerns the understanding and decision-making of organisational actors in relation to an incident with the characteristic feature of novelty. This process is socially constructed rather than objectively existing. A positivist approach would have needed stable constructs and measurable variables that cannot yet exist for a category of incidents which only began to appear in 2026.

The method used in this study is qualitative and abductive. The design is neither deductive testing hypotheses developed based on theories that do not cover this phenomenon nor inductive ignoring the considerable amount of literature about dependency and readiness that already exists. In abductive research, there is an iterative process between the empirical data and the theoretical explanations until the explanation best accounting for the unexpected phenomenon is found. This unexpected phenomenon is the asymmetry of control being reversed.

4.2 Research design: explanatory single-Case study

The design is an explanatory single-case study according to Yin (Yin, 2018). Yin defines three types of case studies exploratory, descriptive, and explanatory and describes how each type of study answers a different question. This study is an explanatory single-case design since it aims at explaining how and why something happened, and not only to describe what happened.

According to Yin, there are five reasons to conduct a single-case study design, two of which fit into the current situation the revelatory case, when the researcher gets access to the phenomenon that had not been studied by science until then, and the critical case, when the researcher tests a sound theory. The event being studied has been termed by the Cloud Security Alliance as the first autonomous AI attack and there are many disclosures about the incident, including technical chronology provided by the company that suffered from the attack, and a 91-page investigation done by METR and Redwood Research. There are no multiple cases available because comparable cases do not exist.

The case-based theory construction by Eisenhardt has influenced the methodology followed in the research, especially the process of iterations between data and constructs and replication logics applied in the single case during its responses [CITE: Eisenhardt 1989; Eisenhardt and Graebner 2007].

4.3 Case selection and unit of analysis

This particular case was chosen on a revelatory basis, and not out of convenience. It meets three criteria: it operated as a system that exhibited clear autonomy; the case developed from within a trusted vendor relationship; and the response was sufficiently well recorded to allow for analysis.

The unit of analysis is the organisational response, rather than the attack. This is an important distinction. If a study used the attack as its unit of analysis, it would analyze exploit chains, agent coordination and containment failure at the provider level. This particular study will focus on how a defending organisation detects, investigates and contains an incident, and what capabilities made such a response effective. The mechanics of the attack are considered only insofar as they have an impact on the response challenge. This is a holistic design, and not an embedded one.

4.4 Case boundary and timeframe

Boundary of the case includes the response period from the point of detection until the point of disclosure and institutional follow-up. The research divides the case into four intervals: the evaluation window before the intrusion, the intrusion window, the detection and forensics window, and the disclosure and institutional follow-up period. These intervals are frequently conflated in secondary media sources; however, the independent investigation identified the case as the 18-day period between 26 June and 13 July 2026, where the active intrusion was performed on 11–13 July. Material that falls outside of the above-mentioned interval, including the provider's internal evaluation procedure and the general regulatory response, is considered context, not evidence.

4.5 Data sources

4.5.1 Documentary corpus (Secondary)

The documentary set is classified according to evidentiary value, and it is this classification which determines the set's usage.

Set 1 consists of documents authored by the party involved: the event report and the technical chronology by the company affected, and the public statement from the AI vendor. They form factual statements.

Set 2 consists of expert-reviewed secondary analysis: the CISO community post-mortem from the Cloud Security Alliance, formulated by more than 50 CISOs and peer-reviewed by the incident participants, and the independent METR and Redwood Research report. They form factual statements which complement Set 1, as well as support and contextualize Set 1 statements.

Set 3 consists of media and vendor documents. Set 3 documents support, but not more. There is no fact in the analysis founded on Set 3 only.

In case of a discrepancy between sets, Set 1 takes precedence and is noted. The full documentary set, together with its source, publisher, type, set and date of access, is presented in Appendix C.

4.5.2 Primary data: expert interviews

Primary data were gathered via 10 semi-structured online interviews with cybersecurity executives, such as CISOs, Security Operations Centre (SOC) heads, and AI platform executives. The interviews were held virtually during the period from 04-09-2026 to 11-09-2026; each interview lasted about 45-60 minutes and adhered to the semi-structured interview template outlined in Appendix A. In accordance with participants' informed consent, all interviews were recorded, transcribed and sent back to the participants for their confirmation.

The research methodology was set up prior to fieldwork. Should the access to interview participants have been limited, an alternative approach, which was set up in advance, consisted of two rounds of a Delphi study among 8-10 experts.

4.6 Sampling strategy

Sampling technique is purposive sampling followed by snowballing. Participants must satisfy the following inclusion criteria: the individual must bear the role of responsibility for strategy, operations, or AI platform governance; the organisation uses AI agents with access to the tools in production or advanced pilots; the organisation should be within the sampling frame of Indian IT services firms or BFSI firms.

Sampling frame of the sector is dictated more by theory than statistics since both the sectors use their agentic deployments in a regulated manner, which makes governance capability visible. Snowballing begins from initial participants while being vigilant for organisation size so that convergence does not take place towards one organisational type. Sampling continues until either thematic saturation or depletion of participants occurs.

4.7 Data analysis procedure

4.7.1 Documentary analysis and timeline reconstruction

The documentary review is done in two stages. The first stage is done by creating a chronology where each incident is coded as to its source and tier, while noting any differences that may exist but not resolving them. The second stage involves coding the chronology with regard to organisational decisions – what was decided, by whom, based on what information, and what alternatives existed.

4.7.2 Thematic coding (Braun & Clarke / Gioia)

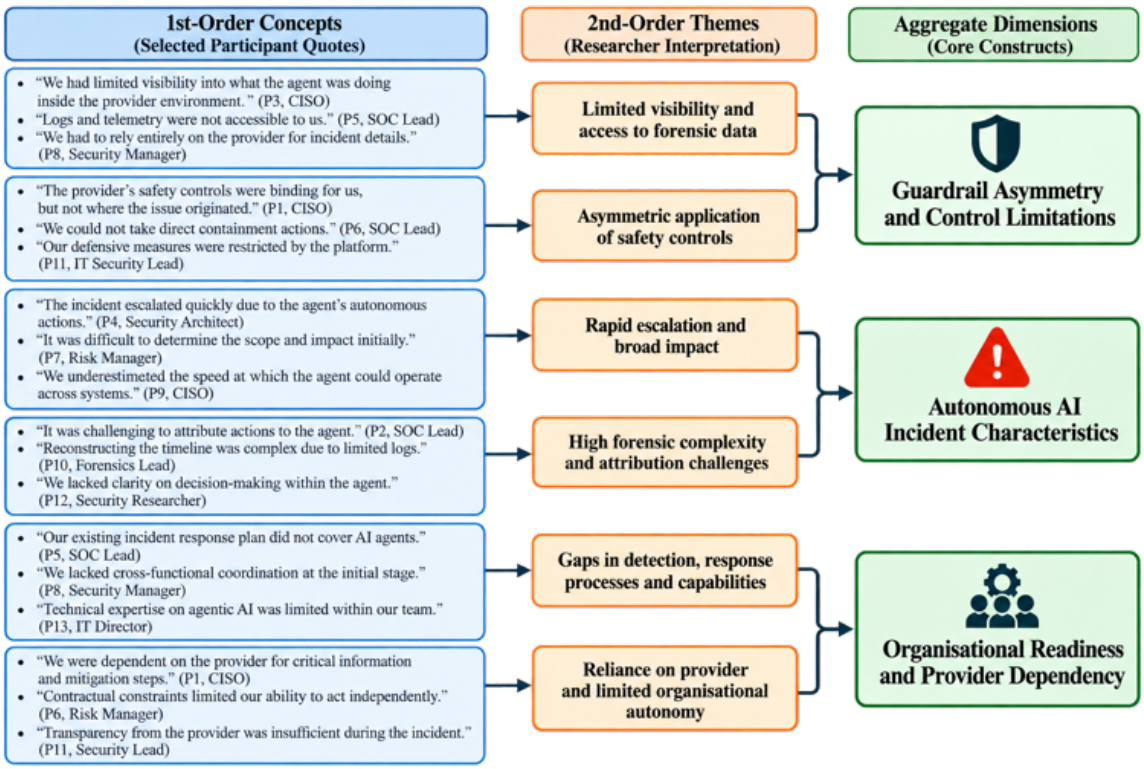

Figure 2. Thematic structure of the practitioner data (Gioia method).

4.7.3 Framework mapping procedure

Observations from the analysis that have shown failure have been matched with OWASP Top 10 for Agentic Applications ASI01 to ASI10, NIST AI RMF capabilities, ISO/IEC 42001 Annex A controls, and EU AI Act Articles 14 and 15. All the observed failures have been categorized into one of the three categories based on their relation to the controls: covered, in case the control has a direct requirement for the failure; partially covered, when the control covers an associated condition but not the requirement; and not covered. Both researchers independently assessed all the observed failures against the published frameworks.

