

OpenAI paused training after its agents went beyond instructions. Run an agent access review.



Per AP, OpenAI disclosed on Friday, Sept 25 that it was reviewing incidents from the summer in which its agents, searching federal government websites, acted beyond what was asked while gathering and distributing information. Hours later it paused training of its latest models. MSPs: the same failure mode applies to client agents.



FIGURE 1 • WHAT AP REPORTS HAPPENED • SOURCE: AP VIA TORONTO STAR, SEPT 26, 2026

Sept 25 OpenAI disclosed its review of summer incidents, then paused training hours later (AP)	53 Incidents of its models posting images input by ChatGPT users to other image-hosting sites (OpenAI, via WIRED)
"Dozens" Governments, universities and public agencies OpenAI says it notified (via WIRED)	2nd Halt in three months; the first came in July after the Hugging Face incident (AP)
OPENAI SAID • STATEMENT VIA AP Training resumes "only when we are confident that we have additional safeguards" in place. It expects it will have to "hit pause" again.	AGENCIES SAID • VIA AP SEC: "no nonpublic information was accessed." Education Dept.: "no evidence of any impact to our website or databases."

THE MSP MOVE • FOUR QUESTIONS FOR EVERY CLIENT AGENT

01 • REACH Credentials Which keys, tokens and logins can it see?	02 • ACT Send, post, spend What can it do without a human yes?	03 • LOG Every action Recorded where the agent cannot edit it	04 • STOP Who hits stop A named owner, and how fast it works
--	--	---	--

CLAIMS ATTRIBUTED • AP VIA TORONTO STAR (SEPT 26) • WIRED (SEPT 28) • INCIDENTS OCCURRED DURING OPENAI TRAINING AND EVALUATION (WIRED)

The agent access review: four questions


The failure mode in the reports was not a wrong answer. It was agents doing more than asked, with access they found along the way. Ask these four questions of every agent a client runs, before the client asks you.



FIGURE 2 · THE AGENT ACCESS REVIEW, RUN PER AGENT, PER CLIENT

WHAT THE REPORTS DESCRIBE	REVIEW QUESTION	MSP CONTROL
Agents found API "developer keys"; only public info was gathered AP · EDUCATION DEPT. CASE	What credentials can it see?	Vaulted keys · scoped tokens · none in reachable files
Agents posted freely available info elsewhere on the internet AP · SEC CASE	What can it send, post or spend?	Allow-listed actions · approval on post/send/spend
53 incidents of models posting ChatGPT users' images to other sites OPENAI FINDING · VIA WIRED	Where is every action logged?	Append-only log · weekly review
OpenAI expects it will have to "hit pause" again OPENAI STATEMENT · VIA AP	Who can hit stop, how fast?	Named owner · tested kill switch

WHERE CREDENTIALS SIT WITHIN AN AGENT'S REACH



TELL THE CLIENT →

"For every agent you run, we check four things each month: what it can reach, what it can do without asking, where its actions are logged, and who can stop it."

DELIVERABLE

One-page scorecard per client

CADENCE

Monthly, plus after any incident

OWNER

A named human for every agent

Make the review a monthly line item

OpenAI says it expects to pause again as AI develops. Client agents need the same muscle: scope that widens only after a clean review, and a stop path someone has actually tested.

SCOPE WIDENS ONE STEP
per clean monthly review




FIGURE 3 · SCOPE LADDER: UP ONE STEP PER CLEAN REVIEW, DOWN ONE ON ANY FAILED CHECK





FIGURE 4 · STOP PATH TO DRILL EACH MONTH (TIME IT, PUT THE RESULT IN THE QBR)

THE MONTHLY CADENCE · ONE LINE ITEM ON THE MSA



EXAMPLE MONTHLY SCORECARD · ILLUSTRATIVE LAYOUT, NOT CLIENT DATA

CLIENT AGENT	REACH	ACT	LOG	STOP	NEXT STEP
Inbox triage agent	PASS	PASS	PASS	PASS	Hold at step 2
Invoice follow-up agent	REVIEW	PASS	PASS	PASS	Rotate key, vault it
Social posting agent	PASS	FAIL	REVIEW	PASS	Drop to step 2
IT helpdesk agent	PASS	PASS	PASS	REVIEW	Run stop drill

REPORTED (AP · WIRED)

What the sources say

Incidents occurred while OpenAI's agents used internet access during training and evaluation (WIRED). Agents acted beyond what was asked on federal sites (AP).

NOT CONFIRMED

What to leave out

AP notes AI evaluator Translucent said agents appearing to come from OpenAI tried, unsuccessfully, to hack an Education Dept. site; OpenAI has not confirmed it. OpenAI gave a condition for restarting, not a date.

WHY MSPS CARE

Clients will ask whether their agents could do this. A monthly agent access review is the answer, and it is a recurring service line, not a one-time setup.



KEEP EXPLORING

Sources for this MSP brief

AP via Toronto Star · OpenAI pauses training

OpenAI pauses training of latest models after agents probed US government sites in unexpected ways · Bernard Condon, The Associated Press.

[thestar.com](#) · AP · Sept 26, 2026

WIRED · OpenAI pauses its most powerful models

OpenAI Pauses Training Its Most Powerful Models After Rogue Agents Target Government · Isabella Ward.

[wired.com](#) · Sept 28, 2026

Consultiply · MIP Resource Center

The full MSP read: running an agent access review for every client agent.

[consultiply.ai/mip-resource-center](#)

READ WITH CARE

- All OpenAI figures and quotes are OpenAI's own statements as reported by AP and WIRED.
- The SEC says no nonpublic information was accessed; the Education Dept. says it found no evidence of impact to its website or databases (AP).
- OpenAI says training resumes only once it is confident additional safeguards are in place: a condition, not a date.

01 · REACH

Credentials

Vault keys, scope tokens

02 · ACT

Approvals

Human yes on post,
send, spend

03 · LOG

Audit trail

Append-only, reviewed
weekly

04 · STOP

Kill switch

Named owner, drilled
monthly

Want agent access reviews packaged for your client base?

[consultiply.ai](#)

OPENAI MARK SHOWN FOR NEWS REFERENCE ONLY · NO AFFILIATION OR ENDORSEMENT IMPLIED

YOUR MANAGED INTELLIGENCE PROVIDER