

Kiteworks

AI Data Security and Compliance Risk Report

Industry Study Reveals 83% of Organizations Operate Without Basic Controls

```
def process_input(input_value):  
    # Placeholder code for processing input  
    return "Processed input: " + str(input_value)  
  
def calculate_output(input_data):  
    # Placeholder code for calculating output  
    return "Output calculated for data: " + input_data  
  
def analyze_data(data):  
    # Placeholder code for data analysis  
    return "Data analyzed: " + data  
  
num_list = [105, 204, 458,  
            701, 958, 1002,  
            1136, 1345, 1560,  
            1823, 2150, 2356,  
            2579, 2793, 2945,  
            3086, 3280, 3404,  
            3611, 3912, 4119,  
            4327, 4567, 4789,  
            4953, 5272, 5412]  
  
for num in num_list:  
    print(process_input(num))  
    print(calculate_output(num))  
    print(analyze_data(str(num)))  
    print("-" * 20)
```

010101001101
00110101010
101010010110
10011010101
010101010101
010101010101
0110100110101

```
def process_input(input_value):  
    # Placeholder code for processing input  
    return "Processed input: " + str(input_value)  
  
print(process_input("data1")) # Processed input: data1  
print(process_input("data2")) # Processed input: data2  
print(process_input("example")) # Processed input: example  
  
def calculate_output(input_data):  
    # Placeholder code for calculating output  
    return "Output calculated for data: " + input_data  
  
print(calculate_output("data1")) # Output calculated for data: data1  
print(calculate_output("data2")) # Output calculated for data: data2  
print(calculate_output("example")) # Output calculated for data: example  
  
def analyze_data(data):  
    # Placeholder code for data analysis  
    return "Data analyzed: " + data  
  
print(analyze_data("data1")) # Data analyzed: data1  
print(analyze_data("data2")) # Data analyzed: data2  
print(analyze_data("example")) # Data analyzed: example
```

Table of Contents

- 3 Executive Summary
- 4 State of AI Security Implementation
- 6 Data Exposure
- 9 Compliance Blind Spots
- 11 Industry Paradoxes
- 13 Key Takeaways: The Window for Action
- 15 References



Executive Summary

Companies want to use AI tools but can't protect their data. Our survey of 461 cybersecurity, IT, risk management, and compliance professionals, conducted by Centiment, reveals that most organizations operate without basic safeguards while employees freely share sensitive information with chatbots and other AI services.

83% of companies rely on training sessions, warning emails, or nothing at all to prevent data exposure to AI tools.



The numbers tell the story: Only 17% of companies can automatically stop employees from uploading confidential data to public AI tools. The other 83% rely on humans—training sessions/audits, warning emails, or guidelines—or nothing at all—that present a serious gap. Meanwhile, employees routinely paste customer records, financial data, and trade secrets into ChatGPT and similar AI services—often from personal devices the company can't even see.


This creates a perfect storm. As organizations rush to embrace the scale and efficiencies AI offers, many think they're protected when they're not. Security incidents involving AI jumped 56% last year.¹ Hackers now specifically target the file transfer systems companies use to move data to AI platforms. And employees increasingly use unauthorized "shadow AI" tools, exposing passwords and company credentials without anyone knowing.

The C-Suite Gap

Companies face shrinking time to fix these issues as threats multiply and regulators prepare enforcement actions. The disconnect between what executives believe and what's actually happening creates unprecedented risk.

State of AI Security Implementation

Security Control Maturity Pyramid





Most companies lack the basic tools needed to protect their data from AI exposure. When we asked how they prevent employees from uploading private information to public AI tools, the responses revealed widespread vulnerability.

The security breakdown looks like this: At the top, only 17% have technology that blocks unauthorized AI access and scans for sensitive data—the bare minimum for protection. The largest group (40%) depends on employee training and occasional reviews, hoping people follow the rules. Another 20% send warning emails but never check if anyone listens, while 10% simply have issued guidelines to follow. The remaining 13% have no policies at all.

Reality Check: Companies overestimate their AI governance capabilities—only 9% have working systems despite 33% claiming full control.

www.kiteworks.com





Independent research exposes dangerous self-deception:

Gartner®

Only 12% have dedicated AI governance structures²

Deloitte.

