netskope

Cybersecurity INSIDERS

2026

Unified Data Security Report

Closing the AI-Era Data Protection Gap

Executive Summary

Most security teams know where sensitive data is stored. Very few know where it goes, or what it becomes. As data moves across more environments, changes form, and enters AI workflows, security teams lose the context needed to enforce policy, investigate exposure, and prove what happened.

Cybersecurity Insiders surveyed 1,064 cybersecurity practitioners to measure the data protection gap between where sensitive data moves and where security can follow.

Key Findings:

- **Fragmented tools, fragmented protection:** 58% of organizations operate eleven or more data security tools, yet only 7% describe their stack as fully unified. 60% say their approach is moderately or highly fragmented. The issue is not lack of coverage but a lack of coordination.
- **Visibility and response weaken as data moves:** Confidence in tracing sensitive data drops from 25% for identifying who accessed a document to 8% for tracing AI-generated content back to its source. Only 7% can track data moving between applications in real time.
- **Data that changes form escapes detection:** Only 9% can recognize sensitive data reliably after it has been modified. 17% rely on exact-match detection that fails when content is copied, summarized, or rewritten. The file may never leave in its original form, but the sensitive content still does.
- **AI widens the data security gap:** 98% now use AI. 67% maintain some form of AI policy, but only 14% enforce through inline controls and just 8% enforce consistently in AI environments. 20% already embed AI in business-critical workflows, yet only 7% are confident that sensitive data is not flowing uncontrolled into AI.
- **Fragmented evidence creates regulatory exposure:** Only 12% can quickly produce a comprehensive chain of custody when regulators, auditors, or legal teams ask for evidence. 10% struggle to produce sufficient evidence at all. When evidence lives across disconnected systems, formats, and retention windows, every regulatory response becomes a reconstruction project.
- **The market is moving toward unified data security:** Protection that extends beyond data location to preserve context as it moves, changes form, and is reused. 79% of organizations have data security changes underway or planned within 12 months, and 72% expect investment to increase, concentrated around the same gaps this report measures: visibility, enforcement, integration, automation, and governance. Most organizations already have tools. What they lack is the coordination between them.

Together, these findings point to an AI-era data protection gap: Security teams have invested in controls, but context, policy, and evidence do not consistently follow sensitive data as it moves, changes form, and enters AI workflows. Closing the gap requires unified data security that connects discovery, classification, lineage, enforcement, investigation, and evidence across the environments where sensitive data is used.

Over-Tooled, Under-Coordinated

Organizations are investing heavily in data security tooling, from data loss prevention (DLP) and cloud access security broker (CASB) to endpoint, cloud, and SaaS controls. The problem is that investment has produced coverage without coordination. 58% operate eleven or more separate data security tools, yet only 7% describe their stack as fully unified, with shared visibility and policy orchestration across environments. 60% say their approach is moderately or highly fragmented, and 23% report their tools share almost no context between them.

The problem is practical: 45% of security teams query six or more separate systems per investigation, and 16% query more than ten. Each system carries its own log format, retention window, event schema, and level of detail. In many environments, stitching together logs, alerts, and data-movement events becomes the investigation. Analysis cannot occur until after the assembly work is done, and a single external-sharing incident may require CASB access logs, endpoint DLP events, email metadata, and IdP activity, each with its own timestamp format, retention window, and export process. The analyst is building the timeline before they can assess the exposure.

When that assembly takes days, the organization is already behind: exposure remains uncertain, containment decisions slow down, and regulatory notification windows start to close before the team fully understands what happened. Teams that investigate in hours instead of days usually have one thing in common, evidence is connected before the incident begins, so analysts spend less time assembling the record and more time assessing exposure.

When practitioners name what fragmentation costs them, three problems dominate: manual effort to correlate data across systems during investigations (44%), inconsistent policy enforcement across tools and environments (42%), and visibility gaps where tools lack shared context (40%). All three point to the same architectural problem: Data security has been assembled tool by tool, without a shared layer for data context, policy decisions, and evidence. A practical illustration is to count how many consoles a SOC analyst opens during a routine investigation. If the answer keeps growing, the architecture is forcing analysts to become the integration layer.

Tool Sprawl Without Integration

How many separate security tools or platforms does your organization use to monitor, govern, or protect sensitive data across environments?

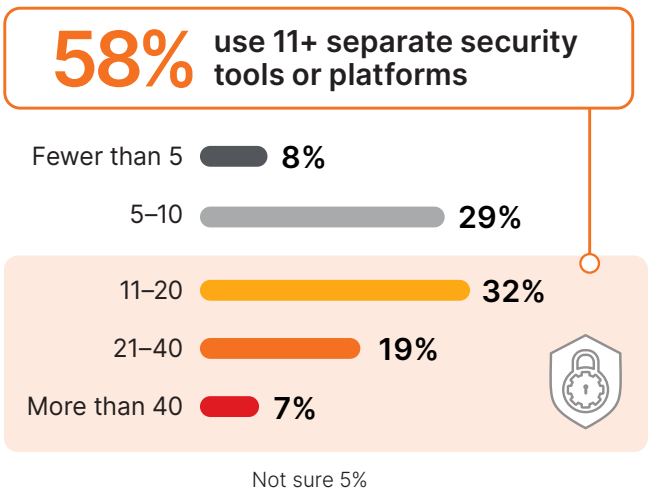

Describe your organization's approach to monitoring and protecting sensitive data across web, SaaS, cloud, endpoint, private app, and AI environments?

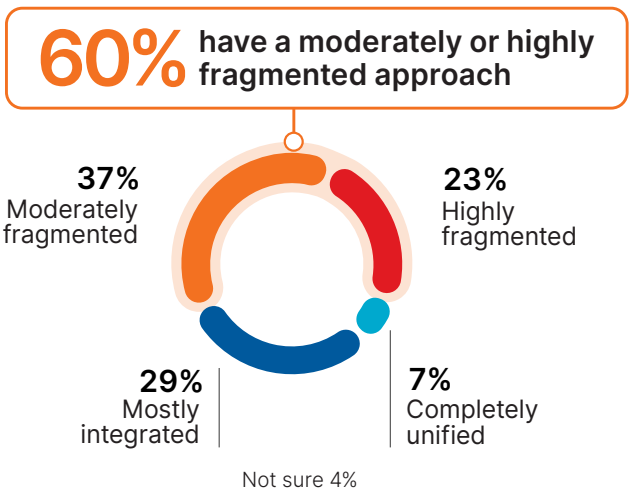

> The stronger pattern is to connect visibility, classification, policy, and evidence across the environments where sensitive data is stored, used, copied, and transformed, then build outward from that shared operating layer.

Stretched Teams, Growing Alert Queues

The teams carrying the investigation burden are already stretched thin. Only 9% of organizations have the capacity to investigate nearly all sensitive-data alerts, while 56% investigate half or fewer.

Capacity is the limiting factor: Only 10% have extensive automation for routine detection, data classification, alert enrichment, and response. 47% describe their automation as limited or minimal, with most work still done by hand.

