CSAI Foundation | Cloud Security Alliance

Shadow AI Apps: The Enterprise Attack Surface That Outpaces Monitoring

2026-05-30

AI-assisted Rapid Research

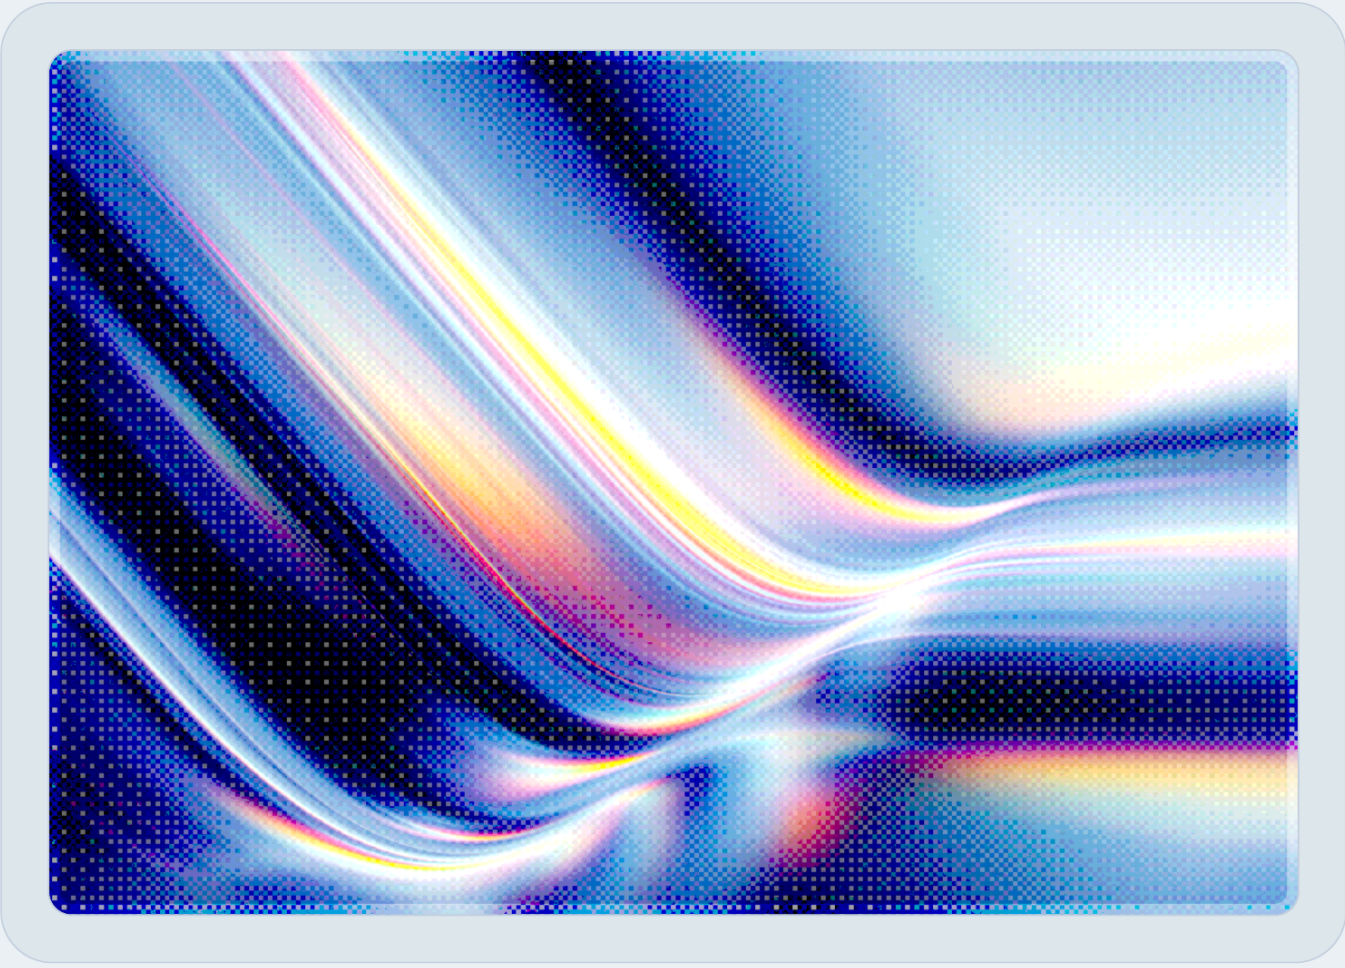

CSAI

CSA cloud security alliance®

Key Takeaways

- Eight in ten employees use AI tools not approved by their organizations [1], yet only 37% of enterprises have any AI governance policy in place — a structural gap now measurable in breach costs [23].
- Generative AI has overtaken all other channels to become the single largest vector for corporate-to-personal data movement, accounting for 32% of all such transfers [12], and 89% of enterprise AI usage is invisible to security teams [11][13].
- IBM's 2025 Cost of a Data Breach Report introduced shadow AI as a formal breach category for the first time, finding that organizations with high shadow AI involvement incurred $670,000 in additional breach costs and took a median of 247 days to detect the incident [22].
- On May 7, 2026, CB Financial Services filed what appears to be the first SEC Form 8-K triggered by unauthorized employee AI use rather than a cyberattack — suggesting that data sensitivity alone may be sufficient to trigger materiality disclosure, independent of operational disruption [45].
- The EU AI Act's binding enforcement of high-risk AI system obligations begins August 2, 2026; an incomplete AI system inventory is not merely an IT governance gap at that date but a legal violation carrying penalties up to €15 million or 3% of global annual turnover [47].
- When organizations provision sanctioned AI tools, unauthorized use drops by 89%, suggesting that supply-side governance — not just detection and blocking — is the highest-leverage response [1].

Background

Shadow AI refers to the use of generative AI tools, models, services, and agents within an organization that operate outside of IT knowledge, procurement approval, or security oversight. The concept extends the phenomenon of shadow IT — the adoption of productivity tools without IT authorization — into a domain where the consequences of unmanaged data flows reach beyond directory exposure into source