Just 9% achieve AI “Ready” level maturity³

Result: Over 3x overestimation of actual capabilities

This overconfidence gap creates a perfect storm—organizations make strategic decisions based on imaginary protections while their actual security posture deteriorates.

The visibility problem runs deeper. While 33% of executives claim they track all AI usage, independent studies show only 9% actually have working governance systems. This means companies overestimate their capabilities by over threefold, making critical decisions based on false confidence.

Shadow AI makes this worse. With 86% of organizations blind to AI data flows, employees download AI apps, browser extensions, and mobile tools without IT approval.⁴ They share company login credentials with AI assistants to “help with tasks.” They upload entire databases to get quick analysis. All invisible to security teams. When credentials leak, the median remediation time stretches to 94 days—over three months of open access to your systems.⁵

Consider what this means: A sales rep uploads customer lists to analyze buying patterns. An HR manager shares employee reviews to draft performance summaries. A finance analyst pastes quarterly results to create presentations. Each action creates permanent risk—once data enters an AI system, it becomes embedded in the model itself, potentially accessible to competitors or malicious actors. You can’t delete it. You can’t retrieve it. It’s there forever.

The scale of shadow AI proliferation is staggering. According to a recent report by Varonis, 98% of companies have employees using unsanctioned applications, with each organization averaging 1,200 unofficial apps—creating an enormous attack surface that security teams cannot monitor.⁶ More alarming, 52% of employees actively use high-risk OAuth applications that can access and exfiltrate company data. This shadow IT sprawl means that even organizations claiming comprehensive AI governance likely have over a thousand backdoors they don’t know exist, each one a potential pipeline for sensitive data to flow directly into AI training models or competitor hands.

Data Exposure

The most disturbing finding involves what employees actually share. The largest group of companies—27%—admit that more than 30% of information sent to AI tools contains private data. This includes Social Security numbers, medical records, credit card information, strategic plans, and protected intellectual property.

27% of organizations report that over 30% of their AI-bound data contains private information—customer records, employee data, and trade secrets.




Another 17% have no idea what their employees share. They can't answer basic questions: What customer data left the building today? Which employees sent financial records to AI tools? How many times did someone share passwords or system credentials?

This differs from traditional data breaches in critical ways. Instead of one big leak, it's thousands of small ones every day. Instead of hackers stealing data, employees give it away voluntarily. Instead of clear audit trails, the evidence scatters across personal devices and shadow IT tools.

The exposure extends beyond just ChatGPT queries. Platform-specific risks multiply the danger: 90% of organizations have sensitive files exposed to all employees through Microsoft 365 Copilot, with an average of 25,000+ sensitive folders accessible to anyone who asks the right prompt. In Salesforce environments, 100% of companies have at least one account capable of exporting all data, while 92% allow users to create public links that AI crawlers can index. These aren't theoretical vulnerabilities—they're active exposure points where a single misguided query can surface years of confidential data.

Breaking down attempts to control AI usage reveals why exposure stays high. Every department—from IT to legal—relies on the honor system rather than technology. They watch what comes out of AI tools instead of controlling what goes in. By then, sensitive data already lives in the AI system, potentially training future versions that competitors might access.

Private Data Exposure Distribution



The industry table exposes just how widespread and systemic AI-driven data risk has become across every major industry. Despite regulatory mandates and the highly sensitive nature of the information at stake, nearly four in 10 organizations in each sector admit that more than 16% of the data their employees send to AI tools is private. The most severe exposure is found in the technology sector, where a staggering 27% of companies report that **over 30%** of AI-processed data is private or sensitive—the highest of any industry. Sectors like healthcare, finance, and manufacturing are only a step behind, with 27% of organizations in each category reporting the same extreme level of exposure. Even the legal sector, whose very existence depends on confidentiality, shows more than one in five firms pushing their most sensitive data into AI tools with little oversight.

Equally troubling is the remarkable uniformity of this risk. No industry can claim a safe harbor. The proportion of organizations reporting significant private data exposure remains virtually identical across sectors—from government to life sciences to professional services. Meanwhile, roughly **17% of organizations in every vertical openly admit they have no idea how much sensitive data employees are sharing with AI platforms.** This blind spot reveals the true scope of the epidemic: It's not just a few reckless actors or under-resourced fields, but a universal problem driven by a lack of technical controls, over-reliance on trust, and the rapid pace of AI adoption. Until organizations can measure and govern what leaves their four walls, the risk will only escalate.