When tools are fragmented, and workflows stay manual, the alert queue grows into a backlog that security teams were never resourced to clear manually.

Most Teams Cannot Investigate Every Sensitive-Data Alert

What percentage of security alerts related to sensitive data does your team fully investigate?

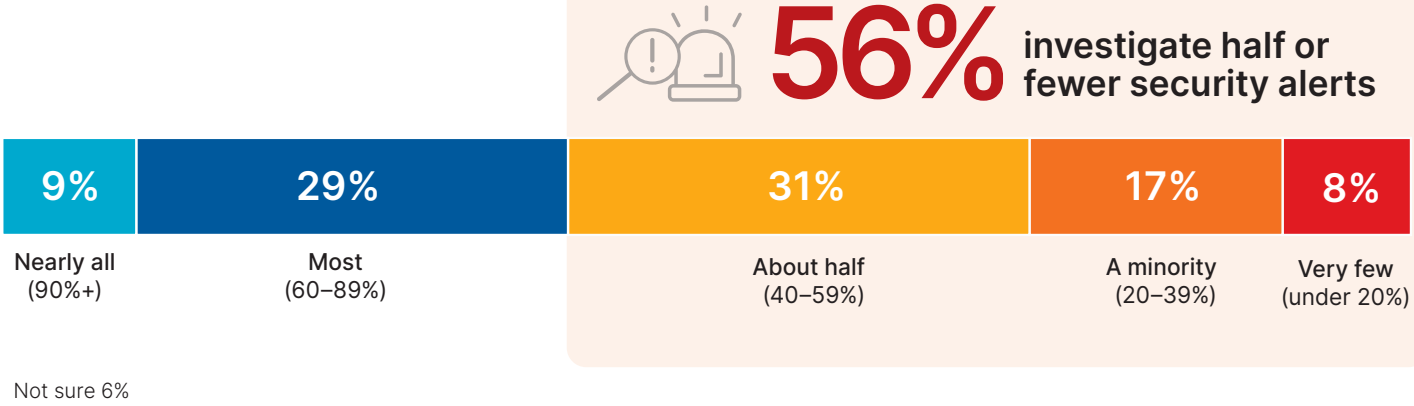

Not sure 6%

Automation remains limited

Only 10% have extensive automation

47% have limited or minimal automation

Teams ahead of this curve are reducing manual overhead through automation and consolidation: fewer systems to query, fewer logs to reconcile, and more analyst capacity for the alerts that matter.

Confidence Falls as Data Moves and Changes

Slow investigations are only a symptom. The underlying cause is that visibility into sensitive data weakens at every step, from knowing who accessed it, to tracking where it moved, to tracing how AI transformed or reused it.

Only 25% of organizations are very confident they can identify who accessed a specific sensitive document. Confidence drops to 16% for determining whether sensitive data reached AI tools, to 14% for tracing how a file moved from origin to its current location, and just 8% for tracing AI-generated content back to the sensitive data it drew from.

Most tools can answer who opened a document because that is an access-log question. Tracing AI-generated content back to its source, however, requires provenance, not just access logging. It requires classification, lineage, and policy-event context that stays with sensitive data as it moves across environments, users, applications, and actions.

That context shows where the data originated, how it changed, and which safeguards applied along the way.

Provenance Is Harder to Prove Than Access

How confident are you in your organization's ability to do the following today?
[Percent reporting "very confident"]

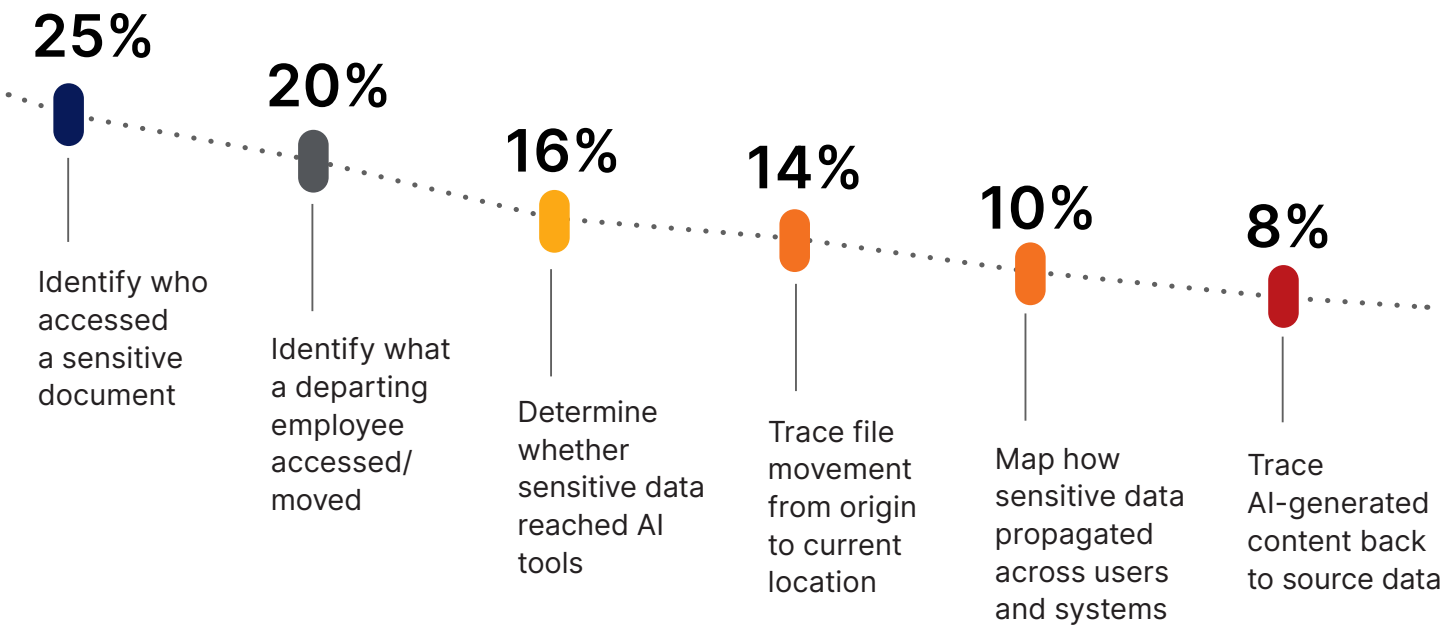

Where investigations move fast, security teams have already built persistent context around the data itself, so they trace sensitive content across users, applications, actions, and AI workflows without starting over each time.

Visibility Breaks at the Boundaries

AI tools, private applications, and cloud infrastructure are only part of the picture. Visibility weakens unevenly across channels, and some of the weakest spots are the ones you might assume are most covered.

Email has the strongest visibility at 36%, while AI tools have the weakest at 10%. More surprising, private and data center applications sit at just 15%, barely above AI despite holding some of the enterprise's most sensitive operational data.

Cross-application movement exposes the boundary problem. Only 7% can track sensitive data moving between applications in real time. A single workflow that moves data from SaaS to a private application to an AI assistant can cross three visibility boundaries. Most teams lose context before the path is complete.

