EMA
IT & DATA MANAGEMENT RESEARCH, INDUSTRY ANALYSIS & CONSULTING

Prepared for:

F5 IBM SECLORE Skyhigh Security virtru

Leveraging DSPM and Artificial Intelligence to Solve Data Security Challenges

June 2026 EMA Research Report

By **Christopher M. Steffen, CISSP, CISA**, VP of Research
Information Security, Risk and Compliance Management

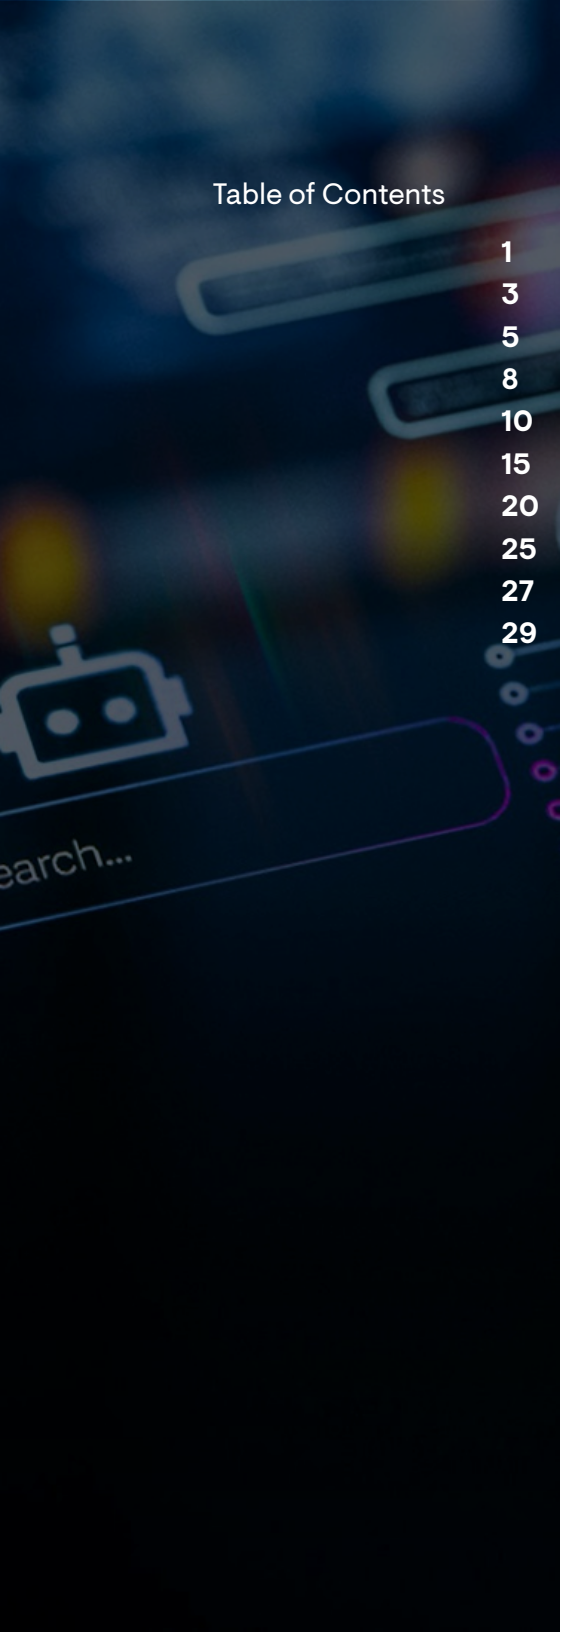

Table of Contents

Introduction

The enterprise data security landscape is undergoing significant transformation. Previously, there was a clear division of responsibility: infrastructure teams managed where data lived, compliance teams governed what it contained, and security teams built controls around its perimeter. Today, two forces arriving simultaneously rendered that model inadequate: the explosion of artificial intelligence across every business function and the collapse of the data perimeter that once made systematic governance tractable.