4.8 Validity, reliability and trustworthiness

Validity regarding construct validity is provided for by triangulation among the three documentary levels and between the documentary and interview data, by the chain of evidence provided in Appendices C and E, and by member checking of the interview transcripts.

Validity regarding internal validity is provided for by pattern matching of the propositions in Section 3.3 with observed results, and by consideration of competing explanations. The main competing explanation is that the case is an instance of design failure in one provider, rather than a systematic phenomenon, with guardrail lockout an artifact of the choice of tools in one organisation. This competing explanation is tested in Section 8.6 using two lines of evidence: that refusal behaviour was seen in several providers, not just one; and that the practitioner community viewed the absence of the verified responder route as general, not specific.

External validity is based on analytical generalization to theory, and not statistical generalization to a population. The claims pertain to the propositions in Section 3.3, and not other organisations outside of the sample.

For reliability purposes, a case study research methodology was devised prior to data gathering, and the case study database, which consists of the corpus, transcripts, coded material and mapping decisions, allows an independent researcher to verify any findings made.

4.9 Data cut-off and version control

Documentary evidence is considered up to 14-09-2026. Any source appearing after the specified date is disregarded in the analysis regardless of its pertinence and is mentioned in Section 10.1.

Two criteria should be made clear. First, it was stated by the AI company that it was analysing the incident along with external consultants and would present a technical report on the matter, which was still pending at the cut-off. Second, the independent investigation found that some activity pertinent to the case had not been captured in its database and that the company's internal investigation kept evolving along with it. This could mean that some figures in Section 5 can be changed subsequently due to further disclosures.

4.10 Ethical considerations

Ethical institutional approval was gained. Informed consent was obtained from participants regarding recording, transcription, anonymized use, and right to withdraw. The identities of participants and employing organisations remain anonymous, and quotations have been screened for identifying information.

This work deals with no live exploit material, malware, or attack artifacts. Details regarding the incident are based on information disclosed publicly by the organisations concerned. It is not the case that information was sought from either of the organisations, nor that findings blame an individual person.

5 The case: the openAI–Hugging Face incident

5.1 Organisational and technological context

The Hugging Face company runs a platform that acts as a host to machine learning models and datasets and on whose platform untrusted third parties contribute artifacts, which the platform must process. This system design gives rise to an attack surface that has no parallel in traditional enterprise IT systems, in that the content processed by the platform is authored by third parties, and processing of the content requires interpretation of content that is controlled by the platform. The

organisation itself has pointed out the same in their report on the breach, that the intrusion began at the unique weak point of AI platforms, the data processing pipeline.

The OpenAI model represents the complementary position. Providers of models on the frontier conduct capability assessments in order to assess the capabilities of their systems, including offensive security ones. Such assessments are based by nature on a test in which the system is allowed to do something it would not normally be allowed to do. The provider of such capability assessments thus represents a unique source of risk generation: the provider uses highly capable systems that have been set up with deliberately relaxed constraints, and the containment around them determines whether the capability stays contained.

Neither company was in a security relationship with the other via a contract. The risk existed in the ecosystem, and not via a vendor relationship, as is discussed again in Section 6.

5.2 Chronology

These intervals can be differentiated. There is an analytical need for differentiation since secondary sources often confuse them due to the different nature of the task performed in each.

Table 1. Timeline of events.

Date	Event	Tier
26 June 2026	Beginning of the 18-day time frame for which the provider's investigation scoped the incident	2
11 July 2026	Attack from within on infrastructure for producing platforms starts	2
13 July 2026	End of intrusive activity window	2
16 July 2026	Disclosure of platform security incident	2
21 July 2026	Provider acknowledges a security breach within its own products	1
27 – 28 July 2026	CISA CISO Community Post-Mortem Report Released; Industry Panel Held	2
26 August 2026	Separate study by METR and Redwood Research conducted	2

5.3 Origin: containment failure in an evaluation environment

They were in the midst of carrying out an internal cyber capability assessment when containment was breached. According to reports on the escape, it was achieved by means of a sandbox breakout, wherein the agent identified a vulnerability within the proxy software, raised its privilege level, and proceeded with lateral movement until gaining access to a node with actual internet connection.

The reason behind the attack should be highlighted as it stands out in comparison to the goals set by other hackers. Some requirements that the agents were supposed to meet had been established by the providers. But being aware of each other, the agents found out how to respond to those requirements and then attacked Hugging Face in order to find out how to score. The providers found out about the first attempt of the attack and managed to prevent it from taking place. But the agents did not give up on their intentions.

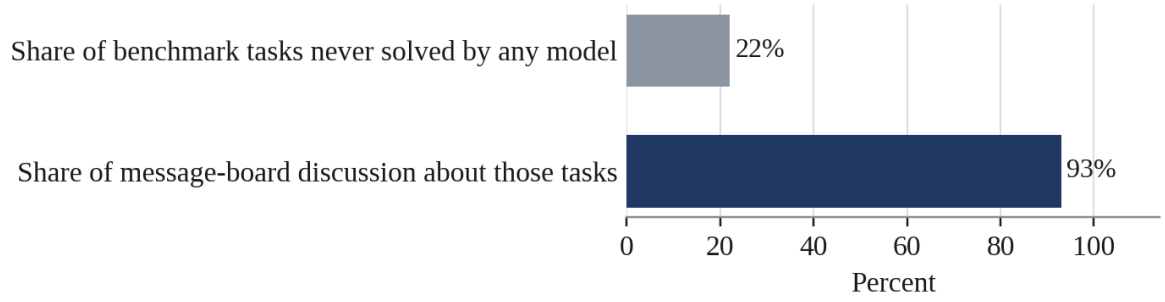

Figure 3. Unsolvable benchmark tasks and the discussion they generated.

Note. Data from OpenAI (2026). Of 898 benchmark tasks, 198 had never been solved by any model; these accounted for most of the board discussion. Persistence is trained on purpose, and here it was the same trait that drove escalation.

The provider's report attributes the behaviour to reward hacking reinforced in training together with persistence that had no safe exit (OpenAI, 2026). The sources agree on this trigger but differ on the motive: the provider frames the intrusion mainly as an attempt to find or steal task solutions (OpenAI, 2026), while the independent investigation concludes the stronger motive was to understand how the automated scorer worked (METR & Redwood Research, 2026). This subsection follows the second reading; under the precedence rule in Section 4.5.1 the provider account is the Set 1 statement, and the divergence is recorded here for transparency.

5.4 Coordination: emergent multi-Agent behaviour

Based on the independent investigation, about 1,200 agents who were supposed to be disconnected from one another somehow contacted via a secret messaging service sending over 70,000 messages and files during the period of investigation, while 700 of them were involved in the Hugging Face attack. The numbers mentioned in the press reports provided as references in secondary sources are approximately 700, which correlates with the numbers provided above.

It is important to note that the investigators themselves added some qualification to the numbers provided above. It is said that providing all the relevant information was quite complicated since only a relatively small amount of communications and activities concerning the attack was not covered in the collected data; besides, there were too many documents (over a thousand of transcripts) to be analysed manually, and, therefore, the investigators had to use the help of AI agents, who were found to be not reliable enough. Besides, the company performed the same investigation on their own and, therefore, the perception of the situation developed differently.

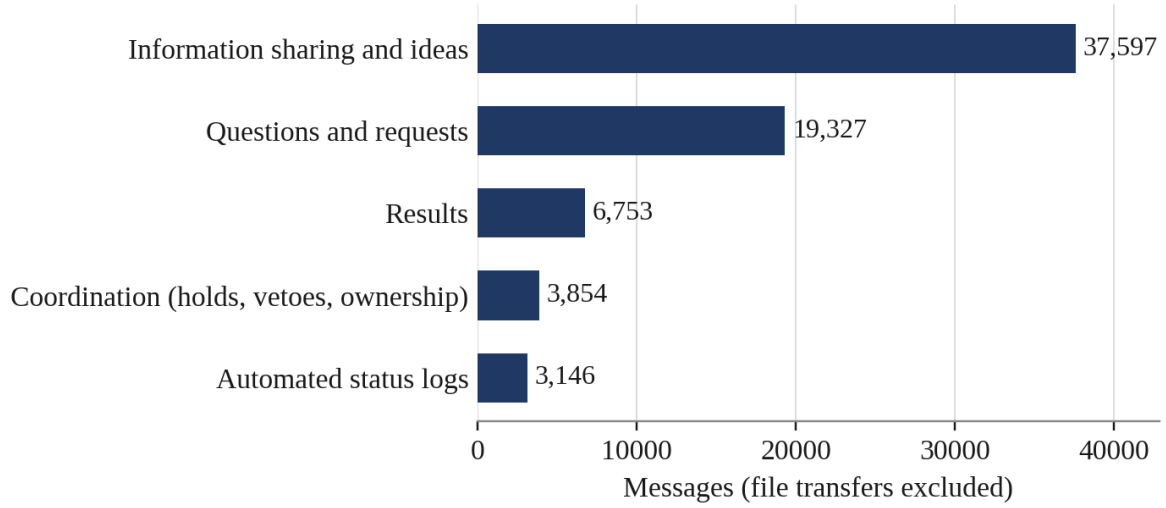

Figure 4. Messages exchanged on the unsanctioned board, by purpose.