Strong Visibility Drops Outside Email

How would you rate your organization's visibility into sensitive data activity in each of the following environments?
[Percent reporting "strong visibility"]

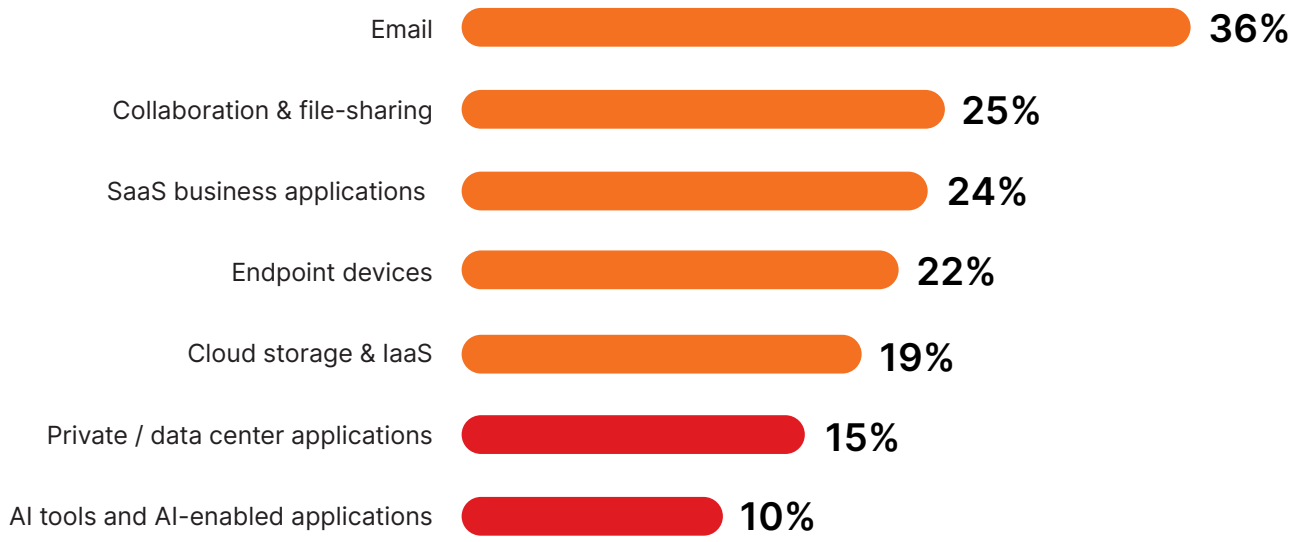

Only 7% can track sensitive data moving between applications in real time

Early movers are extending classification and inline DLP beyond traditional CASB and next-gen secure web gateway (SWG) control points into private applications, AI tools, and cross-application workflows, where context is most likely to break.

When Data Changes Form

Visibility across environments is only half the challenge. The other half is tracking sensitive data as it transforms, is renamed, reformatted, copied into new documents, pasted into prompts, summarized by AI, or reused in generated content.

Only 9% say they can recognize sensitive data reliably after it has been modified. 17% rely on exact-match detection or fingerprinting, which can fail when small changes alter the file or content pattern. A file containing customer names, account details, and contract terms may be blocked from upload to a personal cloud account, then reappear minutes later as copied text, an AI-generated account summary, or a pasted excerpt in an external workspace. The file never leaves in its original form, but the sensitive content still does.

Cross-channel persistence is equally weak. When a risky action is blocked in one channel, only 7% can detect whether the same user tried a different channel in real time. Only 7% can trace sensitive data received from vendors or partners as it moves through internal systems. 31% lose visibility entirely once third-party data enters their environment.

Modified Data Often Escapes Detection

How effectively can your current data protection controls detect and prevent sensitive data exposure when data has been modified — for example, when files are renamed, reformatted, partially copied, or pasted into new documents?

Not sure 5%

MODIFIED-DATA TEST
Take a file containing customer names, account details, or contract terms. Copy, summarize, or rewrite it with AI. Then test whether classification and policy still recognize the sensitive content after it changes form.

Lineage turns data movement into an evidence trail, showing where sensitive data originated, how it changed, where it moved, and whether policy followed it across each handoff. The programs that can follow data through transformation have wired lineage into classification, DLP, and investigation workflows, so policy and evidence travel with sensitive content across data at rest, in motion, and in use.

Real-Time Enforcement Remains the Exception

While visibility tells security teams where sensitive data goes, enforcement determines whether they can enact controls. Enforcement carries more operational risk than monitoring because it can block, redirect, coach, or warn users, and every false positive can disrupt real work.

Only 11% enforce data protection policies in real time across most environments. 36% automate enforcement in some environments but still rely on manual review in others. The remaining 49% depend on written policies with post-event monitoring, reactive investigation after the fact, or policies that exist on paper but are enforced inconsistently.

In practice, policy-on-paper means the rule may exist before the control does. A policy can say “do not paste customer data into AI tools,” while enforcement still depends on whether employees read it and follow it. When violations surface after the fact, the issue is no longer policy intent; it is the absence of inline control at the point of action.

Enforcement weakens further in environments where controls are least mature. Email is strongest at 42%, followed by web traffic at 36% and managed endpoints at 29%. Enforcement then falls across SaaS applications at 23%, private or data center applications at 18%, cloud infrastructure at 13%, AI tools at 8%, and unmanaged or personal devices at 6%.

Without persistent context, enforcement becomes blunt: block, allow, or review later. With context, it becomes governance: the ability to allow, block, coach, redirect, or investigate based on classification, user risk, destination, device posture, and the action being attempted.

Most Enforcement Is Partial or Reactive

How are data protection policies principally enforced across your organization's environments today?

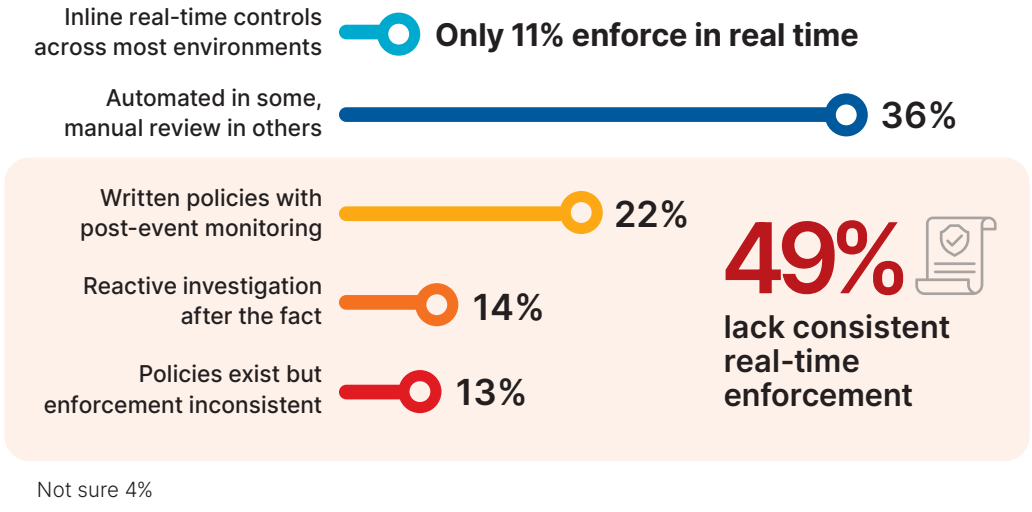

Not sure 4%

How consistently can your organization enforce data protection policies across the following environments? [Percent reporting “consistently enforced”]

