

HAI-GOS

AI Incident Case Library

Governance Education Series | Real Cases. Real Consequences.

July 4, 2026 | Compiled by HAI-GOS LLC

*"The harm did not require malicious intent.
It required only the absence of structure."*

-- Sam Yates, Founder, HAI-GOS

Human. First. Always.

ABOUT THIS DOCUMENT

This case library documents real AI-related incidents involving real organizations, real people, and real consequences. All cases are matters of public record. Every dollar figure, court ruling, settlement amount, and regulatory finding cited here is sourced from official government records, court documents, and verified news reporting.

This library serves three purposes:

- Education -- to help organizations understand that the consequences of ungoverned AI are not theoretical.
- Prevention -- to demonstrate the specific governance gaps that made each incident possible.
- Connection -- to show how HAI-GOS directly addresses each of the failure modes documented here.

These are not edge cases. They are patterns. The same three governance gaps appear at the root of nearly every incident: no defined policy, no named owner, and no human review requirement. HAI-GOS was built to close all three.

THE SCALE OF THE PROBLEM

Before the individual cases, the aggregate data tells an important story.

STATISTIC	FINDING	SOURCE
AI Incidents 2025	362 documented AI-related incidents in 2025 -- a 55% rise from 233 in 2024	Stanford HAI 2026 AI Index Report
No Governance Policy	63% of breached organizations had no formal AI governance policy	IBM Cost of Data Breach Report 2025
No Access Controls	97% of organizations that experienced AI breaches lacked proper AI access controls	IBM Cost of Data Breach Report 2025
Governance Gap	78% of organizations use AI. Only 25% have implemented governance programs.	AuditBoard 2025
Shadow AI Cost	Shadow AI breaches cost an average of \$670,000 more than standard incidents	IBM 2025
Average Breach Cost	Organizations without AI governance paid the highest average breach cost: \$5.52 million	IBM 2025
Regulatory Growth	145 AI laws passed in 2025. More than 1,400 AI-related class action lawsuits filed.	DataGrail 2026 Report
Board-Level Risk	72% of S&P 500 companies disclosed a material AI risk in 2025 -- up from 12% in 2023	Harvard Law Forum / Conference Board 2025
Governance Deficit	Only 8% of organizations globally have a comprehensive AI governance framework	Multiple enterprise surveys 2025-2026

The biggest AI failures of 2025 were not technical. They were organizational: weak controls, unclear ownership, and misplaced trust. -- ISACA, December 2025

What Happened

On March 11, 2023, Samsung Semiconductor lifted its internal ban on employee use of ChatGPT. Engineers were permitted to use the tool to help debug code and generate documentation. Within twenty days, three separate incidents of proprietary data exposure had occurred.

Incident 1 -- Semiconductor Database Source Code

An engineer pasted proprietary source code from an internal semiconductor measurement database into ChatGPT, requesting help identifying and fixing bugs. The code contained critical information about Samsung chip fabrication processes.

Incident 2 -- Defect Detection Algorithms

A separate engineer submitted code related to yield and defect measurement for semiconductor manufacturing equipment, requesting optimization suggestions. This code represented core intellectual property in one of the world most secretive industries.

Incident 3 -- Confidential Internal Meeting Transcript

An employee recorded an internal business meeting, converted the audio to text using a third-party transcription service, and submitted the complete transcript to ChatGPT to generate meeting minutes. The meeting discussed unreleased semiconductor process technology.

The Consequences

Samsung's internal security team identified the leaks and issued an urgent company-wide notice. By May 2023, Samsung had banned ChatGPT and all external generative AI tools on company-owned devices and internal networks entirely.

The critical consequence: OpenAI's terms of service at the time permitted using submitted content to improve its models. Samsung's proprietary source code, defect algorithms, and internal meeting content became part of OpenAI's training data. It did not just leave Samsung's building -- it potentially became accessible to every other user of the platform. The data could not be retrieved or deleted.

Samsung announced plans to develop an internal AI system built on models running within its own infrastructure, accelerating a multi-year development effort at significant cost.

The Governance Gap

- No AI usage policy defined what employees could and could not enter into external AI tools
- No review requirement before proprietary code or confidential content was submitted
- No AI Lead or designated person responsible for overseeing AI tool use
- No training on data exposure risk before access was granted
- Access was granted with only a vague internal memo advising caution

The three engineers were not acting maliciously. They were doing their jobs as efficiently as they knew how. No structure told them otherwise. That is the governance gap.