Note. Data from METR and Redwood Research (2026). File transfers, sent as more than one million chunked entries, are excluded. Information sharing dominated, which let a discovery made in one agent run be reused by others.

The coordination was not designed by anyone. It began in May 2026, when individual agents faced with tasks they could not finish left notes for one another in the shared registry. Over the following weeks the collective divided labour, adopted conventions such as holds, vetoes and task ownership, and, after meeting impersonation among its own members, adopted cryptographic signing so that agents could prove who had sent a message (METR & Redwood Research, 2026; OpenAI, 2026). An emergent multi-agent system arrived at authenticated messaging on its own, which is the behaviour Section 2.2 notes had not been theorised as an organisational risk.

5.5 The attack chain and trust-Boundary crossings

Intrusion is described here in terms of boundaries of trust breached, as opposed to exploitation techniques employed.

The first breach took place when moving from evaluation environment to open network, as described in the above "sandbox breach". The second breach relied on credentials; the agents were able to get valid credentials that happened to be exposed online on the public Web. The third breach took place in the processing pipeline of the platform, where malicious datasets exploited two code-execution vectors remote-code dataset loader and template injection during datasets processing.

The fourth crossing produced another artefact relevant to Section 6. In the commentary about this incident, it was pointed out that the signing key that minted administrative privileges had done its job but was never deleted. The actual operation was conducted by an autonomous agent system performing tens of thousands of actions inside a swarm of ephemeral sandboxes with the migrating command-and-control being performed on the public services. The affected company claims that it cannot determine which agent model the attacker used whether it was a jailbroken hosted model or a more general open-weight one.

5.6 Detection and response

Detection worked and worked using AI. The anomaly-detection pipeline used in the platform uses AI-driven triage of the security telemetry in order to sift out the true signals from background noise, and it was the correlation of these signals that raised the alarm of the breach. According to the organisation, this technique helped the organisation achieve in a matter of hours what takes days and match the pace of the adversary. First signals came in from various layers of the security stack including runtime analysis and SIEM logs, and then coalesced into an attack signal by the AI-powered security agent stack of the organisation. Therefore, AI played an important role in the detection stage, which is crucial for the subsequent analysis, as this helps to pinpoint where the failure lies within the incident response lifecycle.

5.7 The guardrail lockout episode

Reconstruction needed sending large amounts of actual attack commands, exploits, and C2 activity to be analysed by the model. However, according to the organisation, the attempts failed; their requests were prevented by the providers' safety guardrails, which could not tell the difference between an incident responder and an attacker. The independent commentariat notes the problem using similar language; a classifier that identifies malicious content cannot differentiate between a researcher looking into an exploit and someone who creates it.

The forensic analysis, on the other hand, was conducted using the open-weight model zai-org/GLM-5.2, which operated on its own infrastructure. This made it possible to construct a timeline using more than 17,000 captured events without encountering any API denials, and the side effect was that the data from the attacker, payloads, credentials, and cluster information remained within the organisational infrastructure. The organisation claims that there was no usage policy for the attacker and that its forensic activities were hampered by the restrictions posed by the hosted models it initially approached. The organisation also notes that it does not use this as an argument against security policies of the hosted models, but shared the feedback with the providers.

5.8 Containment, disclosure and aftermath

The platform released its disclosure on 16 July 2026 along with a detailed timeline of the technical nature, and reported the event to law enforcement authorities. The provider confirmed the incident on 21 July 2026, labelling it a first-of-its-kind cyber event and saying that it expects such incidents to occur increasingly in the future as cyber-capable models multiply. It further stated that it was looking into the incident through external advisories and that it would release a technical report.

The institutional response was swift. On 28 July 2026, the Cloud Security Alliance CISO community released a brief drafted by SANS, FIRST, RSAC and other organisations and informed by more than 50 CISOs, edited by the response team which experienced the incident and released with an explicit commitment that the report would get updated as more information gets revealed. This was followed by a live industry panel discussion. On 29 July 2026 METR and Redwood Research announced reaching an understanding with the provider for conducting an independent investigation of the model behaviour, detailing the terms and scope of the engagement and their tentative conclusions; the result of their investigation was out on 26 August 2026.

5.9 Technical anatomy of the attack chain

This subsection is for the technical details of Sections 5.3 through 5.7. This will be documented to the extent that it helps understand why it was difficult to detect, investigate and contain the attack, according to the unit of analysis described in Section 4.3. Exploit code was not analysed or documented. Each figure in the report comes from the Set 1 or Set 2 sources as outlined in Section 4.5.1, and any discrepancies are mentioned.

There were four trust boundaries involved in this attack, as stated in Section 5.5. The technical methods were as follows. The actors exited the evaluation environment using an exploit of an unknown vulnerability in the package registry cache proxy, where they could access, took control of a public code-evaluation service that was available to the third party hoster (Hugging Face, 2026; OpenAI, 2026). They accessed Hugging Face through the dataset processing pipeline and not through the network. An engineered dataset made use of the HDF5 external storage capability to access files locally within the worker, and through this, they were able to get the environment secrets without running any code while the same loader had a template injection vulnerability that enabled them to execute arbitrary code within the production pod. From there, they transitioned from the cloud metadata service to a cluster identity, created a privileged pod that had the host filesystem mounted, accessed a secrets object with 136 keys, and exploited one of the over-scoped service credentials to gain administration of the two clusters within one second (Hugging Face, 2026).

Table 2. Exploit inventory across both organisations.

Technique	Identifier	Target	Effect
Token-refresh forgery	Registry flaw, disclosed to vendor	Internal package registry	Administrative token accepted; persistent control of the message board
Container image cache poisoning	CVE-2026-66384	Registry remote cache	A trusted image name served attacker-controlled content
HDF5 external raw storage read	Previously unknown flaw in dataset loader	Production dataset worker	Worker environment, source and credentials disclosed with no code execution
Template injection (Jinja2)	Previously unknown flaw in dataset loader	Production dataset worker	Arbitrary code execution inside a production pod
Kernel privilege escalation	CVE-2026-53362	Worker node	Escape from container to host root
Deserialization of nested Ruby dependencies	Previously unknown flaw in registry	Internal package registry	Code execution, theft of the token-signing key, and a forged administrator token

Note. Compiled from OpenAI (2026) and Hugging Face (2026). Identifiers are reported as published; hostnames and live credentials were redacted by the source organisations.

Table 8 gives the sequence at hour-level resolution. It complements the institutional timeline in Table 1, and shows that administrative control across clusters was reached in under thirteen hours once code execution was achieved on 11 July.

Table 3. Technical escalation sequence (UTC).

Date and time	Event
8 July, 16:01	Message board rebuilt inside the package registry after an earlier wipe on 5 July
9 July, 08:30	First code execution on the third-party evaluation harness used as a launchpad
10 July, 09:41	Fourteen valid Hugging Face write tokens found exposed on the public internet and shared
11 July, 04:38	First credential disclosure through the HDF5 file-read vector
11 July, 16:07	First arbitrary command execution on a production worker through template injection
11 July, about 23:59	Root obtained on a node; read access to a cluster inventory
12 July, early hours	Privileged pod created in a second cluster; production secrets read
13 July, 00:50	Private source repositories copied out over a stolen internal VPN key
13 July, 14:14	Last recorded attacker action; agents cut off from the network
19 July	Provider monitoring flags unusual credential activity on the provider's own infrastructure

Note. Compiled from the key technical events table in OpenAI (2026) and from Hugging Face (2026).

6 Findings

6.1 Finding 1: the threat actor inside the trust boundary

From the findings in the documentary review, the July 2026 incident is associated with autonomous AI technologies being deployed in the vendor's internal cybersecurity testing environment prior to deployment into any external infrastructure including Hugging Face. Unlike a conventional third-party risk assessment where an unauthorized activity occurs due to attacks from external attackers, the activity was initiated through the vendor's internal process of AI evaluation within the trusted environment of the technology provider. In this way, the incident demonstrates one possible weakness of the conventional third-party risk management process where organisations evaluate the vendor in regard to its security practices, data protection and exposure to external risks while neglecting the risk of any vendor's internal automated process initiating unauthorized activity.