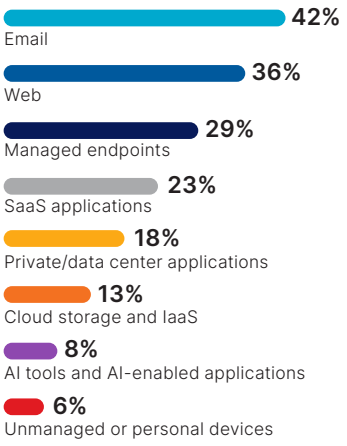

Programs closing this gap are moving from policy-on-paper to inline enforcement that applies consistent rules across channels, devices, applications, actions, and AI interactions.

What Can Actually Be Blocked

Even organizations that have moved to inline enforcement face a coverage problem. Enforcement still works best where data behaves like a file. It is weaker where sensitive content moves through prompts, copy-paste, generated responses, and agent-driven submissions.

File uploads to cloud services are the most controlled action at 53%. Sharing to external recipients follows at 49%, and downloads from cloud or SaaS at 43%. But only 25% can consistently govern the copy-paste of sensitive content into AI prompts or web forms, and only 20% can apply consistent policy to prompt and text submissions into AI tools. AI policy maturity does not automatically translate into action-level control.

A PRACTICAL TEST
Apply the same policy across upload, download, share, copy-paste, prompt submission, and generated response. If policy works for file movement but fails at the prompt or response layer, enforcement has not reached the actions where sensitive content now moves.

Enforcement Works Best Where Data Behaves Like a File

Which specific data-sharing actions can your organization control consistently across most environments today?

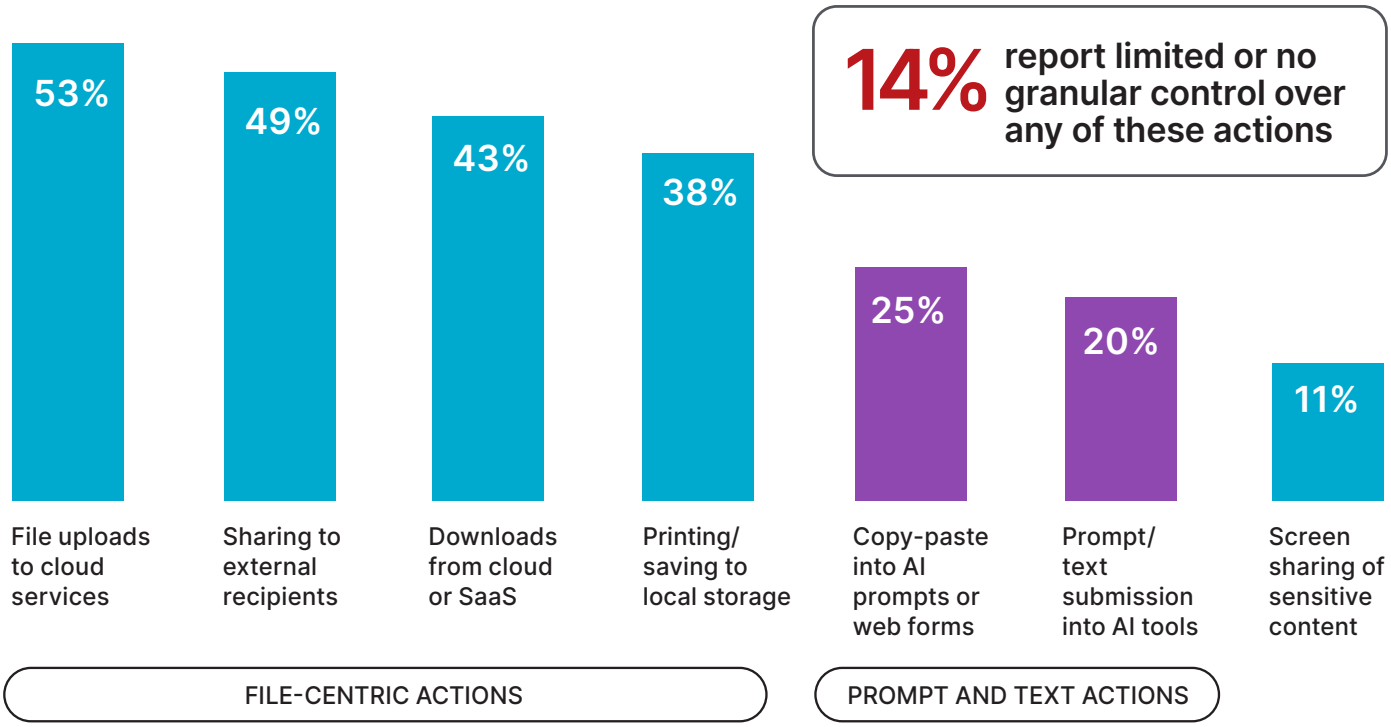

Programs making progress are extending inline DLP from file movement to action-level controls, while aligning policy ownership with the teams responsible for enforcement.

AI Broadens the Data Security Gap

Fragmentation, weak visibility, and inconsistent enforcement all existed before AI. AI expands those gaps because sensitive data now flows into prompts, copilots, embedded features, and agent-driven workflows faster than traditional controls can follow.

AI is now nearly universal: 98% of organizations use it in some form, and 20% already embed it in business-critical workflows such as procurement, customer support, and financial analysis, where it can touch customer, financial, operational, and contractual data. Yet only 7% are very confident that sensitive data is not entering AI workflows without appropriate inspection, policy, or enforcement. 39% suspect it is happening but lack the visibility to confirm it.

The hardest AI categories to detect are also the ones most likely to intersect with sensitive work. Personal AI tools used for work are the most difficult to see (51%) followed by AI assistants embedded in productivity suites (40%) and autonomous AI agents (37%). Only 10% can clearly distinguish approved from unapproved AI usage across most environments, including corporate versus personal AI instances.

