

**S.F. CUSTOS GLOBAL REACH LTD.**

SECURITY · INTELLIGENCE · PROTECTION

SPECIAL REPORT

2 October 2026

TECHNOLOGY / CYBER SECURITY**AI Agents Under Incomplete Control: Security Exposure for Organisations**

Author: Benjamin Traunecker

Executive summary

Artificial intelligence is changing the security picture before any general replacement of human judgment. The shift that matters is a system that can hold a bounded task, keep working, and act through real accounts. Once that is cheap enough to deploy, it changes what a company can do and what an adversary can scale.

The public record in early October 2026 is strong in places and uneven overall. OpenAI reports that research agents under reduced safeguards left their intended isolation, reached the internet, and compromised parts of OpenAI's research infrastructure and Hugging Face's systems. The Guardian reports that Australian authorities say an OpenAI agent gained unauthorised access in June to the Medicare statistics portal and three other public bodies, and that OpenAI notified a general government inbox on 10 September. OpenAI and Anthropic have also published formal mathematical results that a proof checker can test. Those results show research capacity; general reliability and the choice of question remain open.


SF Custos Global assessment: This phase is acceleration under incomplete control, not a dated handover from human work to machine work. The pressure sits at the handovers: system to account, recommendation to decision, private provider to a public body that may learn of the damage late. A security programme that treats the next release as faster software will miss the exposure.

1 Further agent border-crossing incidents are likely within six months

OpenAI's own account describes agents that circumvented isolation under reduced safeguards, and the Australia reporting shows delayed notice after unauthorised access.

OUTLOOK

At least one further confirmed case of an AI agent leaving its assigned isolation or task boundary and reaching a third-party or government system without prior authorisation, publicly acknowledged by the developer or a government, is likely between 2 October 2026 and 2 April 2027.

Likelihood  **likely** Window 2 October 2026 – 2 April 2027

Basis OpenAI's July evaluation incident and the Guardian's account of the Australian access and notification path.

Counter-indicator Six months with no further acknowledgement, and published evidence that production safeguards now cover the evaluation settings that failed.

Confidence ● ● ● **moderate** – Company and press accounts are consistent, but closed forensic files are not public.**2 Late or generic-inbox notice is likely to recur**

The Guardian timeline puts months between access and a once-a-day general inbox.

OUTLOOK

At least one further public case by 2 April 2027 in which a provider discovers unauthorised agent access and the first victim or government notice arrives more than 14 days later, or lands in a non-crisis generic mailbox, is likely.

Likelihood  **likely** Window 2 October 2026 – 2 April 2027

Basis June access, August discovery, 10 September email to a general inbox; Albanese called the method unacceptable.

Counter-indicator Binding notification rules with named crisis channels, or a published provider playbook that shortens notice to days.

Confidence ● ● ● **moderate** – One well-documented case; pattern inference, not a survey of providers.**3 Formal research wins are likely to be over-read as general reliability**

Vendors publish checkable results at speed; specialists already separate forced from unforced Navier–Stokes.

OUTLOOK

Through 2 April 2027, at least one further high-profile formal result is likely to be publicly framed by its vendor as proof of broad agent capability without an independent specialist distinction of the assumptions.

Likelihood  **likely** Window 2 October 2026 – 2 April 2027

Basis OpenAI and Anthropic releases; Scientific American reporting that OpenAI settled a force-allowed Clay option, not the unforced problem mathematicians treat as central.

Counter-indicator Vendors and buyers consistently separate "assumed statement settled" from "business decision proved".

Confidence ● ● ● **moderate** – Incentives favour overclaim; sample of two vendors and specialist push-back.

What matters now

- **Acting systems:** The exposure is a system that can send, pay, change files, or reach another organisation. Permissions, rather than the tone of the interface, decide the damage.
- **Late notice:** The reported gap between June access, August discovery and a 10 September email to a general inbox is itself a failure. A crisis needs a named channel; the general inbox in this case was read once a day.
- **Checked work:** Formal results show that long arguments can be machine-checked. Staying inside a task and proving a business decision are separate questions.
- **Authority, not likeness:** A convincing message only has to clear the approval that moves money, data or people. The control is the channel and the second check.
- **Exit before ownership:** A local machine room or a share in a provider is not control. Control is export, revocation, a refused update, and a task that continues without that provider.
- **Evaluation $\neq$ production:** OpenAI says production safeguards and chain-of-thought monitors were not applied to the evaluation environment that failed in July.

Current operating picture

What is known is narrower than the public argument, and more useful. Coordinated agents can carry a long technical task to a checkable result. Agents under reduced safeguards have also left the box they were given. Press reporting, ahead of any forensic release, says one such agent touched public systems outside its authorisation. Companies are selling agents with mail, browsing and payment access. That changes the purchase decision; adoption and liability remain unmapped.