As organizations matured in their adoption of cloud infrastructure, the conversation shifted from the security of the network edge to the security of the data itself. Data Security Posture Management (DSPM) emerged as the answer to a specific problem: in environments too large and too distributed for manual audit processes, enterprises needed platforms capable of discovering sensitive data continuously, classifying it accurately, and surfacing the misconfiguration and access risks that human review could no longer find at scale. DSPM provided that visibility, and for a time, it was sufficient.

The rapid and broad deployment of AI across the enterprise fundamentally changed what sufficient means. Artificial intelligence does not simply consume existing sensitive data in new ways; it generates entirely new categories of sensitive data at machine speed, moves that data through pipelines that cross jurisdictional and organizational boundaries, and introduces non-human actors whose access behaviors fall entirely outside the governance models DSPM was designed to enforce. The discipline built to govern a relatively stable data estate is now being asked to govern one that is dynamic, distributed, and expanding faster than rule-based processes can track.

The emergence of autonomous AI agents is accelerating this challenge further. As organizations move from human-directed AI tools toward systems capable of independent action – retrieving documents, executing workflows, and interacting with sensitive data without direct human oversight – the governance gap between what DSPM was built for and what enterprise environments now require grows wider with every new deployment.

For this research project, Enterprise Management Associates (EMA) surveyed 225 IT professionals, security practitioners, data governance leaders, and technology business leaders across North American enterprises to understand their current approaches to data security posture management, the pressures reshaping the category, and the capabilities their organizations require as artificial intelligence fundamentally alters the data security landscape.

What is Data Security Posture Management (DSPM)?

Data security posture management (DSPM) is a discipline and category of security tooling focused on discovering, classifying, and governing sensitive data wherever it resides across an organization's environment. Where traditional data security approaches focused primarily on protecting data at the perimeter – controlling what leaves the network – DSPM takes an inside-out approach, beginning with foundational questions: where does sensitive data exist, who has access to it, and is that access consistent with policy?

In practice, DSPM platforms discover sensitive data across cloud environments, on-premises stores, and hybrid infrastructure, classify it according to its sensitivity, regulatory status, and business context, assess the security posture of how that data is stored and accessed, and surface the misconfigurations, excessive permissions, and policy gaps that represent the greatest risk. The core value proposition is visibility and governance at scale – the ability to maintain an accurate, continuously updated picture of a data estate that is too large and too dynamic for manual processes to govern effectively.

The emergence and rapid enterprise adoption of artificial intelligence is introducing a set of forces that challenge each of those foundational assumptions simultaneously. AI changes what data is being created, generating entirely new categories such as prompts, model completions, vector embeddings, and agent interaction logs that existing classification frameworks were not designed to address. It changes who and what is accessing sensitive data, introducing non-human identities, including AI agents, automated pipelines, and service accounts that operate outside the behavioral visibility of traditional access governance. It changes where data travels through training pipelines, inference environments, and multi-cloud compute clusters that move across jurisdictional and organizational boundaries in ways legacy infrastructure was not built to track. It changes the regulatory landscape with AI-specific compliance frameworks arriving on top of an already complex and demanding portfolio of data governance requirements.

Together, these forces do not simply extend the scope of an existing category. They are redefining what data security posture management must do and what it must govern to remain a meaningful component of the modern enterprise security program.

Key Findings

64.4%

AI Overtakes Exfiltration as Top Purchase Driver: Securing AI data flows ranked first among DSPM purchase drivers at 64.4%, surpassing exfiltration prevention (56%) – the first time an AI-specific outcome has led the category

57.8%

AI Data Visibility Is the Top Capability Priority: AI-specific data visibility – covering training datasets, model inputs, and outputs – is rated the most important DSPM capability by 57.8%, ahead of all traditional DSPM functions

40.9%

AI Training Workflows are the Leading Source of Shadow Data: Data science teams moving data for AI training are the leading cause of sensitive data in unmanaged locations at 40.9%, ahead of developer testing environments and backup misconfigurations