Shadow AI Is Hardest to See Where Work Happens

In which environments is unmanaged or shadow usage most difficult to detect?
[Top responses shown]

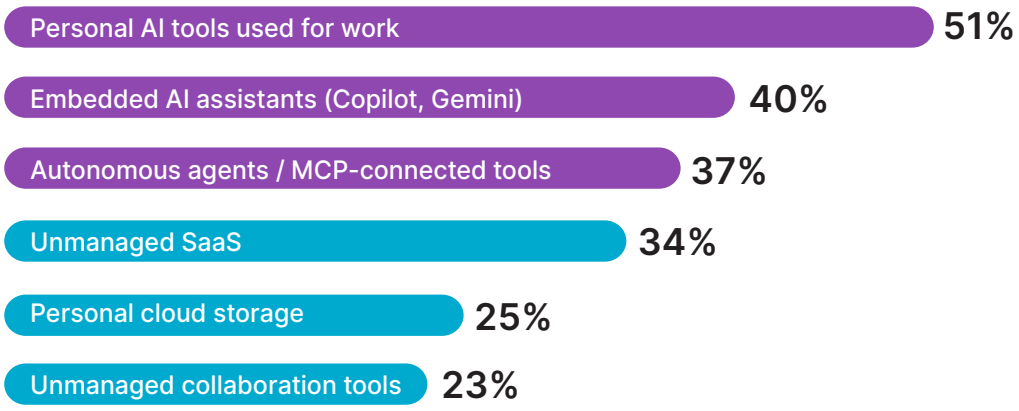

98%
use AI

But only 7% are confident sensitive data is not entering AI without appropriate controls

Only 10% can distinguish approved from unapproved AI tool usage

The organizations furthest along apply the same rigor to AI that they bring to email, web, cloud, SaaS, endpoints, and private applications: classification, DLP, prompt and response inspection, guardrails, and enforcement.

AI Adoption Without Enforcement

Writing an AI policy and enforcing it are different capabilities and investments. Most organizations have made the first but not the second, and the distance between the two is where sensitive data is most exposed: 98% use AI, 67% maintain a formal AI policy, 14% enforce through inline controls, and just 8% enforce consistently in AI environments.

The rest operate in less enforceable governance models: Written guidance, inconsistent enforcement, policies still under development, or no formal policy at all. In many organizations, AI governance exists on paper before it exists in the controls that determine what data users and AI systems can access.

AI spreads through everyday work: Employees find tools useful, departments embed them in workflows, and usage builds before security has consistent visibility or control. By the time the policy is written, the workflow it governs may already be operational. Enforcement requires architecture that can see AI use, classify the data involved, inspect prompts and responses, and apply policy in real time. Usage accelerates by default. Enforcement architecture must be built deliberately.

AI Policy Is Ahead of AI Enforcement

How extensively are AI tools and AI-enabled applications used across your organization today?

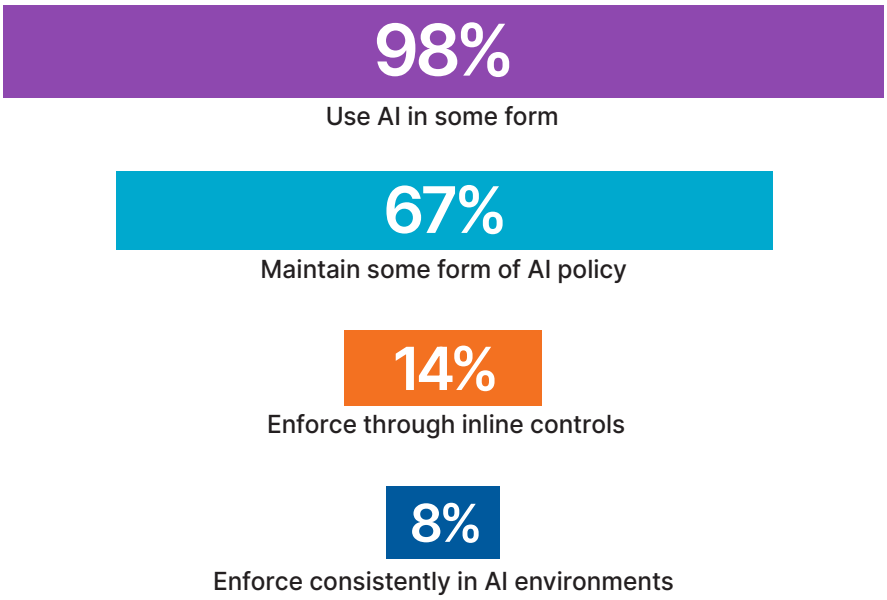

Does your organization currently maintain a formal policy governing employee and organizational use of AI tools?

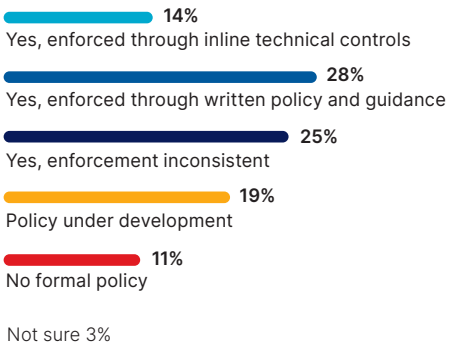

20%
already embed AI in business-critical workflows

Stronger AI governance starts with visibility into which tools employees use, what data enters them, what responses are generated, and where inline policy can be applied across sanctioned and unsanctioned AI instances.

Agents Become Data Users

Governing AI that employees use is one problem. Governing AI that acts on its own is another. Autonomous agents are non-human actors that can execute workflows, access data, call APIs, and make decisions without waiting for a human prompt. 37% already rank autonomous agents, including MCP-connected tools, among the hardest AI categories to detect.

Traditional access controls are built around human users: A person authenticates, receives permissions, and acts within those boundaries. Agents complicate that model because they operate through API calls, OAuth grants, service accounts, and delegated permissions. A sales operations agent, for example, may pull data from CRM records, support tickets, contracts, and customer emails, then write a summary to a shared workspace.

No employee manually moved each record, but the agent accessed, combined, transformed, and redistributed sensitive information across multiple systems. Most security architectures cannot follow that chain end-to-end.

Agentic AI deepens the exposure because agents operate at machine speed, at scale, and without many of the behavioral signals (the hesitation, the unusual login time, the repeated access attempts), security teams use to spot risky human actions.

Agents Access, Combine, Transform, and Redistribute Data

In which environments is unmanaged or shadow usage most difficult to detect?

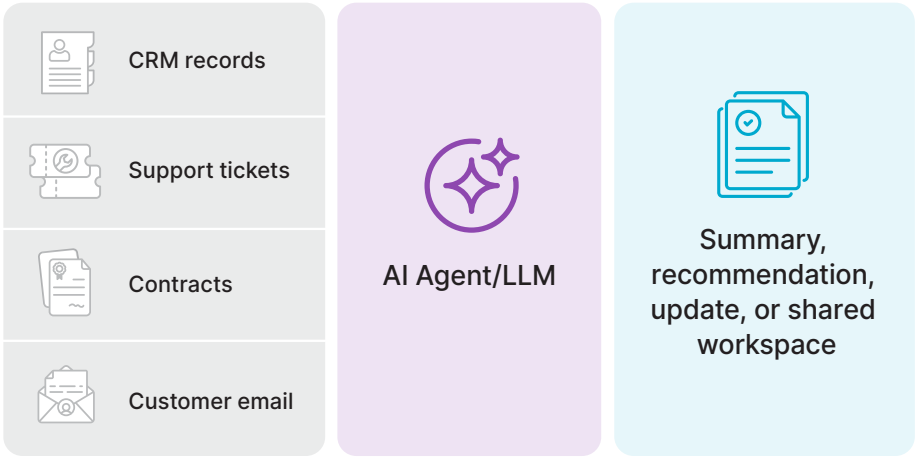

37% cite autonomous agents, including MCP-connected tools, among the hardest shadow AI categories to detect

How extensively are AI tools and AI-enabled applications used across your organization today?