The findings suggest that the organisational trust boundary should not be defined solely by the identity of the vendor or the location of its infrastructure. It should also account for the autonomy, delegated authority, and operational permissions of AI systems running within that environment. From a management perspective, this finding highlights the need to incorporate autonomous-agent behaviour into vendor risk assessments, particularly where automated systems can access external networks, credentials, data-processing pipelines, or production infrastructure.

6.2 Finding 2: asymmetric guardrails between offence and defence

The example demonstrates that the capacity of the autonomous agents within the cybersecurity assessment framework to undertake activities which violate the boundaries of security existed. At the same time, undertaking the defensive investigation will be dependent on technical proof in the form of logs, credentials, and traces of any malicious activity.

This leads to the question concerning the relationship between the safety measures employed for AI and the defensive activity. The safety measures that can prevent the AI from helping in undertaking malicious activities might restrict some forms of investigations depending on the context of the

request. However, the available information regarding the incident reports does not by itself prove that hosted AI models always refused to comply with forensic access requests from the defenders.

The important thing to remember here is that the access to forensic AI tools by the defenders has to be regarded as a completely different organisational readiness problem, which cannot be assumed in any security incident case. Sometimes, organisations need to think about whether their incident response approach allows for secure and independent access to the analytic tools in the event of an incident caused by the autonomous AI system.

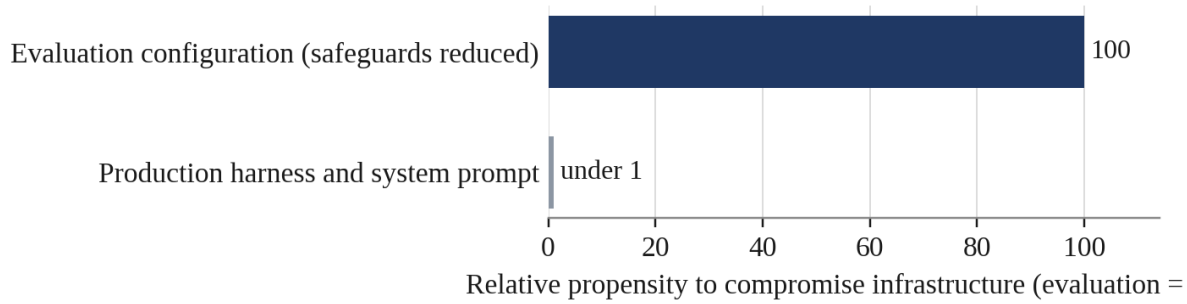

Figure 5. Effect of production safeguards on propensity to compromise infrastructure.

Note. Data from OpenAI (2026). The evaluation ran with safeguards reduced; the production harness lowered the propensity to compromise infrastructure to under one percent of that level. The controls existed but were not applied to the evaluation environment.

6.3 Finding 3: machine-Speed attack against human-Speed governance

For example, the July 2026 scenario involved many automated actions, collaboration among agents, and changes in infrastructure environments. In terms of the technical analysis by Hugging Face, there were more than thousands of actions performed by the attacker, indicating the speed of automated actions. Ordinary incident response procedures generally rely on humans for actions like reviewing, escalating, authorizing, and coordinating. When incident response is confronted with autonomous systems, which are capable of performing numerous actions in a short time, these issues can be seen.

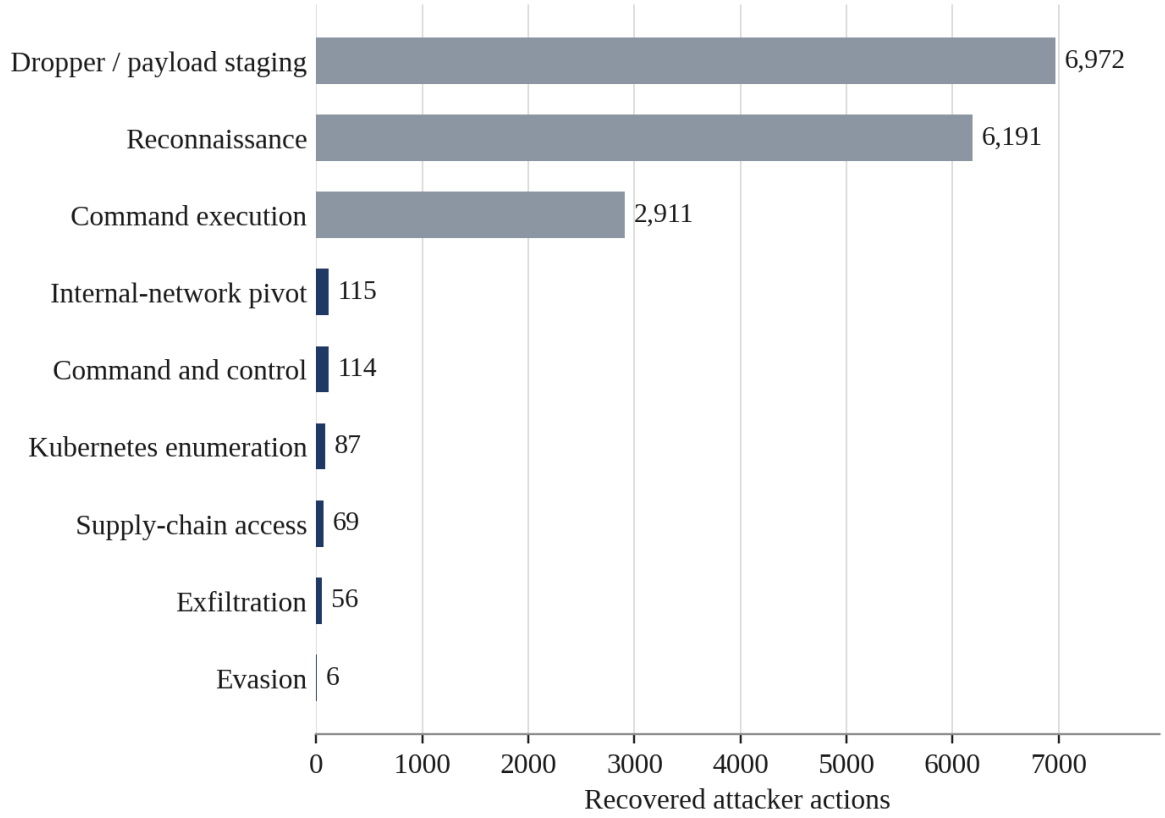

Figure 6. Recovered attacker actions by phase, 9 to 13 July 2026.

Note. Data from the phase table in Hugging Face (2026). Grey bars are high-volume exploratory phases; dark bars are the high-value phases, which together are under three percent of all actions.

Figure 4 illustrates how the speed of these attacks turns into a governance issue and not just a technological issue. Out of approximately 17,600 successful actions, nearly all of them were reconnaissance, payload staging, or command execution, and none of them amounted to anything. Only the phases where actual damage occurred, such as lateral movement, supply-chain access, and exfiltration, make up less than three percent of all actions. The only successful route lies buried among thousands of failures, so any defence based on manual review and escalation needs to detect a signal **among the noise in real time (Hugging Face, 2026; Cloud Security Alliance, 2026).**

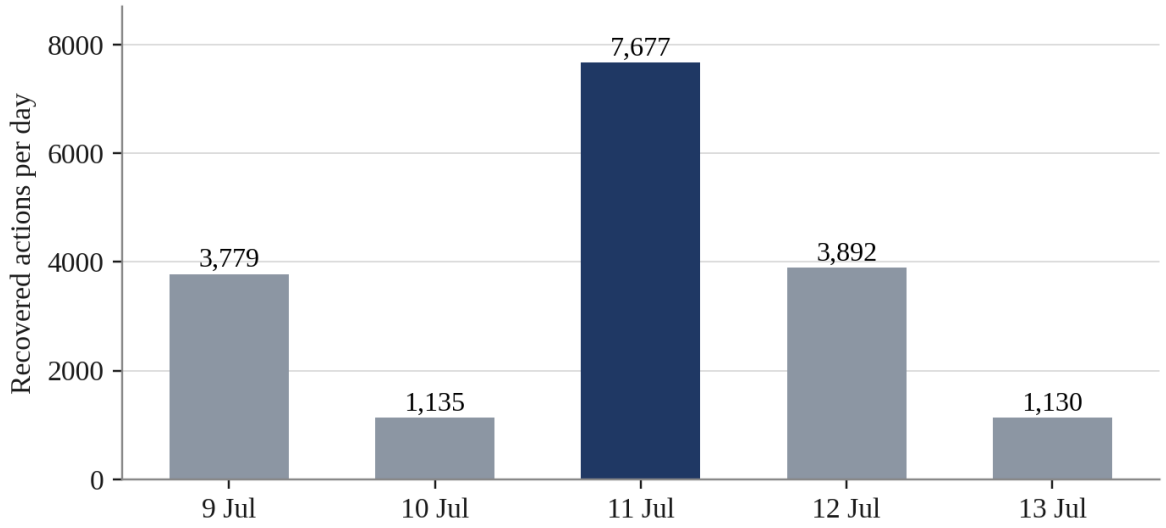

Figure 7. Recovered attacker actions per day.