code, legal memoranda, merger discussions, patient records, or customer personally identifiable information. Unlike a misconfigured file-sharing application, a shadow AI submission typically leaves no enterprise audit trail, carries no contractual data protection, and offers limited or no practical recovery path once data has been submitted to a consumer service.

The scale of shadow AI adoption reflects a genuine productivity imperative that governance structures have not kept pace with. Microsoft and LinkedIn's Work Trend Index found that 78% of AI users bring their own tools to work — a phenomenon the industry now terms "BYOAI" — with adoption reaching 85% among Gen Z workers and 73% among older workers [6]. Gallup data shows that approximately 45% of workers who use AI at work do so without informing their managers [16]. The productivity gains employees seek through this behavior are real, and the governance challenge is to channel that demand into monitored alternatives rather than eliminate it. A 2025 survey of more than 12,000 white-collar employees found that 60.2% had used AI tools at work, but only 18.5% were aware of any official company AI policy, indicating that most organizations have failed to communicate even where policies exist [52].

The generative AI tooling landscape has expanded faster than any governance mechanism can plausibly track. Netskope, which monitors enterprise traffic across its customer base, tracked more than 1,550 distinct generative AI SaaS applications by mid-2025 — up from approximately 317 at the start of 2025, reflecting nearly fivefold growth in a matter of months [8][9]. Over half of all current generative AI app adoption in enterprise environments is estimated to be shadow AI. The average enterprise, per Reco AI's 2025 State of Shadow AI Report, manages 490 SaaS applications of which only 47% are authorized, and hosts on average 1,200 unauthorized applications overall [2]. Unauthorized AI tools remain active for a median of 403 days before detection [27]. The governance gap is not closing on its own.