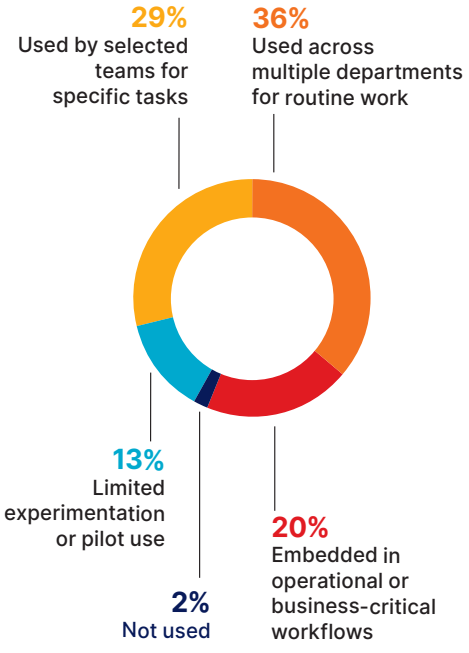

Mature agentic AI governance treats software agents as governed actors, with least-privilege permissions, monitored API activity, and policy applied at every data handoff.

Proving What Happened

Even well-defended organizations face incidents. When one involves sensitive data, the first question regulators, auditors, and legal teams ask is the same: Can you prove what happened, where the data went, who accessed it, and which protections were in place?

Visibility gaps obscure where data moved, enforcement gaps leave risky channels open, and uninvestigated alerts become missing evidence. Fragmentation slows response and weakens the organization's ability to prove what happened.

Only 12% can quickly produce a comprehensive, auditable chain of custody when regulators, auditors, or legal teams ask for evidence. 36% require significant manual effort across multiple systems. 10% struggle to produce sufficient evidence at all. The evidence often exists. The problem is that it lives across systems, formats, retention windows, and export processes that were never designed to answer the question quickly.

Data subject requests expose the same weakness. Only 9% can quickly locate all instances of an individual's data across systems, while 38% require extensive manual work. 9% cannot confirm that all instances have been found.

Audit readiness depends on the same visibility, classification, lineage, and policy-event records created through daily data protection workflows. When those foundations are fragmented, every regulatory response becomes a reconstruction project.

Evidence Is Still Assembled Manually

How well can your organization currently prove chain of custody for sensitive data during a regulatory audit or compliance review?

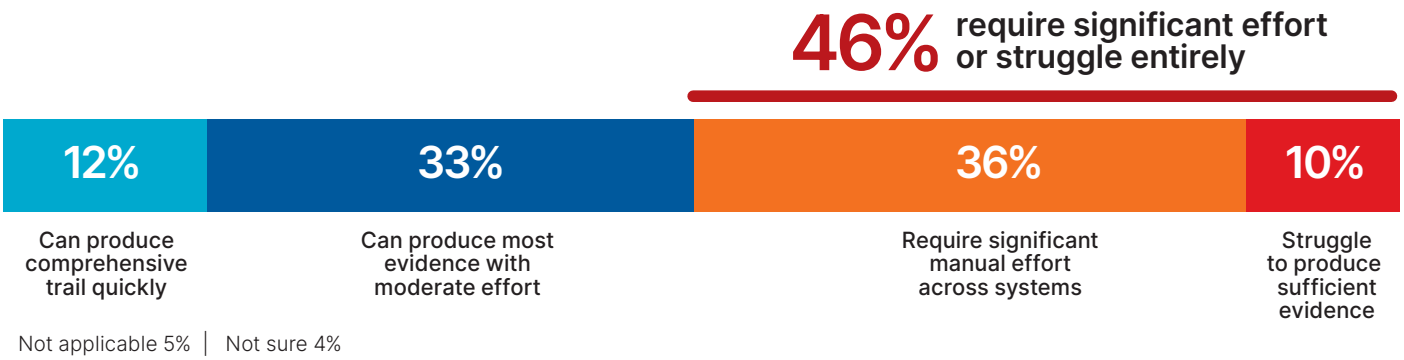

Only 9% can quickly fulfill data subject requests across systems

27% of organizations cannot reconstruct a sensitive data path before a 72-hour breach-notification deadline closes

The organizations that respond quickly have a unified evidence trail across environments, with chain-of-custody records generated through daily protection workflows instead of assembled after the fact.

Slow Evidence Collides With Regulatory Deadlines

Producing evidence eventually is not fast enough. Regulatory response depends on producing it before the notification window closes.

GDPR requires breach notification within 72 hours, and many US state breach notification laws impose tight reporting timelines of their own. That puts 27% of organizations in severe timing trouble: 19% take weeks or longer to reconstruct a sensitive data path, and 8% rarely have enough data to reconstruct the path at all. Even the 41% who reconstruct in days face a practical deadline problem. Breach response does not end when the security team assembles the technical record.

The legal, privacy, and communications teams still need time to assess exposure, determine notification obligations, and prepare messaging. A breach discovered Friday afternoon can trigger a 72-hour clock. By Monday morning, much of the window is gone. If the security team is still pulling logs from separate systems and reconciling timestamps, legal and privacy teams may have to make notification decisions before the full data path is clear.

Third-party data adds to the timing problem. When vendor or partner data is involved, the organization must trace provenance, downstream distribution, and whether the data was copied, transformed, or further shared. Only 7% can do this effectively, while 31% lose visibility once third-party data enters their environment.

Most Investigations Take Days or Longer

When investigating a potential data exposure incident, how long does it typically take your security team to determine where the data originated and how it moved across applications?