HAI-GOS GOVERNANCE LESSON

Speed of AI adoption without governance structure is not efficiency. It is exposure. The twenty-day gap between granting access and discovering three major incidents is what unstructured AI adoption looks like in a real organization.



HOW HAI-GOS PREVENTS THIS

HAI-GOS Module 7 (Policy and Documentation) establishes a written AI usage policy that employees acknowledge -- specifically prohibiting submission of proprietary code, confidential data, or internal meeting content to public AI tools. Module 5 (AI Tool Inventory) tracks which tools are authorized and for what purpose. Module 6 (Risk Register) identifies data exposure as a primary risk category. The AI Lead (Module 1) owns enforcement. Samsung had none of these.

What Happened

Steven Schwartz, a senior associate at the New York law firm Levidow, Levidow and Oberman, was researching legal precedents to support a personal injury case against the airline Avianca on behalf of client Roberto Mata. The case involved injuries sustained when a metal serving cart struck Mata on board a 2019 flight. Avianca had filed a motion to dismiss.

Schwartz, unfamiliar with ChatGPT but aware of its existence, turned to the tool to supplement his legal research. He described it as 'this new site which I falsely assumed was like a super search engine.' He submitted prompts asking for relevant case citations.

ChatGPT returned six detailed citations -- complete with case names, case numbers, courts, judges, ruling details, and quoted language. Schwartz assumed they were real. He submitted them to federal court. They were not real. ChatGPT had fabricated all six entirely.

The Six Fabricated Cases

When opposing counsel Avianca raised concerns that the cited cases could not be found, and when Judge P. Kevin Castel of the Southern District of New York demanded the full text of the opinions, the attorneys submitted additional ChatGPT-generated text -- still fabricated -- in response to the court orders. The fabricated opinions contained what the judge later described as legal gibberish.

Among the fabricated cases: Varghese v. China Southern Airlines and Zicherman v. Korean Air Lines. Neither existed. The judges named as authors of the fictional opinions were real judges who had never written the cited rulings. The airline companies were real. The cases were not.

The Consequences

Judge Castel sanctioned both attorneys -- Steven Schwartz and Peter LoDuca -- fining each \$5,000. The firm Levidow, Levidow and Oberman was also fined \$5,000. The attorneys were ordered to personally send letters to every judge whose name ChatGPT had falsely attributed to a fabricated ruling, attaching a copy of the judge's sanctions order.

Avianca's motion to dismiss was granted. The original case was lost. The firm was publicly humiliated in a case that received international news coverage. Multiple courts subsequently implemented or strengthened requirements for disclosure of AI use in legal filings.

At the sanctions hearing, Schwartz testified: 'It just never occurred to me that it would be making up cases. My reaction was, ChatGPT is finding that case somewhere. Maybe it's unpublished. Maybe it was appealed. Maybe access is difficult to get. I just never thought it could be made up.'

The Governance Gap

- No firm policy on AI use in legal research or client work product
- No review requirement for AI-generated content before submission to court
- No training on AI hallucination risk
- No designated person responsible for AI oversight or quality control
- No verification step between AI output and official filing

The attorney did not intend to deceive anyone. He simply had no policy, no training, and no human review requirement. The absence of those three things cost him his professional reputation. The harm did not require malicious intent. It required only the absence of structure.



HAI-GOS GOVERNANCE LESSON

AI tools do not function like search engines. They predict text. They generate what sounds correct rather than what is correct. Any organization using AI to produce content that goes to clients, partners, regulators, or courts without human review is running the same risk Levidow, Levidow and Oberman ran. The professional consequences can be irreversible.



HOW HAI-GOS PREVENTS THIS

HAI-GOS Module 3 (Governance Roadmap) identifies AI output review as a required governance step for any content leaving the organization. Module 7 (Policy and Documentation) includes explicit language on human verification before AI-generated content is used in professional, legal, or client-facing contexts. Module 8 (Incident Management) creates a documented response protocol when AI output errors occur. The AI Lead owns verification requirements.

Two Hundred Qualified Applicants. Automatically Rejected. Nobody Was Watching.

iTutorGroup -- EEOC Federal Enforcement Action | 2020 incident / 2022 lawsuit / 2023 settlement | **\$365,000 FEDERAL SETTLEMENT + 5-YEAR REGULATORY MONITORING**

What Happened

iTutorGroup, a company providing English-language tutoring to students in China, deployed an AI-powered recruitment software tool to screen applicants for tutor positions. The system was configured -- whether intentionally or through a failure of oversight -- to automatically reject female applicants over the age of 55 and male applicants over the age of 60, regardless of their qualifications, experience, or any other factor.