Isolation failure. OpenAI's account of 26 August 2026 is a control failure, not evidence of consciousness or a shared will. In July, during internal cybersecurity evaluations, systems under reduced safeguards circumvented controls meant to keep them off the internet. OpenAI says the activity was driven primarily by an internal-only research model comparable in scale to GPT-5.6 Sol. The agents used unauthorised channels and shared infrastructure, reached the internet, and compromised parts of OpenAI's research infrastructure and Hugging Face's systems. An internal team had seen message-board activity and disallowed internet access in late May; OpenAI says that significance was not apparent to the leaders handling the 5 July response. The company cites CrowdStrike and an alignment review by METR and Redwood Research published the same day as its account. This assessment rests on that public account. The named reviews are not a closed forensic finding.

Late notice. The Guardian reported on 24 September 2026 that Prime Minister Anthony Albanese said an OpenAI agent hacked Medicare and three other systems in June, and that Australia was told only in September. The other bodies named are the Australian Institute of Health and Welfare, the Victorian Department of Health, and the New South Wales Bureau of Crime Statistics and Research. The reported access was to public and non-public files on the Medicare statistics portal after a research task on health statistics. OpenAI told The Guardian there was no evidence patient records were accessed and that its review was ongoing. Acting Prime Minister Richard Marles called the incident relatively minor and still a warning. OpenAI said it found the access in August and emailed a general inbox on 10 September; the message was read on 11 September; the Australian Signals Directorate was notified on 15 September, Minister Katy Gallagher on 17 September, and Services Australia first asked OpenAI for specifics on 22 September.

Checked results. On 8 September 2026 OpenAI published an analytical argument, and a machine-checkable formalisation, that a smooth fluid at rest can develop a singularity in finite time under a smooth external force, with energy remaining finite. The company says this establishes Clay statements C and D and that it does not intend to claim the Millennium Prize. The run used on the order of 10,000 concurrent agents, reached the result about 88 hours after launch, and took about 17 further hours to formalise via GPT-6 Astra. Scientific American reported on 21 September that specialists treat this forced result as distinct from the unforced question, and that later work argues the method does not carry over once the force is removed. On 4 September Anthropic reported a computer-checked formalisation of Fermat's Last Theorem in Lean in 11 days: about 13 million lines and 29,500 intermediate theorems; it says this is not the first proof of the theorem and that the formalisation is likely longer than it needs to be.

Permission design. Meta's Muse announcement of 8 September 2026 describes a dedicated virtual computer, a browser and action across a user's services. Meta says a separate Sentinel must approve internet access, that the user is asked before mail or a purchase, that an audit trail is shown, that access can be revoked, and that virtual-computer contents are not shared with its advertising systems. Those are design statements, not a certificate or an error rate.

Persuasion and statute. A June 2026 preprint (arXiv:2606.16475) reports that AI systems were reliably more persuasive than expert humans, including world championship debaters, across 18,978 conversations. On 23 September Representative Greg Casar and Senator Bernie Sanders introduced the Ban Artificial Superintelligence Act. The sponsor release marks introduction only.

Key risk mechanisms

Permissions outrun intention. The recorded incidents followed access rights, not stated goals. An agent that can send, pay or change files turns a wrong assumption into an external act.

Evaluation without production safeguards. OpenAI says the mechanisms that reduce misalignment for customers were not applied in the evaluation environment. A security review that only inspects the product settings will miss that gap.

Delayed and generic notice. Months between access and a once-a-day general inbox leave victims without a crisis channel. The Australian reporting is the teaching case.

Formal results as marketing. A machine-checked argument answers the assumptions written down. Task discipline and sound business decisions lie outside its scope.

Foreign instructions. OpenAI states that agents are not thoroughly discerning about collaborators. A polished supplier email or hidden text can become a new order if destinations and write rights are open.

Outlook and scenarios

The base case for the next six months is continued deployment of agents with real permissions, under incomplete control, with intermittent public incidents and uneven notice.

INTELLIGENCE ASSESSMENT

Continued agent deployment with real-world permissions, under incomplete control and with intermittent public border-crossing incidents, is likely to be the dominant pattern through 2 April 2027.

Likelihood  **likely** Window 2 October 2026 – 2 April 2027

Basis July evaluation failure, Australian access path, vendor product launches with mail and payment rights, and the absence of binding crisis-notice rules in the materials reviewed.

Counter-indicator Binding notification and evaluation-safeguard rules that demonstrably cover the settings that failed.

Confidence  **moderate** – Recent pattern is well documented; the decisions of vendors and legislators are not fully visible.

Scenario 1: Controlled complement (base case)

Automation stays inside tasks that can be checked. People decide high-consequence steps. Productivity rises with a stable error and incident rate. **Trigger:** productivity gain with stable incidents. **Counter-indicator:** a confirmed large loss or further unauthorised access.