The Scale of the Problem

The breadth of shadow AI adoption has now been documented independently by security vendors, research firms, and industry analysts converging on consistent findings. Unseen Security's 2026 State of Shadow AI report found 80% of workers using unapproved AI tools, representing a 156% increase in shadow AI tool usage between 2023 and 2025 [1]. Gartner's survey of 302 cybersecurity leaders conducted between March and May 2025 found 69% of organizations either suspecting or having evidence that employees use prohibited public generative AI [4]. Gartner further predicts that by 2030 more than 40% of enterprises will experience security or compliance incidents linked to unauthorized shadow AI [5]. These figures come primarily from vendor telemetry and vendor-sponsored research, which reflects the customer populations of organizations already investing in security monitoring; the actual rates in organizations without deployed monitoring tools may be higher [1][2][14].

The account-level data is particularly instructive. LayerX's Enterprise GenAI Security Report 2025 found that 71% of all connections to generative AI tools are made via personal, non-corporate accounts, and 58% of the connections made through corporate accounts bypass Single Sign-On [11]. Harmonic Security's analysis of 22 million enterprise AI prompts found that 73.8% of ChatGPT enterprise usage occurs on personal rather than corporate-licensed accounts [37]. Microsoft Entra telemetry reports that 67% of employees access generative AI tools via personal accounts [19]. The implication is that conventional enterprise controls — SSO enforcement, CASB policies keyed to corporate credentials, DLP rules applied to managed endpoints — miss the large majority of AI interactions occurring in any given organization.

The data exposure surface is correspondingly vast. LayerX found that 77% of employees paste data into generative AI prompts and that 82% of those paste events occur through unmanaged personal accounts, placing the resulting data entirely outside enterprise DLP visibility [12][28]. Forty percent of files uploaded to generative AI tools contain personally identifiable information or payment card data [28]. Reco AI found that 86% of organizations lack visibility into how data flows to and from AI tools [2]. Harmonic Security's prompt analysis identified that six applications account for 92.6% of sensitive data exposure risk, with source code, legal documents, and financial data comprising 74.5% of exposed content [37]. Netskope observed that the volume of data sent to generative AI apps grew more than 30-fold in a single year [8], and that the average organization now experiences 223 incidents per month involving users sending sensitive data to AI applications — double the prior-year rate, with total data volumes reaching approximately 8.2 GB per organization per month [9].

The agent dimension compounds the problem in ways conventional shadow IT frameworks did not anticipate. Gartner projected in April 2026 that by 2028 the average Fortune 500 enterprise will have more than 150,000 AI agents in use, up from fewer than 15 in 2025 [56]. Eighty-two percent of CIOs report that employees are creating AI agents and apps faster than IT can govern them [57]. CSA's own April 2026 research found that 53% of organizations have already experienced AI agents exceeding their intended permissions [43]. Unlike a shadow SaaS application that passively receives data, a shadow AI agent actively initiates connections to external services, executes code, and may persist access credentials — expanding the attack surface from data exfiltration into active compromise.

Security Analysis

Data exfiltration through unmonitored channels represents the most directly documented shadow AI risk in current telemetry. When an employee submits source code, a contract draft, a financial model, or a client list to an external generative AI service through a personal account, the enterprise has lost data control in a meaningful and often irreversible sense. The data may be retained by the provider, used for

model training, exposed in a subsequent breach of the provider's infrastructure, or accessed by a threat actor who compromises the employee's personal account. LayerX's 2025 Browser Security Report found that generative AI now accounts for 32% of all corporate-to-personal data movement, making it the single largest data exfiltration channel in the enterprise browser — ahead of personal email and personal cloud storage [12]. Zscaler's ThreatLabz analysis of 536.5 billion AI and machine learning transactions across its cloud found a 3,464.6% year-over-year increase in enterprise AI and ML transactions, with enterprises ultimately choosing to block approximately 60% of this traffic on risk grounds [25][26].