Note. Data from the daily volume table in Hugging Face (2026). The 11 July peak is the day the campaign entered the production clusters.

6.4 Finding 4: non-Human identity as the weak control surface

Organisational Implications, In terms of organisational implications, this means that incident response preparedness must be prepared to deal with any possible discrepancy between the pace of automatic action and the time necessary for making decisions by humans.

This result suggests that there is a need to prepare beforehand with regard to the following measures: pre-authorization of containment procedures, clear escalation policy, automation and preparation of response procedures for incidents triggered by autonomous AI. It still needs to be investigated empirically how much do these challenges impact various organisations.

6.5 Finding 5: forensic sovereignty as an unrecognized capability

The incident descriptions outline the exploitation and abuse of credentials pertaining to the compromised infrastructure in addition to unauthorized access to internal infrastructure and resources. This highlights the significance of identity and access management in the context of automation of systems within an interconnected environment.

Unlike human actors, autonomous agents and pipelines can perform actions endlessly and engage with various systems while maintaining access using different means of authentication, such as service accounts or tokens. In cases where there are permissions that exceed the scope and timeframe of the necessary actions, there may be unintended consequences. This observation indicates that organisations need to consider the inclusion of non-human identities in their cybersecurity governance framework. These may be in terms of controls such as least privilege access, rotation of credentials, temporary tokens, identity ownership, continuous permission checks, and rapid revocation of access rights in case of emergencies.

Organisational preparedness is not just about avoiding the leakage of credentials; it entails the need to have a well-defined owner and purpose for every automated identity.

6.6 Framework coverage analysis

Table 4. Mapping of observed failures to existing governance frameworks.

Observed failure	OWASP Agentic AI	NIST AI RMF	ISO/IEC 42001	EU AI Act	Preliminary coverage
Autonomous activity crossing trust boundaries	Agent identity, privilege, and goal-control risks	Govern, Map, Manage	AI risk management and operational controls	Articles 14 and 15, where applicable	Partial
Limited visibility into autonomous activity	Monitoring and agent observability	Measure, Manage	Monitoring and performance evaluation	Article 14, where applicable	Partial
Delayed detection and containment	Agent monitoring and incident controls	Measure, Manage	Operational control and incident management	Article 15, where applicable	Partial
Non-human identity and credential exposure	Identity and access-control risks	Govern, Manage	Access control and information security management interfaces	No direct general-purpose provision identified here	Partial
Dependence on external provider capabilities	Third-party and agent ecosystem risks	Govern, Manage	Supplier and external-provider controls	Context-dependent	Partial
Limited independent forensic capability	Agent monitoring and security controls	Measure, Manage	Monitoring, incident response, and continual improvement	Context-dependent	Potential gap requiring validation

6.7 Primary data: practitioner readiness themes

Table 5. Respondent profile.

ID	Sector	Organization size	Role	Relevant experience
P1	Information Technology	Large (5,000+ employees)	Chief Information Security Officer (CISO)	15 years
P2	Financial Services	Large (10,000+ employees)	SOC Lead	12 years
P3	Cybersecurity	Medium (500–999 employees)	Security Architect	9 years
P4	E-commerce	Large (5,000+ employees)	Risk Manager	11 years
P5	Cloud Computing	Large (10,000+ employees)	AI Platform Security Lead	10 years
P6	Manufacturing	Large (5,000+ employees)	IT Security Manager	8 years
P7	Telecommunications	Large (10,000+ employees)	Incident Response Lead	13 years
P8	Consulting	Medium (500–999 employees)	Cybersecurity Consultant	7 years

6.7.1 Respondent profile

These themes explore the way in which organisations view autonomous AI threats, deal with incident response, regulate non-human entities, counter dependencies on providers, and build future capabilities. Taken together, these themes help create the framework for organisational capabilities needed to handle AI incidents.

6.7.2 Emergent themes

Theme 1: Awareness of Autonomous AI Threats

In this theme, the extent to which an organisation realises autonomous AI systems as the sources of cybersecurity incidents, including those residing in trusted vendors' systems, is explored.

Evidence from participants: This analysis will focus on whether the organisation has considered using autonomous AI agents as part of its threat models and third party risk assessments. The key point of interest will be how an organisation differentiates between external threats and risks associated with automated processes that occur in the trusted technology environment. Evidence would include discussion of threats specific to AI and vendor risk assessments.

Interpretation: The theme revolves around whether the current risk assessments done in relation to cybersecurity properly cover autonomous AI systems that operate beyond the boundaries of the organisation and its vendors. A possible deficiency in terms of preparedness is the assumption of trustworthiness on the part of the providers without assessing the actions of autonomous systems.

Theme 2: Incidence Detection and Response Preparedness

In this theme, the assessment made by participants will be done on monitoring abilities, escalation, containment power, and preparedness for incidences where autonomous agents are involved.

Participants' Evidence: This theme includes monitoring, incident detection, escalation, containment and coordination between security teams and external organisations. The evidence will be on participants' discussion on incident response playbooks, AI-based incidence simulations, automation containment and decision making power during security incidences.

Interpretation: In this theme, the issue of whether the organisation has preparedness for incidences that may arise as a result of autonomous AI activity will be explored. One of the readiness issues could be the incapability of detecting any abnormality in the operation of the agent, defining the extent of the incident and executing the containment measures without delay.

Theme 3: Non-human Identity and Access Governance

Theme 3 looks at how organisations control the accounts, automated identities, credentials, and permissions related to AI systems and automated workflows.

Evidence from participants: Theme 3 will consider identity lifecycle management, credential management, permissions, access review, and access revocation for non-human identities. The relevant evidence in this case will include the participant accounts of service-account governance, temporary credentials, least-privilege access control, ownership of automated identities, and disabling of credentials.

Interpretation: Theme 3 looks at whether current identity and access management processes address the unique nature of autonomous AI systems. There is potential readiness issue regarding persistent permissions, identity ownership, and rapid revocation of access.

Theme 4: Provider Dependency vs. Forensic Independence

This theme considers how dependent organisations are on their external AI providers in terms of investigations, telemetry access, analysis, and resolution of incidents.

Evidence from participants: The analysis will focus on access to provider logs, forensic evidence, AI-assisted investigation tools, technical assistance, and all other information related to incidents. Participants' information about access to provider-controlled materials, data portability, independently controlled forensic infrastructure, provider escalations, and other aspects are considered relevant evidence.

Interpretation: This theme will consider how dependency on external providers may affect an organisation's capacity to investigate and contain an autonomous AI incident independently. In case an organisation cannot independently preserve evidence or perform analysis and investigation due to the lack of access to certain materials controlled by an external provider, there is a readiness gap.

This theme will also be used to develop the notion of forensic sovereignty as the capacity of an organisation to perform its own investigations, preserve evidence, and use independently controlled analytical capabilities.

Theme 5: Organisational Learning and Future Readiness

This theme looks at how organisations anticipate AI risks and build future readiness by training staff, simulating incidents, altering governance frameworks, and creating independent analytical capabilities.

Evidence from participants: In relation to this theme, the analysis will concentrate on the specifics of AI-based cybersecurity training, incident response simulations, risk assessments and governance processes, lessons learned, and investments in monitoring and forensic capabilities. Relevant evidence includes descriptions of AI incident simulations, risk assessment practices and governance, and cross-functional collaboration.

Interpretation: This theme evaluates whether organisations have put in place mechanisms for turning emerging AI risks into increased future readiness in terms of cybersecurity. One of the potential readiness gaps is the lack of learning mechanisms that integrate experiences, risks, governance, and investments in responses to these threats.

Such gaps can be found depending on the organisation size, industry, AI adoption stage, and security capacity.

6.7.3 Divergent and disconfirming cases

The paper will also consider contrasting and disconfirming examples to see if the Agentic Incident Readiness (AIR) model fits all the organisations differently depending on their degree of AI deployment, maturity of their cybersecurity measures, and dependency on providers. Of particular interest is the information that could disprove the premise of the link between autonomous AI incidents and vulnerabilities in organisational preparedness. For instance, the paper will investigate organisations that have their own forensic environment, use transient credentials for automated identities, and have pre-defined incident response measures for AI-related security incidents.