AI-Processed Private Data Exposure by Industry

Industry	0%	1%-5%	6%-15%	16%-30%	Over 30%	Don't Know
All Industries	7%	13%	23%	14%	26%	17%
Technology	7%	13%	23%	13%	27%	17%
Healthcare	6%	14%	23%	13%	26%	17%
Financial Services	7%	13%	23%	13%	26%	17%
Legal/Law	8%	15%	23%	15%	23%	16%
Government	7%	13%	24%	13%	26%	17%
Life Sciences/ Pharmaceuticals	9%	13%	22%	13%	26%	17%
Manufacturing	7%	13%	24%	13%	26%	17%
Professional Services	9%	13%	22%	13%	26%	17%

Credential Exposure:
Employees routinely share usernames, passwords, and access tokens with AI assistants—median remediation time: 94 days.



The credential exposure problem deserves special attention. Employees routinely share usernames, passwords, and access tokens with AI assistants to “streamline workflows.” Each exposed credential becomes a backdoor into company systems, available to anyone who knows how to extract it from the AI platform. With remediation taking 94 days on average, attackers have months to exploit each leaked credential.

Compliance Blind Spots

Companies dramatically underestimate regulatory risk. Only 12% list compliance violations among their top AI concerns—the same priority as exotic threats like “data poisoning” that rarely occur. Meanwhile, real enforcement accelerates.

59 AI regulations issued by U.S. agencies in 2024—more than double the previous year.



U.S. agencies issued 59 AI regulations in 2024, more than double the previous year. Globally, 75 countries increased AI legislation by 21%.⁷ These aren't suggestions—they carry million-dollar penalties and criminal liability for executives.

Current practices violate specific regulatory provisions daily:

- **GDPR Article 30** requires maintaining records of all processing activities—impossible when you can't track AI uploads
- **CCPA Section 1798.130** mandates ability to track and delete personal information upon request—but companies don't know which AI systems have it
- **HIPAA § 164.312** demands comprehensive audit trails for all ePHI access—unachievable with shadow AI
- **SOX** requires financial data controls that AI usage completely bypasses


Without visibility into AI interactions, companies can't respond to customer data requests, prove compliance during audits, investigate breaches, implement deletion requirements, or demonstrate data protection to regulators.

The shortcuts driven by ROI pressure make this worse. Executives demand quick AI wins to justify investment, leading to specific compromises:

- Security assessments get abbreviated to “high-level reviews”
- Governance frameworks become “future roadmap items”
- Risk evaluations shrink from comprehensive audits to checkbox exercises
- Compliance reviews get deferred until “after we see results”

Each shortcut multiplies compliance exposure while regulators sharpen enforcement tools.


The compliance gap becomes critical when paired with identity management failures. Organizations average 15,000 “ghost users”—stale but enabled accounts that retain full access to systems and data. With 176,000 inactive external identities and over 31,000 stale permissions in the typical enterprise, the audit trail regulators demand becomes impossible to produce. Most critically, only 10% of companies have properly labeled their files—a fundamental requirement for GDPR Article 5 and HIPAA Privacy Rule compliance. Without proper data classification, organizations cannot demonstrate lawful processing, respond to deletion requests, or prove they’ve protected sensitive information according to its risk level.



Industry Paradoxes

Industries with the most to lose show the weakest protections, creating dangerous contradictions between requirements and reality.

Across every industry, less than one in five organizations rely on technical controls, while a staggering 68% to 71% depend primarily on human-centric measures like training/audits, warnings, and policies. Even more concerning, 13% to 15% have no formal AI data policy in place at all—a gap that extends equally to financial services, healthcare, manufacturing, and professional services. Following are some key takeaways.



Healthcare's Compliance Fiction: HIPAA requires tracking 100% of patient data access, yet only 35% of healthcare organizations can see their AI usage. Every untracked ChatGPT query containing patient information violates federal law. Only 39% of healthcare executives even recognize AI security threats—the lowest awareness of any industry. The same organizations trusted with life-and-death data operate with less security than retail stores.

Further, 90% of organizations in the Varonis study admitted that sensitive files are accessible through AI copilots, with an average of 25,000+ unprotected folders containing PHI. Plus, only 39% of healthcare executives even acknowledge AI as a security threat—the lowest of any industry.