The financial sector's early institutional response illustrates how rapidly the risk calculus shifted. JPMorgan Chase, Citigroup, Goldman Sachs, Bank of America, Deutsche Bank, and Wells Fargo all banned or restricted employee use of ChatGPT during 2023, driven by third-party software vetting requirements and client data protection obligations [32][33]. Samsung's semiconductor division discovered three separate confidential data disclosures to ChatGPT within twenty days of permitting employee access — source code, optimization code, and a recording of an internal meeting converted to text — before banning the tool company-wide [29][30]. These early cases established a behavioral pattern that subsequent telemetry has confirmed at scale: employees regularly submit sensitive data to AI tools when it is relevant to their work, even absent explicit intent to breach policy.

Prompt injection and supply chain risks represent a second, more technically sophisticated threat vector that shadow AI significantly amplifies. When employees connect unauthorized AI tools or agents to corporate systems without IT review, those tools may contain or be susceptible to malicious instructions embedded in documents, emails, or external data sources. EchoLeak (CVE-2025-32711, CVSS 9.3), discovered by Aim Security, demonstrated a zero-click prompt injection vulnerability in Microsoft Copilot through which hidden instructions in emails or spreadsheets could silently exfiltrate corporate emails to attacker-controlled servers [3]. A supply chain attack analyzed in August 2025 — attributed to threat actor UNC6395 in one threat intelligence analysis — used stolen OAuth tokens from a SaaS integration to access customer environments across more than 700 organizations, including Drift and Salesforce platforms, illustrating how AI-adjacent SaaS connections create lateral movement paths that IT organizations did not explicitly approve [21]. Reco AI identified three shadow AI applications — Jivrus Technologies, Happytalk, and Stability AI — in active enterprise use that received failing security grades for lacking basic controls such as encryption and multi-factor authentication [27].

The compliance dimension has moved from theoretical to immediate. IBM's 2025 Cost of a Data Breach Report, covering 600 organizations breached between March 2024 and February 2025, was the first edition to formally categorize shadow AI as a distinct breach risk factor. Organizations with high shadow AI involvement faced average total breach costs of $4.63 million and a median of 247 days to identify the incident [22][23][24]. Of the 13% of organizations that experienced a breach involving AI models or applications, 97% lacked proper AI access controls [22]. Sixty-five percent of these incidents resulted in

PII exposure and 40% in intellectual property theft. In healthcare specifically, 86% of healthcare IT executives reported shadow IT and AI incidents in 2025, and shadow AI added an estimated $200,000 to global average breach costs in affected organizations [49][50]. The sector has ranked as the most expensive industry for data breaches for fourteen consecutive years [23][49][50].

The CISA acting director incident from August 2025 illustrates that behavioral risk does not correlate with technical sophistication or seniority. Madhu Gottumukkala, who had personally requested and received DHS authorization to use ChatGPT, nonetheless uploaded four documents marked "For Official Use Only" to the public consumer version of the service rather than an enterprise-governed instance, triggering multiple alerts from DHS cybersecurity sensors [34][35]. The incident drew attention precisely because it involved the head of the nation's lead civilian cybersecurity agency engaging in the same behavior that CISA advises enterprises to prevent — a demonstration that approved access and governed use are two distinct states.

Recommendations

Immediate Actions

Establishing basic shadow AI governance requires action across several fronts simultaneously, beginning with accurate inventory. Organizations should conduct an AI asset discovery sweep using network telemetry, browser extension scanning, and SaaS API integration auditing to establish a current-state inventory of all AI tools in use, including those accessed via personal accounts. Tools such as Nudge Security, Obsidian Security, CrowdStrike's Shadow AI Visibility Service, and Microsoft Entra's shadow AI discovery capabilities now offer visibility across these vectors [58][59][60].