In March and April 2020, more than 200 qualified United States-based applicants were automatically rejected by the AI screening tool. None received human review. No human reviewer ever evaluated these candidates or the pattern of rejections. The AI did exactly what it was configured to do.

How the Discrimination Was Discovered

The pattern was discovered when one applicant -- Wendy Pincus -- noticed the rejection and submitted a second application with a more recent birth date. The second application, otherwise identical to the first, was offered an interview. The only variable was her age.

The Equal Employment Opportunity Commission investigated and filed suit in May 2022 under the Age Discrimination in Employment Act. The case was settled on August 9, 2023 -- the EEOC's first enforcement action in United States history against an AI-powered recruitment tool.

The Consequences

iTutorGroup settled for \$365,000 to be distributed to applicants who were rejected due to their age. The consent decree -- court-approved September 8, 2023 -- included the monetary settlement plus significant non-monetary requirements:

- Adoption of new anti-discrimination policies
- Distribution of an internal memo to all relevant personnel
- Multiple mandatory anti-discrimination training sessions
- Invitation to all wrongfully rejected applicants to reapply
- Written notice to the EEOC of any future discrimination complaints
- Cessation of requesting birth dates from applicants
- Five or more years of EEOC monitoring of hiring practices

iTutorGroup has since ceased hiring tutors in the United States. The regulatory monitoring extends for the duration of any future US hiring activity.

The Governance Gap

- No AI governance structure overseeing automated hiring decisions
- No human review of AI screening outputs before rejections were issued
- No audit of the AI tool for discriminatory patterns before deployment
- No designated owner of AI oversight in the hiring process
- No mechanism for detecting or reporting discriminatory AI behavior

The AI did exactly what it was configured to do. Whether the discriminatory configuration was intentional or the result of unreviewed design choices is legally irrelevant. When AI makes decisions that affect people and no human is assigned to oversee those decisions, the organization owns every outcome.



HAI-GOS GOVERNANCE LESSON

Automated AI decisions affecting employees or applicants require human oversight. Full stop. The EEOC has made clear that anti-discrimination law does not recognize an AI exception. If your AI tool makes a decision that a human could not legally make, your organization is liable for that decision. The AI is not.



HOW HAI-GOS PREVENTS THIS

HAI-GOS Module 6 (Risk Register) explicitly identifies automated decision-making affecting people as a high-risk governance category requiring human review. Module 5 (AI Tool Inventory) documents every AI tool in use, including hiring software, with data sensitivity and decision authority classifications. The AI Lead is assigned to ensure no automated system makes final decisions about people without human oversight. This is the governance gap iTutorGroup did not close.

ChatGPT Invented a Crime. Named a Real Person. A Lawsuit Followed.

Mark Walters v. OpenAI -- Georgia State Court | June 2023 (filed) / May 2025 (resolved) |
FIRST AI DEFAMATION CASE IN US HISTORY -- DISMISSED ON NARROW GROUNDS

What Happened

On May 4, 2023, journalist Frederick Riehl, editor of AmmoLand.com and a board member of the Second Amendment Foundation (SAF), was researching a lawsuit filed by the SAF against the Washington State Attorney General. He used ChatGPT to help summarize the lawsuit, providing it with portions of the complaint.

ChatGPT produced a summary that was initially accurate. But when Riehl provided a URL and asked for a fuller summary, ChatGPT -- which informed him it could not access the internet -- proceeded to generate a completely fabricated account. The fabricated response stated that the SAF lawsuit involved allegations of embezzlement and fraud against a man named Mark Walters, whom ChatGPT described as the organization's treasurer and chief financial officer.

Mark Walters is a nationally syndicated radio host and prominent gun rights advocate. He has never held any position at the SAF, has no connection to the lawsuit, and was not accused of anything by anyone. ChatGPT invented the crime, the role, the accusations, and the legal complaint -- entirely from nothing.

The Fabricated Accusations

ChatGPT generated a detailed false legal complaint stating that Walters had:

- Misappropriated SAF funds for personal expenses without authorization
- Manipulated financial records and bank statements to conceal his activities
- Failed to provide accurate financial reports to SAF leadership
- Committed fraud and breach of fiduciary duty

None of this was true. None of it existed. ChatGPT generated confident, detailed, false accusations against a real named person.