Financial Services' Awareness-Action Gap: Banks and investment firms show the highest concern about data leaks (29%) but match the lowest implementation of controls (16%). Despite handling account numbers, transactions, and financial records, 39% admit sending substantial private data to AI tools. They know the risk but choose speed over security.

Government Paradox:
Only 17% of agencies
have technical safeguards
for citizen data—the rest
rely on hope.





Government's Public Trust Failure: Agencies protecting citizen data rely entirely on employee goodwill. Only 17% have technical safeguards, while 11% have no governance plans at all—the highest rate of any sector. With 39% reporting significant private data in AI systems, every interaction risks exposing Social Security numbers, tax records, or classified information.



Technology Sector's Credibility Crisis: Companies selling AI security to others can't secure their own usage. While 100% build AI products, only 17% protect against their own employees' AI risks—an 83% hypocrisy gap. The same firms teaching others about AI safety operate without basic controls, undermining their credibility when breaches occur.



Legal Firms' Privilege Paradox: Lawyers show the highest data leak concerns (31%) but lowest technical controls (15%). With 38% exposing private data to AI, every interaction risks breaking attorney-client privilege. State bars require absolute confidentiality, but firms rely on employee promises rather than technology.

AI Data Risks Due to Control Gaps

Industry	Technical Controls	Human-Dependent Controls	No Policies
All Industries	17	70	13
Financial Services	16	71	13
Government	17	70	13
Healthcare	17	68	14
Manufacturing	17	70	13
Technology	17	70	14
Legal/Law	15	70	15
Life Sciences/ Pharmaceuticals	17	70	13
Professional Services	17	70	13

Key Takeaways: The Window for Action

Three realities demand immediate executive attention:

First, security delusion must end. With 83% lacking real controls while 27% hemorrhage data, the fantasy of protection has become dangerous. Companies claiming readiness at three times their actual capability make strategic bets on illusions. This isn't just a technology problem—it's a leadership failure requiring honest assessment and decisive action.

Second, compliance isn't optional. With 59 new regulations last year and enforcement accelerating, the 60% of companies blind to their AI usage face immediate consequences. When auditors arrive asking for AI data trails, "we didn't know" triggers penalties, not sympathy. Demonstrating control over AI data flows has shifted from nice-to-have to survival requirement.

Third, no industry is safe. Healthcare, finance, government, and technology all fail at similar rates despite different regulations. This universal vulnerability means companies can't copy competitors—they must build their own protections. Heavy regulation hasn't driven better security, proving external pressure alone won't save you.

Making matters worse, third-party AI tools multiply exposure exponentially. Third-party involvement in breaches doubled from 15% to 30% in just one year.

Third-Party Time Bomb: Third-party involvement in breaches doubled from 15% to 30% in one year while 44% of zero-day attacks target managed file transfer systems used for AI data exchange.⁸



And as evidenced in the Varonis study, platform-specific exposures compound the risk with 100% of Salesforce deployments and 90% of Microsoft 365 environments having critical data exposed to AI tools. This makes every major enterprise system a potential data hemorrhage point. The ghost users haunting every organization—15,000 stale accounts retaining full access—ensure that even after addressing current employees, yesterday's workforce can still become tomorrow's breach.

THE PATH FORWARD requires four immediate actions:

1. **Flip the Focus: Face Reality First.** Audit actual AI usage, not theoretical frameworks. Close the 300% overconfidence gap before it closes you. Focus on tracking and controlling inputs as much as you focus on monitoring outputs.
2. **Deploy Technology-First Controls.** Human-dependent measures fail across every industry. Automated blocking and scanning represent minimum viable protection. If you can't stop the upload, you've already lost.
3. **Establish Data Governance Command Centers.** File transfer breaches prove that disconnected security creates cascading failures. Organizations need unified governance platforms that track every data movement, enforce classification policies, and maintain audit trails across all AI touchpoints. When Finastra lost 400 GB without detection for 8 days, it wasn't just a security failure—it was a governance blindness that modern AI usage makes fatal.
4. **Unify Risk Management: Gain Total Visibility.** Without knowing what data flows where, compliance becomes impossible and risk management becomes fiction. Real-time AI monitoring across all platforms—cloud, on-premises, and shadow IT—isn't optional anymore. Data lineage tracking must extend from initial creation through AI processing to final outputs, creating the forensic trail regulators now demand.