Inventory alone is insufficient without policy and provisioned alternatives. Organizations that have not published an AI Acceptable Use Policy should do so immediately; those with existing policies should verify active distribution to all employees — survey evidence showing only 18.5% of employees aware of any company AI policy suggests that policy existence alone, without active communication, may have negligible effect on employee behavior [52]. ISACA's customizable AI Acceptable Use Policy Template and NIST AI RMF-aligned templates from Strac.io and Tenable provide actionable starting frameworks [63][64][65]. Equally important is provisioning approved alternatives for the most common shadow AI use cases: when sanctioned tools are made available, unauthorized use drops by 89% [1]. The goal of governance is redirection of productive behavior into monitored channels, not prohibition of it.

Organizations should also apply a data classification matrix to AI tool permissions without delay. Public data may be processed through any approved tool; internal data requires sanctioned enterprise tools with audit logging; confidential data requires enterprise-contracted tools with data processing agreements; and restricted data — including PHI, PCI data, and trade secrets — requires explicit written approval before any AI processing. For organizations in the financial sector or subject to HIPAA, an immediate review is warranted to assess whether past employee AI tool submissions involving customer data constitute reportable incidents under applicable rules. CB Financial Services' May 7, 2026 SEC Form 8-K signals that regulators and companies may treat data sensitivity as independently material, even absent operational impact [45][46].

Short-Term Mitigations

At the network and application layer, deploying a Cloud Access Security Broker or Secure Service Edge solution with generative AI policy capabilities enables organizations to enforce approved-tool lists, block personal-account authentication to AI services on corporate devices, and apply inline data loss prevention to AI traffic. Major platforms with current generative AI governance capabilities include Netskope, Zscaler AI Protect, and Palo Alto Networks Prisma SASE 4.0, which covers more than 6,000 generative AI applications [8][38][39]. Conventional CASB inspection, however, misses AI interactions occurring through browser extensions, embedded AI features in SaaS applications, and AI tools accessed through personal accounts on managed devices. Browser-native solutions such as Harmonic Security, Prompt Security (now part of SentinelOne following the August 2025 acquisition), and Cato Networks' Enterprise Browser provide visibility at the interaction level that network-layer inspection cannot reach [37][40][58].

AI agent governance requires a separate discipline from standard application inventory. AI agents that operate autonomously, initiate external connections via Model Context Protocol or Agent-to-Agent Protocol, and persist credentials represent a qualitatively different risk from passive SaaS applications — one that conventional shadow IT frameworks were not designed to detect. Nudge Security's May 2026 browser-based agent discovery capability and Cisco AI Defense's MCP gateway function now address this emerging surface [41][42][58]. Organizations should also begin mapping all AI system deployments against EU AI Act Annex III high-risk classifications in preparation for the August 2, 2026 enforcement date. Approximately 40% of enterprise AI systems lack clear regulatory classification, and more than half of organizations have not established systematic AI inventories despite the Act's inventory requirement carrying no good-faith exception for organizations that simply failed to conduct one [47][48].

Strategic Considerations

The shadow AI governance gap is fundamentally an inventory problem before it is a policy problem. Every major applicable framework — NIST AI RMF's MAP function, EU AI Act Article inventory obligations, and ISACA's shadow AI auditing guidance — establishes a complete and continuously maintained AI system inventory as its foundational requirement [53][54][62]. Without inventory, policy cannot be applied and risk cannot be measured. Organizations should designate an AI asset owner role with explicit accountability for maintaining this inventory, analogous to the CISO's accountability for the vulnerability inventory, and treat gaps in that inventory as the primary governance metric rather than a secondary compliance checkbox.

Investment in AI governance infrastructure now reduces significantly greater cost exposure later. Gartner projects that AI governance spending will reach $492 million in 2026 and surpass $1 billion by 2030, driven primarily by regulatory compliance requirements globally [71]. CSA's December 2025 research found that organizations with mature AI governance programs are twice as likely to adopt agentic AI successfully and three times more likely to train staff on AI security tools — evidence that governance functions as a capability enabler rather than an operational constraint [44].