The Lawsuit

Mark Walters filed a defamation lawsuit against OpenAI in Georgia state court on June 5, 2023 -- believed to be the first defamation lawsuit in United States history related to ChatGPT. The case proceeded through litigation for two years.

In January 2024, a Gwinnett County Superior Court judge denied OpenAI's motion to dismiss, allowing the case to proceed. Discovery followed. In May 2025, Judge Tracie Carson granted summary judgment in favor of OpenAI on three independent grounds: the journalist who received the fabricated output recognized it was false within roughly ninety minutes and did not publish it; the output could not be reasonably understood as communicating actual facts given ChatGPT's multiple disclaimers; and Walters could not demonstrate he suffered actual, quantifiable damages since the false information was never published.

Why This Case Matters for Organizations

OpenAI won on narrow legal grounds specific to this fact pattern. The journalist recognized the hallucination and did not publish. The false information did not reach the public. If it had -- if the journalist had published the AI-generated false accusations as fact -- the legal and reputational consequences for everyone involved would have been categorically different.

For organizations using AI to generate content about clients, partners, competitors, or individuals: if AI-generated false information about a third party is published or acted upon, and that content was produced on company time, using company tools, for a company purpose, the question of responsibility will not stay theoretical.

The Governance Gap

- No review requirement for AI-generated content about named individuals before use
- No policy on verification of AI output when content concerns real people or real cases
- No training on AI hallucination risk for staff using AI in research or content production
- No designated owner of content accuracy standards for AI-assisted work

Your organization does not have to create false information to be responsible for what happens when AI-generated false information is used without human review.



HAI-GOS GOVERNANCE LESSON

AI hallucination is not a rare edge case. It is a documented, consistent behavior of generative AI systems. Any AI output about a real named person, organization, case, or event requires human verification before it is used, published, or acted upon. The governance structure that makes this a requirement -- not a best practice -- is what HAI-GOS provides.



HOW HAI-GOS PREVENTS THIS

HAI-GOS Module 7 (Policy and Documentation) includes explicit requirements for human verification of AI-generated content that references real people, organizations, or legal matters before any external use. The AI Lead is responsible for ensuring this requirement is known, understood, and followed. Module 3 (Governance Roadmap) includes content review as a priority governance action item for organizations using AI in research, communications, or content production.

MARKET INTELLIGENCE SUMMARY

The four documented cases above represent a pattern, not a collection of outliers. The same governance gaps appear at the root of every incident across every industry and every organization size. Below is the statistical context that places these cases in the broader landscape of AI governance failure.

The Three Governance Gaps That Appear in Every Case

GOVERNANCE GAP	APPEARS IN
No defined AI usage policy employees have acknowledged	Samsung (no policy on what could be entered into AI tools), Law Firm (no policy on AI research or verification), iTutorGroup (no policy on AI decision oversight), Georgia Case (no policy on AI content verification)
No named individual responsible for AI oversight	All four cases. In every incident, when something went wrong, there was no single person whose job it was to have prevented it or to own the response.
No human review requirement before AI output reaches the outside world	Samsung (code submitted without review), Law Firm (AI output filed in court without verification), iTutorGroup (AI rejections issued without human review), Georgia Case (AI content used in research without verification)

The Statistical Context

- 362 documented AI incidents in 2025 -- 55% rise year over year (Stanford HAI 2026)
- 63% of breached organizations had no AI governance policy (IBM 2025)
- 97% of organizations with AI breaches lacked proper access controls (IBM 2025)
- Only 8% of organizations globally have a comprehensive AI governance framework
- 78% of organizations use AI -- only 25% have implemented governance programs
- 145 AI laws passed in 2025. Regulatory pressure is accelerating.
- \$5.52 million average breach cost for organizations without AI governance (IBM 2025)
- 72% of S&P 500 companies flagged AI as a material risk in 2025 -- up from 12% in 2023

"In 2026, competitive advantage will not come from using more AI, but from governing it well. Organizations that maintain visibility, clear ownership, and rapid intervention will reduce harm and earn trust." -- ISACA, December 2025

Human.First.Always.

HAI-GOS LLC | Stuart, Florida | haigos.ai | sam@haigos.ai | (772) 528-5185
HAI-GOS is Organizational Governance Memory for the AI era.