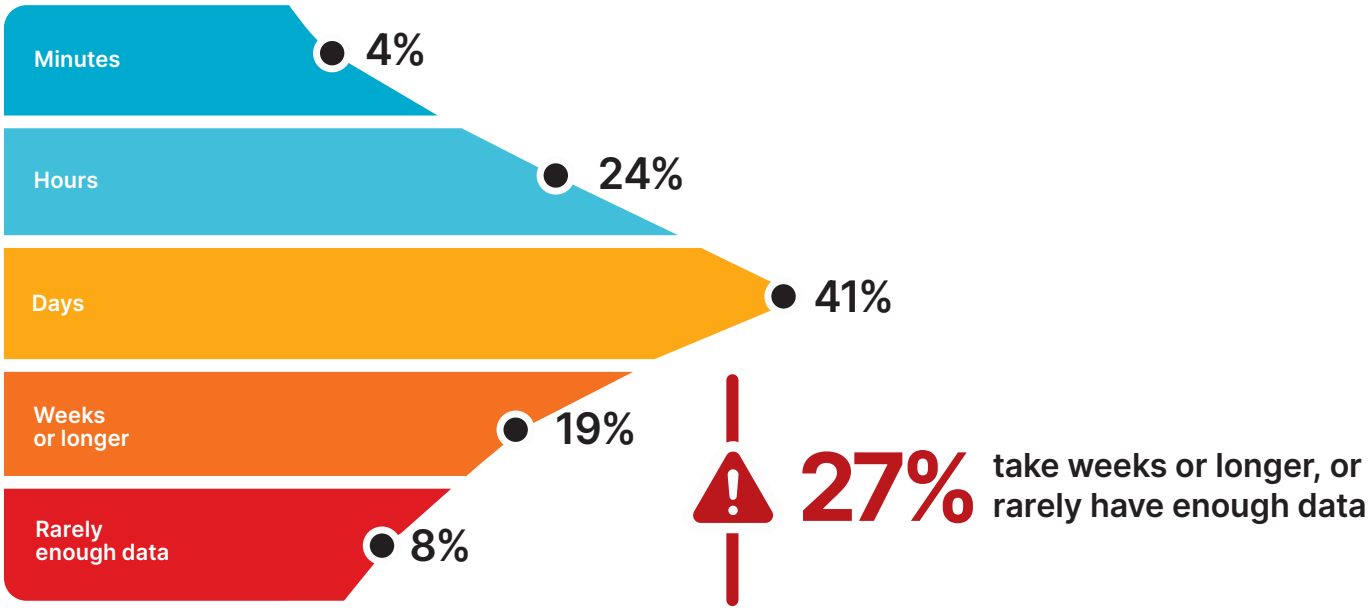

Not sure 4%

Organizations that meet regulatory timelines build persistent context into daily operations, so lineage, classification, policy events, and evidence trails already exist before the regulatory clock starts.

Why the Market Moves Toward Unified Data Security

The market is moving toward unified data security because the pressures are converging. Slow investigations, inconsistent enforcement, AI exposure, and regulatory deadlines all point to the same architectural requirement: shared data context and consistent policy across environments.

79% of organizations have data security changes underway or planned within 12 months, and 72% expect investment to increase. The capability gaps practitioners name most cluster at the top around visibility across environments and consistent enforcement, then thin out toward governance and compliance readiness. Investment priorities track the same gaps, led by visibility (44%), enforcement (39%), and automation (37%).

Integration and automation determine whether the other investments pay off. Visibility that cannot be automated stays slow; enforcement that cannot be integrated stays inconsistent. Integration complexity (35%) and budget constraints (31%) remain real barriers, but only 9% of organizations say consolidation is not a priority. And 75% are actively interested in AI-powered capabilities to unify systems and processes.

Friction is its own risk. 41% say security controls create enough friction that employees often turn to unapproved tools or workflows: personal AI accounts, unmanaged file shares, or copy-paste into channels DLP cannot see. Each workaround adds another stream of unmonitored data movement outside the controls meant to catch it.

Investment Priorities Follow the Gaps

Which capabilities will receive the greatest investment in your organization over the next 12–24 months to improve data security?

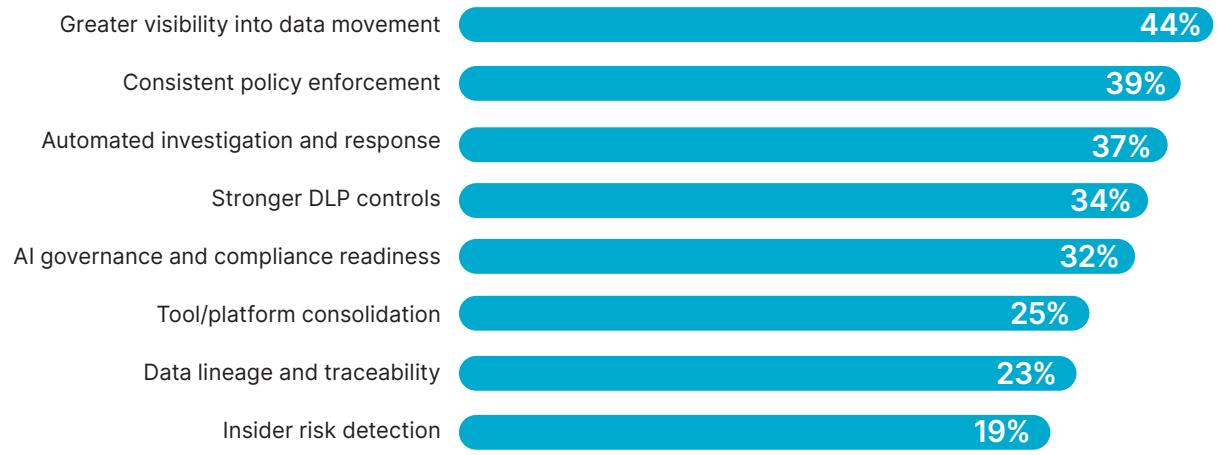

79% have changes underway or planned within 12 months

72% expect investment to increase

41% say controls create enough friction that employees turn to unapproved tools or workflows

> The more mature path to unified data security also protects the user experience, using coaching, adaptive controls, and risk-based enforcement where blunt blocking would drive workarounds. Consolidation that adds friction can recreate the same shadow workflows it was meant to eliminate.

Data Security Maturity Matrix

Data security maturity is rarely consistent across the program. Strong coverage in one area can coexist with weak visibility, uneven enforcement, slow investigations, or fragmented evidence elsewhere. This matrix translates the findings into five capability areas and three maturity tiers, helping teams identify where progress is currently blocked.

CAPABILITY	REACTIVE	MANAGED	ADAPTIVE
Visibility	Primary environments only. Context breaks across cloud, private apps, and AI. Data traced only to last known location.	Major environments monitored. Gaps remain in private apps, AI tools, and cross-app movement. Reconstruction still requires manual effort.	Persistent context follows data across environments, actions, and AI. Lineage tracks provenance, movement, transformation, and reuse.
Enforcement	Policy-driven and manually enforced. Inline controls limited. AI prompts, responses, and copy-paste weakly governed.	Automated enforcement in core environments. Coverage uneven across AI, cloud, unmanaged devices, and cross-channel workflows.	Inline enforcement across high-risk environments and actions. Policy covers uploads, prompts, responses, copy-paste, sharing, and agent workflows.
Detection & Response	Disconnected logs. Manual reconstruction. Multiple systems queried. Many alerts unresolved.	Defined workflows. Partial automation. Cross-environment correlation still manually assembled.	Connected evidence. Automated correlation. Rapid data-path reconstruction, alert prioritization, and response.
Governance & Compliance	Policies on paper. Fragmented ownership. Audit evidence manually collected. AI governance underdeveloped.	Clearer accountability. Some governance workflows tied to controls. AI policy exists, but enforcement varies.	Governance tied to enforcement. Classification, lineage, policy events, and evidence generated through daily workflows.
Integration & Automation	Tools, logs, policies, classifications, and evidence disconnected. Context reconciled manually.	Some shared context and workflow automation. Integration gaps persist across visibility, enforcement, and response.	Connected data security model. Discovery, classification, DSPM, DLP, lineage, enforcement, investigation, and evidence operate from shared context.