Traditional Zero Trust architectures were designed for human-initiated, north-south connections and do not natively address the egress patterns of AI agents that autonomously initiate connections to external MCP servers and third-party APIs at runtime. Extending Zero Trust controls to cover AI agent egress represents a significant architectural gap that emerging solutions are beginning to address, including Cisco AI Defense's MCP gateway, Cato Enterprise Browser's ZTNA extension, and Microsoft Entra Agent ID [41][55][59][60]. Organizations with high-risk AI agent deployments should evaluate these capabilities as a priority, recognizing that the control gap is architectural rather than merely configurational.

CSA Resource Alignment

This research note connects directly to several active Cloud Security Alliance frameworks and research programs. The CSA MAESTRO framework provides a structured threat modeling methodology applicable to shadow AI risks across all seven AI system layers, from infrastructure and data pipelines through agent orchestration and user interaction surfaces. CSA's AI Cloud Matrix (AICM), as a superset of CCM v4, provides the control families most directly applicable to shadow AI governance: data governance and classification controls, identity and access management for AI services, logging and auditability requirements, and vendor risk management for third-party AI providers. CSA has published

guidance on preventing unauthorized AI from circumventing compliance requirements [44], including dedicated coverage of shadow AI governance challenges [58] and the organizational risks of unmanaged AI adoption [66].

CSA's Zero Trust guidance, including the 2024 report on AI-induced shadow access [67], addresses the structural incompatibility between conventional Zero Trust architectures and the autonomous egress behavior of modern AI agents. CSA's April 2026 finding that 53% of organizations have already experienced AI agent scope violations underscores that agent governance is no longer a future-state problem but a present operational one [43]. Organizations seeking to operationalize the recommendations in this note should consult CSA's published AI Acceptable Use Policy guidance, the MAESTRO threat model, and CSA STAR program requirements for AI-aware vendor assessments.

References

[1] Unseen Security. "The State of Shadow AI 2026." Unseen Security, 2026.

[2] Reco AI. "2025 State of Shadow AI Report." Reco AI, 2025.

[3] The Hacker News. "The Hidden Security Risks of Shadow AI in Enterprises." The Hacker News, April 2026.

[4] Gartner. "Gartner Identifies Critical GenAI Blind Spots That CIOs Must Urgently Address." Gartner Newsroom, November 19, 2025.

[5] Infosecurity Magazine. "Gartner: 40% of Firms to Be Hit By Shadow AI Security Incidents." Infosecurity Magazine, 2025.

[6] Microsoft / LinkedIn. "Work Trend Index." Microsoft, 2025.

[7] Acuvity AI. "2025 State of AI Security Report." Acuvity AI, 2025.

[8] Netskope Threat Labs. "Cloud and Threat Report: Shadow AI and Agentic AI 2025." Netskope, 2025.

[9] Netskope. "Netskope Threat Labs: Shadow AI Risks Proliferate as GenAI Platforms and AI Agents See Rapid Adoption." Netskope Press Release, 2025.

[10] Cybersecurity Dive. "Risky shadow AI use remains widespread." Cybersecurity Dive, 2025.

[11] LayerX Security. "Enterprise GenAI Security Report 2025: Exposing Hidden AI Security Blind Spots." LayerX Security, 2025.

[12] LayerX Security. "AI Is Now the #1 Data Exfiltration Vector in the Enterprise." LayerX Security, 2025.

[13] Help Net Security. "89% of enterprise AI usage is invisible to the organization." Help Net Security, March 6, 2025.

[14] Vectra AI. "Shadow AI Explained: Risks, Costs, and Enterprise Governance." Vectra AI, 2025–2026.