IN 18 MONTHS—OR SOONER, companies will divide into two groups: those who secured their AI usage, and those explaining their failures to regulators, customers, and shareholders.

The collision of explosive AI adoption, surging security incidents, and accelerating regulation creates a closing window. Model training contamination is permanent—every piece of sensitive data shared today becomes tomorrow's compliance violation and competitive disadvantage.

Resources

¹ “[The 2025 AI Index Report](#),” Stanford.

² “[AI Governance Frameworks for Responsible AI](#),” Gartner Peer Community, March 20, 2023.

³ Amy Dove, “[AI at a crossroads building trust as the path to scale](#),” Deloitte, February 21, 2025.

⁴ “[The 2025 AI Index Report](#),” Stanford.

⁵ “[2025 Data Breach Investigations Report](#),” Verizon, April 2025.

⁶ “[2025 State of Data Security Report: Quantifying AI’s Impact on Data Risk](#),” Varonis, June 2025.

⁷ “[The 2025 AI Index Report](#),” Stanford.

⁸ “[2025 Data Breach Investigations Report](#),” Verizon.

Legal Disclaimer

The information provided in this report is for general informational purposes only and should not be construed as professional advice. Kiteworks and Centiment make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained in this report. Any reliance you place on such information is strictly at your own risk. None of the sponsoring or contributing organizations shall be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from loss of data or profits arising out of, or in connection with, the use of this report. Readers should consult with qualified legal counsel and cybersecurity professionals when addressing specific compliance requirements.

About Centiment

Centiment is a market research firm specializing in data collection and analysis for the cybersecurity and technology sectors. The company delivers actionable insights through customized survey design, targeted respondent recruitment, and sophisticated analytics. Centiment’s proprietary research platform ensures exceptional data quality through AI-driven verification and expert human oversight. The company serves Fortune 500 enterprises, technology vendors, and government agencies, providing intelligence for strategic decisions in evolving markets. Headquartered in Denver, Centiment conducts research globally to help organizations understand complex technology landscapes and cybersecurity trends.

Kiteworks

Copyright © 2025 Kiteworks. Kiteworks’ mission is to empower organizations to effectively manage risk in every send, share, receive, and use of private data. The Kiteworks platform provides customers with a Private Data Network that delivers data governance, compliance, and protection. The platform unifies, tracks, controls, and secures sensitive data moving within, into, and out of their organization, significantly improving risk management and ensuring regulatory compliance on all private data exchanges. Headquartered in Silicon Valley, Kiteworks protects over 100 million end-users and over 1,500 global enterprises and government agencies.

June 2025

www.kiteworks.com



```
def process_input(input_value):
    # Placeholder code for processing input
    return "Processed input: " + str(input_value)

def calculate_output(input_data):
    # Placeholder code for calculating output
    return "Output calculated for data: " + input_data

def analyze_data(data):
    # Placeholder code for data analysis
    return "Data analyzed: " + data

num_list = [105, 204, 458,
            701, 958, 1002,
            1136, 1345, 1560,
            1823, 2150, 2356,
            2579, 2793, 2945,
            3066, 3280, 3404,
            3611, 3912, 4119,
            4327, 4567, 4789,
            4953, 5272, 5412]

for num in num_list:
    print(process_input(num))
    print(calculate_output(str(num)))
    print(analyze_data(str(num)))
    print("-" * 20)
```

```
def process_input(input_value):
    # Placeholder code for processing input
    return "Processed input: " + input_value

print(process_input("value1")) # Processed input: value1
print(process_input("value2")) # Processed input: value2
print(process_input("example")) # Processed input: example

def calculate_output(input_data):
    # Placeholder code for calculating output
    return "Output calculated for data: " + input_data

print(calculate_output("data1")) # Output calculated for data: data1
print(calculate_output("data2")) # Output calculated for data: data2
print(calculate_output("example")) # Output calculated for data: example

def analyze_data(data):
    # Placeholder code for data analysis
    return "Data analyzed: " + data

print(analyze_data("data1")) # Data analyzed: data1
print(analyze_data("data2")) # Data analyzed: data2
print(analyze_data("example")) # Data analyzed: example
```