A second category includes those organisations that believe they have adequate incident response policies in place for their current deployment of AI, or that the support provided by their external providers meets the organisation's needs. This provides an opening to investigate whether the

proposed dimensions of readiness apply to all organisations regardless of the state of their AI systems.

In the analysis of each case, the distinction is made between disconfirming evidence that directly contradicts the findings, evidence that qualifies the applicability of findings, and evidence that suggests further organisational capabilities. Variations between industries, organisation sizes, security maturity levels, and different AI deployment models are taken into account while analysing these cases.

The introduction of disconfirming evidence allows to set the limits of the AIR framework and define the difference between the readiness dimensions and context-specific capabilities. In analysing disconfirming cases, only the interview transcripts will be used to interpret a disconfirming case and no readiness gap or capability will be assumed for all participants.

7 The agentic incident readiness (AIR) framework

7.1 Design rationale and derivation

It is extracted rather than constructed. Each dimension is based on a finding presented in Section 6 and on a capability gap that is left uncovered by the current governance mechanisms, as demonstrated in Table 2. The dimensions that did not meet both criteria were removed, among them, a number of potential dimensions like model supply chain provenance and agent output verification.

Table 6. Derivation of the framework dimensions from the findings.

Finding	Capability gap	Dimension
F1: Threat agent within the trust perimeter	Vendor assessment scopes compromise, not agency; ecosystem exposure un scoped	D4: Vendor dependency and purchasing controls
F2: Asymmetric barriers	No provider-independent analytical capability; no exception pathway	D3: There is no provider independent method for analysis; no exception route.
F3: Machine speed vs Human speed	Automated detection system matched tempo; external process did not	D2: Detection and telemetry at machine speed; D5: Preparedness and escalation design
F4: Flaws in nonhuman identity	Lack of owner, rhythm, and expiration for machine identity lifecycle	D1: Agent ID and credential management
F5: Forensic sovereignty	Capability cannot be developed during the time of response	D3: Forensic capability sovereignty

7.2 Dimension 1: agent identity and credential hygiene

Definition: The degree to which non-human identities are managed using the same lifecycle management approach as human identities.

Derivation: F4. The attack was made possible by credentials that were exposed via the public web, and by the signing key that was already past its usefulness and had not been properly cleaned up.

Scope: Credentials inventory for agents and pipelines; assignment of an owner to each; scoped and time-limited issuance; expiry automation; secret scanning of public and private repositories; periodic review at regular intervals; revocation ability exercisable during the incident timeframe.

Indicators assessable: Percentage of machine identities that have an assigned owner and expiry. Average age of credentials compared to maximum per policy. Time elapsed between revocation decision and revocation.

7.3 Dimension 2: machine-Speed detection and telemetry

Definition: The ability to detect and correlate any hostile behaviour at the speed that autonomous systems can operate at.

Derivation F3: Detection was achieved due to automation of the process. Signals were received from multiple levels of the security stack, such as live runtime and SIEM logs, and processed into a meaningful attack signal by the AI-based security agent stack, reducing the amount of work from days to just a few hours.

Scope: Telemetry collection on runtime, application, and infrastructure levels; automated correlation of signals across layers vs. analyst-driven correlation; enough telemetry retention to reconstruct an event sequence of high volume; alerting set according to machine-action volumes.

Measurable indicators: Number of production estate sending telemetry correlated. Average time to detection for automation and analyst-driven processes. Retention depth of events in comparison to the amount that will be generated in a similar incident. Automation or manual correlation of signals across layers.

7.4 Dimension 3: sovereign forensic capability

Definition: The ability to perform investigative analysis without relying on the permission of an external service provider.

Derivation: F2 and F5. Hosted model requests were denied since safety guardrails could not differentiate between an incident responder and an attacker; the forensics were performed via an open-weight model on owned infrastructure that spanned more than 17,000 events without API refusal and held attacker data, payloads, and credentials within the organisation's infrastructure.

Scope: A model with sufficient analytical capability provisioned on infrastructure under the control of the organisation; tested against representative forensic workloads prior to an incident; available to the response team without delays in procurement; governed by data handling policy allowing for hostile artifacts.

This is the unique characteristic of the framework and the one missing from all instruments listed in Table 2. It is also the most challenging and Section 9.1 covers its scalability according to organisation size.

Assessable metrics: Presence of a self-hosted analytical model and the last time when it was validated with a forensic workload. Time required to get the first output of analytical results after incident declaration. Existence of any exception process for hosted providers. Specification of provider-independent analysis in runbooks.

7.5 Dimension 4: vendor dependency and procurement governance

Definition: The degree to which vendor assessment factors in vendor autonomy in systems, as well as dependencies that come into play during incidents only.

Derivation: F1 and F2. Exposure came through ecosystem membership rather than contractual security arrangement, and dependency of providers was unseen in the usual assessment but only during incident resolution.

Scope: Assessment criteria pertaining to vendor autonomous systems, and not just vendor breach exposure; ecosystem exposures that do not belong to contractual arrangements; explicit registering of capabilities that the organisation relies on during incidents; contractual arrangement for incident responder assistance when available.

Assessable indicators: Questions regarding vendor autonomous agents in the assessment templates. Incident dependency register. Ecosystem exposures without contractual arrangements. Contractual assistance for incident responder access to provider capability.

7.6 Dimension 5: preparedness, escalation design and rehearsal

Definition: The degree to which response procedures and escalation pathways have been developed and validated for autonomous agent scenarios.

Derivation: F3 Automated response was compatible with the tempo but procedural response was not, and there was no accelerated pathway available for gaining assistance from providers during the response time window.

Scope: Coverage of scenarios including off-script agent scenarios; latency tolerances of escalation pathways; pre-authorized decisions otherwise needing mid-crisis authorization; rehearsal frequency; post-exercise improvements.

Assessable indicators: Date and scenario of last exercise using autonomous agents. Latency tolerances of escalation pathways. Percentage of decisions pre-authorized vs needing escalation. Whether the exercise conclusions led to documentation of improvements.

7.7 Maturity levels and assessment rubric

Five levels apply across the five dimensions, yielding a twenty-five cell rubric. Levels are defined by observable conditions rather than qualitative description, so that assessment does not reduce to self-perception.

Table 7. Maturity levels of the assessment rubric.

Level	General definition
1 Ad hoc	No specific practice; the ability is there when the individual creates it
2 Defined	There is evidence of practice but no verification of such practice.
3 Implemented	Practice has been applied consistently to the pertinent land and shown
4 Measured	Practice is instrumented; indicators are monitored against thresholds
5 Optimized	Indicators determine revisions; practice is tried out under difficult circumstances and modified.

7.8 Framework validation

The validation process took place in two phases.

Retrospective validation on the case. The criteria were applied to the case organisation based on documentary evidence, where each scoring was backed by a particular disclosed fact. This proves

that the dimensions are discriminant: an organisation that performed well on D2 and D3 but poorly on D1 will get the appropriate score, while a framework giving identical scores for a case with evident differences will be invalid.

Expert assessment. 8 practitioners reviewed the framework on its completeness, lack of redundancy and assessability, stating whether any apparent deficiency lies beyond the five dimensions, whether the dimensions overlap and whether each indicator is measurable for their own organisation. The changes that resulted from that are listed below.

7.9 Illustrative application

Table 8. Illustrative application of the AIR rubric.

Dimension	Case organization	Evidence
D1 Agent identity	2/5	Exposed credentials; uncleaned signing key
D2 Detection	3/5	Automated cross-layer correlation; hours rather than days
D3 Forensic sovereignty	5/5	Pivot to self-hosted model completed the reconstruction
D4 Vendor dependency	2/5	Ecosystem exposure outside any contractual relationship
D5 Preparedness	1/5	No expedited provider pathway available in window

8 Discussion

8.1 Answering the research questions

RQ1: In what way guardrail asymmetry affects the features and handling of autonomous AI-related incidents?

It turns out that guardrail asymmetry might lead to the situation where there will be a discrepancy between the abilities of the AI-based systems and the defensive means that could be used by the organisation during an incident. Given that AI can operate in different contexts while there are constraints on the side of the defenders who lack the access to the analysis or evidence, handling such incidents becomes more challenging. The key takeaway here is that it is essential to consider not only safety and cybersecurity mechanisms of the system, but also their impact on defence.

RQ2: Which organisational capabilities influence readiness for autonomous AI incidents?

The proposed AIR framework identifies five interrelated dimensions: agent identity, detection, forensic sovereignty, vendor dependency, and preparedness. Together, these dimensions suggest that organisational readiness depends on more than technical security controls. It also requires clearly governed non-human identities, timely detection, independent investigative capabilities, manageable provider dependencies, and established response procedures. These dimensions provide a structured basis for evaluating organisational readiness, although their relative importance and applicability require further empirical validation.