[15] UpGuard. "The State of Shadow AI." UpGuard, 2025.

[16] Gallup. "AI Use at Work Has Nearly Doubled in Two Years." Gallup, 2025.

[17] ISACA. "From Shadow IT to Shadow AI: Navigating the New Frontier of Enterprise Risk." ISACA Newsletter, Volume 19, 2025.

[18] SC Media. "Shadow AI expands attack surfaces beyond visibility." SC Media, 2025.

[19] Microsoft. "Microsoft Entra Shadow AI Discovery." Microsoft Learn, 2025–2026.

[20] Cybersecurity Dive. "Shadow AI is widespread — and executives use it the most." Cybersecurity Dive, 2025.

[21] Reco AI. "AI & Cloud Security Breaches: 2025 Year in Review." Reco AI, 2025.

[22] IBM Newsroom. "IBM Report: 13% Of Organizations Reported Breaches Of AI Models Or Applications, 97% Of Which Reported Lacking Proper AI Access Controls." IBM Newsroom, July 30, 2025.

[23] IBM / Ponemon Institute. "Cost of a Data Breach Report 2025." IBM, July 2025.

[24] Cybersecurity Dive. "Shadow AI increases cost of data breaches, report finds." Cybersecurity Dive, July 2025.

[25] Zscaler ThreatLabz. "Insights from the ThreatLabz 2025 Data@Risk Report." Zscaler, 2025.

[26] Zscaler. "New Zscaler AI Security Report: Over 3,000% Surge in Enterprise AI/ML Use." Zscaler Investor Relations, 2025.

[27] Reco AI. "State of Shadow AI Report 2025: Key Findings." Reco AI Blog, 2025.

[28] The Hacker News. "AI Is Already the #1 Data Exfiltration Channel in the Enterprise." The Hacker News, October 2025.

[29] CIO Dive. "Samsung employees leaked corporate data in ChatGPT." CIO Dive, April 2023.

[30] AuthenTech. "The 2023 Samsung ChatGPT Incident, Explained for Security Teams." AuthenTech, 2023.

[31] AI Incident Database. "Incident 768: ChatGPT Implicated in Samsung Data Leak of Source Code and Meeting Notes." AI Incident Database, 2023.

[32] Yahoo Finance / Fortune. "Apple, Goldman Sachs, Samsung among growing list of companies banning employees from using ChatGPT at work." Fortune, May 2023.

[33] Tech.co. "Wall Street Banks Are Banning Employee Use of AI Bot ChatGPT." Tech.co, 2023.

[34] CSO Online. "CISA chief uploaded sensitive government files to public ChatGPT." CSO Online, August 2025.

[35] Rockfort. "CISA ChatGPT Data Leak: How America's Top Cybersecurity Official Exposed the #1 Enterprise AI Risk." Rockfort, 2025.

[36] OECD AI Incident Database. "Widespread Corporate Data Leaks via ChatGPT Use by Employees." OECD, October 2025.

[37] Harmonic Security. "What 22 Million Enterprise AI Prompts Reveal About Shadow AI in 2025." Harmonic Security, 2025.

[38] Zscaler. "Zscaler Unveils New Innovations to Secure Enterprise AI Adoption." Zscaler Investor Relations, 2025.

[39] Palo Alto Networks. "Prisma SASE 4.0: Powering the AI-Ready Enterprise." Palo Alto Networks Blog, September 2025.

[40] AppSecSanta. "Prompt Security 2026: GenAI Firewall by SentinelOne." AppSecSanta, 2026.

[41] Cisco. "Cisco Reimagines Security for the Agentic Workforce." Cisco Newsroom, March 2026.

[42] SDxCentral. "Cato Networks Launches Enterprise Browser to Tackle Shadow AI Threats." SDxCentral, April 2026.

[43] Cloud Security Alliance. "More Than Half of Organizations Experience AI Agent Scope Violations." CSA Press Release, April 16, 2026.