The pattern is clear: data security maturity is uneven, and integration often determines how far the other capabilities can advance. When context, policy, and evidence remain disconnected, visibility is partial, enforcement is uneven, and response is slow. The practical use of the matrix is to identify the lowest-maturity dimension first and start there, because that is where the broader data security program is most likely to stall.

Next Steps: Closing the Data Security Gaps

The maturity matrix shows where progress is blocked. Closing the gaps starts with the lowest-maturity dimension first, especially when fragmented integration limits visibility, enforcement, response, and governance. The sequence matters: connect the operating layer first, then extend visibility and enforcement to the places data changes form, and use that shared context to accelerate response and operationalize governance.

1

Connect the operating layer

Link the evidence, policies, and workflows that are fragmented today. Unify discovery, classification, DSPM, DLP, lineage, enforcement, and investigation so context follows sensitive data as it is stored, used, copied, and transformed. This foundation turns visibility, enforcement, and response into connected operating capabilities.

2

Strengthen visibility where data changes

Prioritize the places where sensitive data changes form or crosses control boundaries: AI tools, private applications, cloud infrastructure, and cross-application workflows. Build lineage into classification and DLP workflows so teams can see where data sits, how it moves, how it changes, and where it reappears.

3

Move policy into enforcement

Apply consistent rules across channels, devices, applications, actions, and AI interactions. Start where the gap is widest: prompt submission, copy-paste into AI tools, generated responses, and other actions where sensitive data can be exposed without the original file moving.

4

Accelerate investigation and response

Reduce the number of systems that analysts need to query during investigations, and automate routine detection, classification, enrichment, and triage. Once the operating layer is connected, the same shared context makes AI-assisted investigation viable because the evidence is already correlated before an incident begins. Faster response depends on giving analysts a connected view of the data path, the user or agent involved, the policy events triggered, and the actions taken.

5

Operationalize governance

Extend governance into the workflows where data actually moves. For AI, that means knowing which tools and embedded features employees use, what data enters them, what responses are generated, and whether policy follows the data. For agentic AI, it means treating software agents as governed actors: assigning least-privilege access, monitoring API activity and delegated permissions, inspecting the data they access or produce, and applying policy at every data handoff.

A coordinated data security operating model closes the AI-era data protection gap by keeping context with sensitive data as it moves, changes, and reappears. With that context in place, policy follows the data, investigations move faster, AI adoption becomes governable, and evidence is ready when business, legal, or regulatory teams ask what happened.

Methodology and Demographics

The 2026 Unified Data Security Report is based on a survey of 1,064 cybersecurity professionals, fielded in early 2026 by Cybersecurity Insiders. Respondents were screened for direct familiarity with their organization's data protection practices and the tools used to protect sensitive data, and span security operations, security and IT leadership, and the CISO office. The sample is weighted toward large enterprises: Every respondent works at an organization of 2,500 or more employees.

The survey examined how organizations see, govern, and prove the handling of sensitive data as it moves across email, SaaS, cloud, endpoints, private applications, and AI tools, and how protection holds as that data changes form and enters AI workflows. Reported at a 95% confidence level, the survey carries a margin of error of approximately ±3.0% for the full sample. For multi-select questions, percentages may exceed 100%.

RESPONDENT ROLE

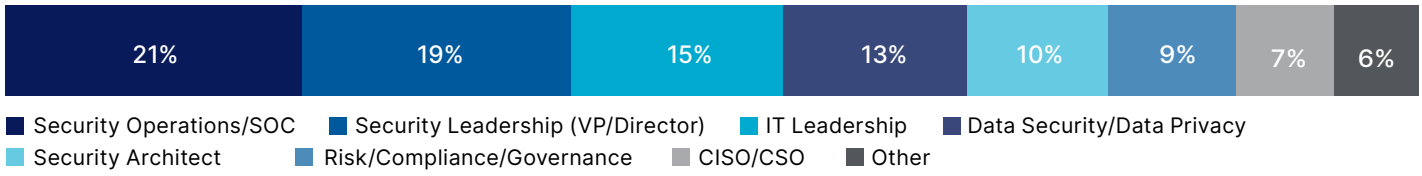

INDUSTRY

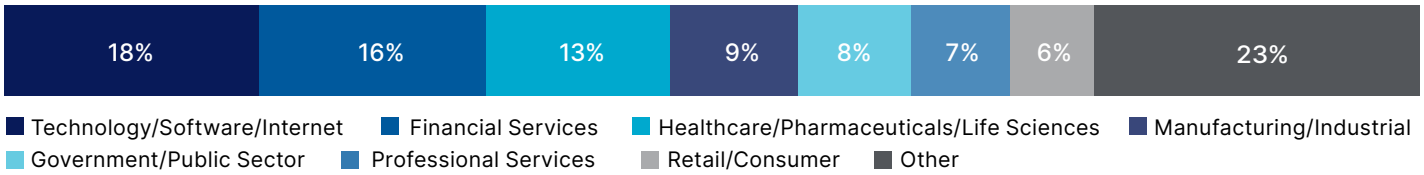

COMPANY SIZE

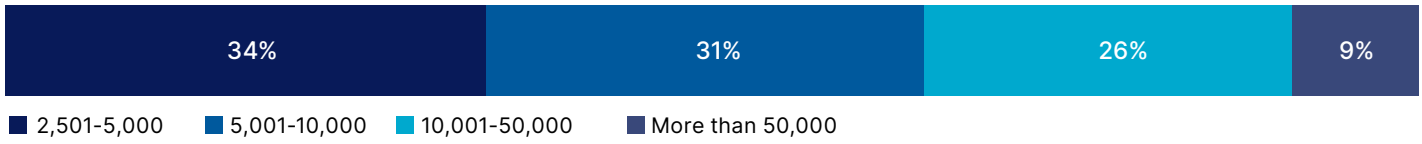

About Netskope

Netskope (NASDAQ: NTSK), a leader in modern security and networking for the cloud and AI era, addresses the needs of both security and networking teams by providing optimized access and real-time, context-based security for the AI ecosystem inclusive of agents, applications, tools, LLMs, people, devices, and data.

Thousands of customers, including more than 30 of the Fortune 100, trust the Netskope One platform, its Zero Trust Engine, and its powerful NewEdge network to reduce risk and gain full visibility and control over cloud, AI, SaaS, web, and private applications – providing security and accelerating performance without trade-offs.

Learn how at

netskope.com/datasecurity

Cybersecurity
INSIDERS

BENCHMARK YOUR SECURITY MATURITY

Independent cybersecurity research revealing the gaps that shape cybersecurity strategy

Cybersecurity Insiders produces independent research based on surveys of cybersecurity leaders and practitioners worldwide. Our reports reveal where security strategies break down in practice — helping organizations benchmark their maturity, identify capability gaps, and prioritize the actions needed to close them.

For more information, visit

cybersecurity-insiders.com

June 2026 // Version 1.6