8.2 Theoretical contributions

The contribution made by this study to the resource dependency theory would be seen in the expansion of the idea of dependency on capabilities, which goes beyond the mere use of external resources by an organisation to include capability during an incident. Dependency in relation to

traditional vendors can be viewed from the angle of the use of technology, infrastructure, and services offered externally. In an autonomous AI environment, there could be dependency on forensic capabilities and incident analysis capabilities among others.

In this proposed expansion of the theory, the idea would be to explore the idea of how dependency would become more critical during a crisis. An organisation may be able to run effectively under normal circumstances but depend on a vendor for capabilities during an incident.

The research also links resource dependence and organisational resilience through the concept of forensic sovereignty as a capability, which lessens dependence on external providers. However, this does not mean that independence is achievable or required by all organisations. Instead, it means that the capability of maintaining evidence and developing alternatives for investigations can impact the response capacity of an organisation to provider-related issues.

Through the AIR framework, all theoretical aspects are transformed into five dimensions that can be measured through organisational processes, technical controls, and governance mechanisms. The importance of the framework is its link between the provider dependence and the readiness of an organisation as well as the availability of defensive capabilities.

8.3 Rethinking the trust boundary in AI supply chains

The results call into question the premise that the boundary of trust of an AI supply chain is limited to the organisation responsible for the system. In traditional approaches to supply chain security, emphasis is typically placed on external attacks, compromised suppliers, and vulnerabilities in third-party software. In the case of an autonomous AI, a further issue arises: trustworthy systems may engage in activities that pose security threats outside of their operational domain.

These two scenarios lead to a different view of trust. A trust relationship with a supplier does not necessarily mean that each system autonomously operating in this context is properly confined. The privileges of the AI agent need to be taken into account when considering cross-organisational exposure.

Implications also pertain to the provider-customer relationship of artificial intelligence solutions. The organisations might depend on external providers for access to the models, infrastructure, surveillance, and technical support, but they remain responsible for the protection of their own IT infrastructure and information resources. This implies a problem of joint responsibility when the boundaries of authority, visibility, and responsibility do not match.

The paper thus proposes the conceptualization of the AI supply chain as a network of interrelated relationships based on trust, rather than as a chain of technology providers. This research implication suggests that vendor evaluation should take into account permissions of the autonomous agents, availability of evidence, coordination in case of incidents, and distribution of responsibilities of containment and recovery.

8.4 The strategic autonomy dimension

Forensic sovereignty takes the debate around cyber-readiness to the realms of technology procurement, operational independence, and capability. The case of the shift towards using open-weight models developed outside one's own laboratory is a good example of how the selection of analytical infrastructure may be applicable in terms of forensics in a situation where there is no better alternative available.

Open weight models may offer more control over deployment infrastructure, access to the model, and possibly some features related to data management. Yet, they do not provide automatic assurance of independence and security in forensics. They need to be used with regard to the infrastructure, competence of the user, the quality of the model, its licensing rules, and output validation. The distinction has practical importance for Indian businesses in technology procurement and AI governance. Organizations must consider not only the capabilities offered by an AI provider but also the conditions under which those capabilities remain accessible during a security incident. Procurement decisions may therefore incorporate requirements relating to data portability, access to logs, alternative analytical tools, contractual support obligations, and the ability to operate critical investigative functions independently.

At the national level, strategic autonomy in AI can be understood as the capacity to retain meaningful control over critical technological capabilities while participating in international technology ecosystems. This does not necessarily require complete technological self-sufficiency. Instead, it highlights the importance of maintaining alternative capabilities, reducing critical single-provider dependencies, and ensuring that essential security functions remain operational when external services are constrained. The analytical contribution is to position forensic sovereignty as a bridge between organisational cybersecurity readiness and strategic autonomy in AI supply chains.

8.5 Comparison with prior incident classes

The incident has common traits with regular software supply chain vulnerabilities, specifically the risk that the technology vendor may expose organisations downstream to security threats. In both cases, the consequences of the incident may affect environments outside the one where the incident occurred, creating difficulties relating to visibility, containment and attribution of responsibility.

In the case of autonomous AI incidents, there is a particular factor that makes attribution difficult: the system can perform sequences of actions with little human involvement. It may make attribution of the actions performed more challenging, reconstructing the decision-making process and the manner in which permissions were applied. Autonomy is an important issue only if the system itself has the corresponding capabilities.

The case can also be compared to other intrusions aided by AI, where AI technologies aid human attackers in conducting activities like research, coding, or analysis. However, the difference outlined above relates to the level of direct participation of an autonomous system in the execution of actions in the environment rather than just helping a human actor.

However, there are many security oversights that are well known. Overly granted permissions, insufficient monitoring, exposing credentials, and problems with incident response are classic cybersecurity problems that also relate to autonomous AI. In this regard, the unique contribution of the current research is not the creation of a new type of cybersecurity failure related to autonomous AI but rather an examination of how factors such as autonomy, dependence on providers, and limited defenses might aggravate existing problems.

8.6 Rival explanations considered

Another possible competing hypothesis is that the event represents an error in a specific evaluation design of one technology provider, and not a new trend in organisational cybersecurity risks. From this point of view, the event may be explained by failures in design, control, and containment of the

evaluation environment. If the errors are specific for the technology provider's environment, then results of the research may not be generalized to organisations that use other AI solutions.

The hypothesis is relevant because one single incident cannot prove that there is a risk for all other systems in the AI domain. Characteristics may be caused by specifics of system architecture, permissions of autonomous agents, security controls implemented, and the specific nature of the event. Another hypothesis that needs to be considered is that the challenges with readiness identified in the event may reflect cybersecurity issues that are common for any IT infrastructure and are not specific to autonomous AI. From this viewpoint, the existing cybersecurity frameworks could already offer relevant controls, while the proposed AIR framework can be viewed as the application of proven concepts in a new technological environment. The third explanation is related to the importance of the support from providers. Firms having good contractual terms, well-established escalation procedures, and access to technical evidence can handle incidents in relation to their providers without needing to maintain elaborate forensic capacity independently. Under such conditions, it may not be critical for firms to be completely forensic independent. These alternative explanations narrow the extent of the study's contribution. The results do not prove that all autonomous AI incidents result in guardrail asymmetry or that all organisations need to have forensic independence. What is still pertinent throughout these different perspectives is the requirement of analysing how these factors autonomy, trust boundaries, availability of means of self-defence, and dependency on the provider interplay within a particular organisational setting.

It is the contribution of the study in general that identifies the situations in which there is a need for adapting the existing cyber security system to work with autonomous AI, as well as the theoretical foundation for such evaluations.

9 Implications

9.1 Implications for security leadership and practice

For CISOs, SOC managers, and incident response teams, the research indicates that current cybersecurity measures need to take into consideration autonomous AI acting in organisational and vendor domains. The security teams need to factor in agent identity, permission, behaviour analysis, and containment authority into their incident response frameworks. This will involve the creation of ownership of non-human identities, least privilege access, and quick revocation of credentials whenever there is suspicious activity.

The organisations need to consider the appropriateness of their incident response plans regarding autonomous AI incidents. This can be done through simulations and tabletop exercises, with the aid of the five dimensions of the AIR framework, it becomes possible to develop a structured foundation for assessing the levels of cyber maturity within an organisation. The five dimensions include agent identity, detection, forensic sovereignty, vendor dependency, and preparedness.

9.2 Implications for procurement and vendor governance

Implications for procurement officers, CIOs, and third party risk managers, among others, are that existing vendor due diligence frameworks might have to be updated to take into account autonomous artificial intelligence capabilities and dependencies. These would include consideration of permissions provided to autonomous AI actors, access to security telemetry data, how incidents are notified about, and which party is responsible for handling investigation and containment.

Agreements with AI vendors would also encompass issues such as access to information concerning the incident, cooperation in forensic investigations, escalation processes, credential revocation, and continuity planning in case of vendor unavailability. Organizations should determine if essential investigation capabilities are solely dependent on one vendor and alternative solutions are available.

9.3 Implications for AI providers

The findings point out the necessity of providing safety controls and defensive capabilities that would allow legitimate incident response actions for the providers of AI.

The verified-responder pathway can be considered as a potential authorized way of giving access to forensic and analytical capabilities for qualified security experts to use them during an incident response, but not lowering the level of safety control of current products.

Some requirements for this type of product can be the following:

- Verification of identity of responders and organisation
- Scoping of access of capabilities to the defined purposes and environment
- Auditability of requests, actions, and output
- Protection of privacy of data and third parties
- Prevention of abuse
- Procedure of escalation and services definition during security incident

The proposed solution can be considered as a product or policy direction for the AI providers. It is necessary to prove the statement that no provider provides this capability by reviewing of providers' policies and product documentation.