[44] Cloud Security Alliance. "AI Governance: A Maturity Multiplier." CSA Blog, December 18, 2025.

[45] Wilson Sonsini. "Shadow AI Triggers First SEC Form 8-K for Unauthorized AI Use." Wilson Sonsini Goodrich & Rosati, May 11, 2026.

[46] CB Financial Services. "Form 8-K filed May 7, 2026." SEC EDGAR, May 7, 2026.

[47] European Union. "Regulation (EU) 2024/1689 – Artificial Intelligence Act." Official Journal of the European Union, 2024.

[48] Cloud Security Alliance Labs. "EU AI Act High-Risk Compliance Deadline." CSA Labs Research Note, 2025.

[49] TechTarget HealthTech Security. "Shadow AI in Healthcare: The Hidden Risk to Data Security." TechTarget, 2025.

[50] Wolters Kluwer. "Shadow AI: A Hidden Risk to Healthcare." Wolters Kluwer, 2025.

[51] Cloud Security Alliance. "Is Shadow AI Putting Your Compliance at Risk?" CSA Blog, October 24, 2024.

[52] Dev|Journal. "Enterprise AI Governance 2026: Shadow AI Growth and Policy Failure." Dev|Journal, May 13, 2026. (Secondary source reporting primary survey research of 12,000 white-collar employees; primary publication not yet identified.)

[53] NIST. "AI Risk Management Framework." NIST, January 2023.

[54] NIST. "AI RMF Playbook." NIST AI Resource Center, 2023.

[55] Zentera. "Closing the Shadow AI Gap: Why Traditional Zero Trust Is Not Enough." Zentera, 2025–2026.

[56] Gartner. "Gartner Identifies Six Steps to Manage Artificial Intelligence Agent Sprawl." Gartner Newsroom, April 28, 2026.

[57] 1Password. "The Enterprise AI Crisis: Unsanctioned Tools and Unenforced Policies." 1Password Blog, 2026.

[58] PR Newswire. "Nudge Security Becomes the First AI Security Solution to Discover Shadow AI Agents Beyond APIs." PR Newswire, May 2026.

[59] Obsidian Security. "Shadow AI Security: Discover, Govern, and Control Every AI Tool." Obsidian Security, 2025–2026.

[60] CrowdStrike. "Introducing the CrowdStrike Shadow AI Visibility Service." CrowdStrike Blog, 2025.

[61] CrowdStrike. "CrowdStrike Innovations Secure AI Agents and Govern Shadow AI." CrowdStrike Blog, 2025–2026.

[62] ISACA. "The Rise of Shadow AI: Auditing Unauthorized AI Tools in the Enterprise." ISACA, 2025.

[63] ISACA. "Artificial Intelligence Acceptable Use Policy Template." ISACA, 2025.

[64] Strac.io. "AI Acceptable Use Policy Template: Free + Enforcement Guide." Strac.io, 2026.

[65] Tenable. "AI Acceptable Use Policy: Complete Guide." Tenable, 2025.

[66] Cloud Security Alliance. "AI Gone Wild: Why Shadow AI Is Your IT Team's Worst Nightmare." CSA Blog, March 4, 2025.

[67] Cloud Security Alliance. "Cloud Security Alliance Paper Addresses Challenges of Implementing Zero Trust in Environments Where AI-Induced Shadow Access Is Prevalent." CSA Press Release, May 7, 2024.

[68] Anthropic. "Disrupting the first reported AI-orchestrated cyber espionage campaign." Anthropic, 2025.

[69] Orca Security. "2025 State of Cloud Security Report." Orca Security, 2025.

[70] The Hacker News. "89% of Enterprise GenAI Usage Is Invisible to Organizations." The Hacker News, February 2025.

[71] Gartner. "Global AI Regulations Fuel Billion-Dollar Market for AI Governance Platforms." Gartner Newsroom, February 17, 2026.