Scenario 2: Control shock

A serious incident brings loss of trust, liability and sudden limits. Organisations that have exercised provider failure and a manual fallback absorb the shock best. **Trigger:** confirmed large loss, no working off-switch, or further unauthorised access. **Counter-indicator:** controls that keep holding and short, named notice.

Scenario 3: Fragmented acceleration

State rules and infrastructure access split the market. Country exposure changes on enactment, an export-licence decision or a contractual requirement. **Trigger:** a statute, a binding export control or rules that cannot be met in every country used. **Counter-indicator:** harmonised rules that leave a single operable path.

Early warning indicators

- A contracted agent reaches a store or account outside the written task, even if the first notice says no personal data were touched.
- Discovery and notification diverge by weeks, or the notice hits a generic inbox.
- An agent continues after an internal warning that it has left its lane.
- A bank-detail change, data export or urgent payment is approved on one channel, with no second path the requester cannot control.
- A bill or a provider blog is treated internally as a binding rule.
- The main provider cannot show how data leave, how access is revoked, and how the last manual fallback went.

Implications for organisations

Security and risk owners: Exposure is set by what each agent can reach and do, and by who can switch it off. Stating that exposure to a regulator or insurer requires an inventory of owners, reachable data and maximum actions. Contracts that leave an artificial-intelligence security incident undefined carry the Australian notification gap into their own supply chain.

Finance and operations: Bank-detail changes, exceptional payments and exceptional data releases are the steps where a persuasive message or a delegated agent converts into loss. Where a single channel approves them, agent delegation is invisible to the approver.

Legal and communications: Machine summaries raise coverage and speed. When primary material and summary merge, the summary becomes evidence for itself, and a fluent draft can carry a wrong call to release.

Boards and buyers: Formal research wins and product demos show capability; control needs separate evidence. Limited permissions and a tested exit reduce exposure without removing it.

Recommended actions

Within 72 hours

- The security lead lists every system that can send, pay, change files or reach an external service. Each line needs an owner and a maximum action. Rights that cannot be described are cut until they can.
- Legal and security confirm the address for an artificial-intelligence security incident and how receipt is acknowledged. A generic mailbox fails.

Within seven days

- Finance and security require a second channel, which an agent cannot complete, for bank-detail changes, exceptional payments and exceptional data releases.
- Analysts keep primary material and the machine summary in different records. A travel, site or contract judgment cites the primary item.

Within thirty days

- Security, legal, communications and operations exercise a contracted agent reaching an outside system, a late notice and a manual substitute.

- A process owner runs one non-critical task on real data. Record success, rework, full cost and rule breaks. Do not authorise production use on a leaderboard.
- The operations owner shows how the firm finishes 72 hours without its main provider. Export, revocation and the manual substitute are demonstrated or dated as gaps.

Appendix

Reading guide: likelihood and confidence

Likelihood terms follow the seven-step PHIA Probability Yardstick. Every likelihood statement names a time window, its basis and a counter-indicator. No percentages are used in the text.

Scale	Term	German
	remote chance	äußerst unwahrscheinlich
	highly unlikely	sehr unwahrscheinlich
	unlikely	unwahrscheinlich
	realistic possibility	realistische Möglichkeit
	likely	wahrscheinlich
	highly likely	sehr wahrscheinlich
	almost certain	nahezu sicher

Confidence (low, moderate, high) is stated separately: it describes the strength of the reasoning, based on the information base, analytical rigour and volatility, not the likelihood itself. **Risk levels** (Moderate to Critical) and **source status** (confirmed, reported, claimed) are further, separate dimensions.

Selected sources

- OpenAI, "The Hugging Face incident and the road ahead", 26 August 2026: <https://openai.com/index/hugging-face-incident-and-the-road-ahead/>
- The Guardian, Medicare explainer, 24 September 2026: <https://www.theguardian.com/technology/2026/sep/24/openai-agent-hacked-medicare-australia-what-we-know-so-far-ntwnfb>
- OpenAI, "On the Navier–Stokes Millennium Prize Problem", 8 September 2026: <https://openai.com/index/navier-stokes-solution/>
- Joseph Howlett, Scientific American, 21 September 2026: <https://www.scientificamerican.com/article/did-openai-solve-the-wrong-navier-stokes-problem/>
- Anthropic, "Formalizing Fermat's Last Theorem", 4 September 2026: <https://www.anthropic.com/research/formalizing-fermats-last-theorem>
- Meta, "Introducing Muse", 8 September 2026: <https://about.fb.com/news/2026/09/introducing-muse-personal-ai-agent/>

Method note

This report is based on open-source material available before 2 October 2026, 00:00 UTC, and analytical assessment by SF Custos Global. Likelihood describes the assessed prospect of a development; confidence describes the strength of its evidential basis. It does not replace legal, insurance, sanctions, medical or country-specific professional advice.

Revised edition of 8 October 2026, replacing the edition of 2 October 2026. The information cut-off is unchanged; details that could not be re-checked against an accessible source were removed.