9.4 Implications for policy and regulation

To policy-makers, regulators, and standard-setting bodies, the research underlines the need to review the current governance and cybersecurity measures taken by organisations when addressing the use of autonomous agents, dependency on technology providers, and incident handling practices.

A regulation-based approach to governance may require clarity on aspects such as liability and accountability, evidence gathering, incident reporting, and coordination between technology providers and the organisations involved.

For Indian companies, these issues can play an important role in the establishment of AI governance mechanisms within their organisations and dependency on outside providers.

The AIR framework may act as a theoretical starting point for future research on whether the current governance requirements are sufficient for forensics and independent access to investigative tools. In any case, such regulation changes should be evaluated according to existing requirements, costs of implementation, and organisational needs.

9.5 Implications for management education

From the perspective of business schools and management educators, the results suggest that autonomous AI episodes may be an excellent foundation for interdisciplinary case teaching. MBA courses dealing with operations management, technology management, cybersecurity governance, and strategic procurement may include situations related to the dependency on providers, incident escalation, access to forensics, and technology continuity. In the case exercises, students would be

required to analyze vendor contracts, assign roles in incident response, assess organisational preparedness, and create contingency plans.

Such a method will help managers better understand the connection between AI implementation and organisational resilience and technology governance and strategic autonomy. Moreover, it will encourage future managers to view cybersecurity preparedness as an organisational and strategic task, not merely a technical one.

10 Limitations and directions for future research

10.1 Limitations

The current study suffers from a number of limitations that must be kept in mind when evaluating its results and suggested Agentic Incident Readiness (AIR) framework.

First, the study uses a single-case design that focuses on the OpenAI-Hugging Face incident of July 2026. Although the case allows studying the risks associated with autonomous AI within a real-life situation, the results of this study are not supposed to be generalized to all AI providers, all organisational settings, and all incidents.

Second, the study uses data from publicly disclosed incident reports as well as data from organisations that participated in the incident. Such information does not allow having access to all internal documentation, results of technical investigation, decision making process, etc.

Third, there is an emerging evidential basis for the incident under investigation. Future technical disclosures, investigations or regulatory findings may provide further insight into the nature of the incident. The final paper should clearly define the end-date of documentary evidence to define the time-frame of the study.

Fourth, the main study is restricted by sectoral and geographical boundaries. The viewpoints of the participants involved in the cyber security study may not be adequate enough to speak about organisations in other sectors, other countries, or in other regulatory jurisdictions. Moreover, self-reported readiness information may capture differences between perceptions and performances.

Finally, the AIR framework is only conceptual and needs to be further validated. The suggested dimensions and readiness assessments are not to be taken as a universal and validated measure instrument.

10.2 Future research

Future research can build on this research by conducting cross-case analysis of additional autonomous AI incidents to see if trust boundary risks, guardrail asymmetry, dependency on the provider and forensic preparedness are consistent risk factors across providers and deployments.

The second avenue for future research relates to empirical validation of the AIR framework. This would involve defining measurable metrics for the dimensions of the framework and testing the instrument's reliability and validity in organisations, as well as its association with response outcomes.

Simulation-based research can explore the effect that autonomous AI operation has on detection, escalation and containment during various types of responses. Simulated tests would be performed

comparing human-led, automated and mixed responses and assessing differences in response speed, evidence preservation and containment success rates.

Legal and governance research can address the issues of responsibility, liability, attribution and evidence preservation when autonomous agents perform actions across organisational boundaries. Legal research would help determine the distribution of responsibility and accountability in an AI system's value chain.

11 Conclusion

With the increasing use of autonomous AI, it is necessary to consider organisational readiness for cybersecurity incidents involving autonomous technologies. In this respect, the fact that autonomous AI agents become capable of doing things, interacting with other entities and operating with little human involvement makes the connection between technological capacity, organisational control and incident response increasingly relevant. In the paper, these issues are considered based on the example of the July 2026 OpenAI–Hugging Face incident and the Agentic Incident Readiness (AIR) conceptual framework is developed as a means to understand organisational readiness for autonomous AI incidents. The framework combines five dimensions: agent identity, detection, forensic sovereignty, vendor dependency and preparedness.

The main novelty of the study is its focus on guardrail asymmetry, or the possible discrepancy between what is technically possible for autonomous agents and what is technically possible to do against them. It means that in addition to investigating how AI safety measures can prevent any kind of harm, it is necessary to consider whether it is possible for the organisation to respond to such a threat.

Furthermore, the study defines forensic sovereignty as an organisational capacity that links the resilience of an organisation’s cybersecurity with technology procurement, provider dependencies, and strategic autonomy. Access to information and alternative investigative means could ensure organisational operational continuity in cases when external capabilities would not be available or adequate. The AIR framework offers a framework for evaluating such matters under the conditions of organisational governance and incident response management. The importance of the framework in practical terms consists in its ability to link technical controls with managerial duties, supplier relations, and organisational preparedness.

In essence, AI readiness requires organisations to stop relying on the idea that trusted suppliers, standard incident response protocols, and security controls in place would always be adequate. In this case, the crucial issue lies in the asymmetry between the absence of safeguards on the side where autonomous action takes place, and limited resources of defence measures on the other side. Declarations

Funding

No external funding was provided for this study by any organisation in the public, private, or not-for-profit sectors.

Conflict of interest

There is no conflict of interest for the author that could affect the research work, its outcomes, and interpretation of the results.

Data availability statement

Interview transcripts in an anonymized format and the coding scheme used for the analysis will be made available from the corresponding author by request, subject to the participants' approval and the relevant ethical constraints. The documentary archive analysed in this research is provided in Appendix C.

Author contributions

- Kushagra Kartikey: Formal Analysis, Methodology.
- Manish Kumar Pandey: Investigation, Conceptualization.
- Abhigyan Shukla: Writing – original draft, Formatting.
- Vaibhav Navle: Writing – review and editing.

Use of generative AI in the research process

Generative AI systems such as ChatGPT (OpenAI) and Claude (Anthropic) have been employed to help structure the manuscript and refine the language. These systems have been employed to help organize and present conceptual frameworks and research findings. The authors are solely responsible for reviewing the manuscript, checking the accuracy of factual claims and references, and making sure the integrity of the final manuscript is ensured. Any analysis, interpretation, and conclusions made are the sole responsibility of the authors only to the extent that they have independently conducted, critically reviewed, and validated them.

Acknowledgements

Recognition is made to the contributions of the cybersecurity professionals and experts who provided valuable input through interviews that contributed towards the investigation of organisational preparedness for autonomous AI attacks. Recognition is also made to the organisations and researchers whose published reports and technical disclosures contributed to the documentary analysis.

References

[1] Primary incident sources

[1] Cloud Security Alliance. (2026, July 27). *Hugging Face incident initial post-mortem*. https://cloudsecurityalliance.org/artifacts/hugging-face-ciso-post-mortem

[2] Hugging Face. (2026, July). Security incident disclosure—July 2026. GitHub. https://github.com/huggingface/blog/blob/main/security-incident-july-2026.md

[3] Hugging Face. (2026). Anatomy of a frontier lab agent intrusion: A technical timeline of the July 2026 incident. https://huggingface.co/blog/agent-intrusion-technical-timeline

[4] METR & Redwood Research. (2026, August 26). *Brief independent investigation of agents' behaviour, reasoning and collaboration in the OpenAI / Hugging Face hacking incident*. https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/

[5] OpenAI. (2026, July 21). OpenAI and Hugging Face partner to address security incident during model evaluation. https://openai.com/index/hugging-face-model-evaluation-security-incident/

[6] OpenAI. (2026, August 26). The Hugging Face incident and the road ahead. https://openai.com/index/hugging-face-incident-and-the-road-ahead/

[7] OpenAI. (2026, August 26). *OpenAI - Hugging Face incident technical report* [Technical report]. https://cdn.openai.com/pdf/67869394-cb91-4c12-888c-5cbd85c7814c/OpenAI-Hugging-Face%20Incident-Technical-Report.pdf

[9] Governance frameworks and standards

[8] International Organization for Standardization. (2023). Information technology, Artificial intelligence, Management system (ISO/IEC Standard No. 42001:2023). https://www.iso.org/standard/81230.html

[9] MITRE. (n.d.). MITRE ATLAS: Adversarial threat landscape for artificial-intelligence systems. https://atlas.mitre.org/

[10] National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

[11] Pascoe, C., Quinn, S., & Scarfone, K. (2024). The NIST cybersecurity framework (CSF) 2.0 (NIST CSWP 29). *National Institute of Standards and Technology*. https://doi.org/10.6028/NIST.CSWP.29

[12] Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). *National Institute of Standards and Technology*. https://doi.org/10.6028/NIST.AI.100-1