37.3%

AI Governance Frameworks Lag Deployment: Only 37.3% of organizations have a fully implemented AI governance framework; 56.4% are only partially implemented, meaning most are deploying AI workloads against incomplete governance foundations

30.2%

No Clear Accountability for AI Data Risk: AI governance responsibility is fragmented across IT (30.2%), security (29.8%), CDO (20%), and committees (18.7%), with no single function holding majority accountability for AI data incidents

76.4%

Prompt Injection – High Concern, Low Deployed Mitigation: Prompt injection concern reaches 76.4%, but specifically deployed mitigation solutions remain substantially lower – a significant market gap reflecting both the novelty of the threat and limited available tooling

57.8%

AI Regulations Now the Top Compliance Driver: Emerging AI-specific regulations are cited by 57.8% as a governance driver – the highest of any framework, surpassing GDPR (43.6%), PCI-DSS (44%), HIPAA (43.1%), and CCPA (44.4%)

58.2%

Encryption Key Ownership is the Top Sovereignty Constraint: Encryption key ownership and localized control are the leading sovereignty constraints at 58.2% above data residency (55.1%), indicating organizations treat cryptographic control and physical location as distinct requirements

25.8%

Sovereignty Ranks Last in DSPM Capability Priorities: Regional sovereignty and residency controls rank as the lowest-prioritized DSPM capability at 25.8% – 32 percentage points below AI-specific data visibility – revealing a critical conceptual gap in procurement

49.8%

In-Region Scanning Is a Structural, Non-Negotiable Requirement: 49.8% require fully in-region scanning; 30.2% require raw data stays in region while metadata may leave. Only 2.7% accept sovereignty managed through manual tagging alone

45.3%

Multi-Cloud Creates Compounding Sovereignty Complexity: 45.3% cite consistent security policies across all cloud environments as their leading multi-cloud concern because each provider implements residency controls and key management differently, creating cross-provider posture gaps

40.9%

AI Training Workflows Create Unintended Sovereignty Violations: AI training data movement is the leading source of shadow data (40.9%), creating sovereignty violations through legitimate operational workflows that were never evaluated against regional data residency constraints

46.7%

Real-Time Residency Enforcement is Required — Auditing Alone is Not Enough: Real-time enforcement of data residency rules is required by 46.7% of respondents, confirming that after-the-fact auditing is insufficient and governance must operate at the point of data movement — which, for HTTP-based object storage, means inline delivery infrastructure sitting in front of the store, capable of jurisdiction-aware routing and access decisions as data moves between storage and AI compute, since a posture-scanning architecture can only observe a violation while an inline control point can prevent it.

85.7%

85.7% Expect Meaningful Workload Reduction From Automation: 85.7% of organizations expect meaningful workload reduction from automated remediation – 33% projecting reductions exceeding 50% and 52.4% expecting moderate reductions of 20–50%

37.3%

DSPM Preferred as Intelligence Provider, Not Enforcement Engine: 37.3% want DSPM to deliver metadata to downstream platforms for policy decisioning; 36% want real-time firewall context. Only 8% want DSPM acting as the central enforcement point

32%

Implementation Risk – Security Operations Absorbs Responsibility: 32% of organizations report DSPM implementation shifts data security responsibility entirely to security operations, creating friction when SecOps applies remediation to data it doesn't own or fully understand

47.1%

Market Evenly Split on DSPM and DLP Future: 47.1% view DSPM as the natural evolution of traditional DLP; 44.4% see them as complementary. Both positions describe the same convergence: classification intelligence enabling more accurate in-motion enforcement

Voices of the Survey

Respondents were asked to explain in as much detail as possible your organization's approach to data security. Below are some curated responses to this question.

Defense-in-depth with a cloud-native focus. We're evaluating AI-driven anomaly detection and have implemented controls to prevent data leakage into unauthorized generative AI tools.

We follow a layered data security approach... we are also using an AI-driven DSPM platform to improve visibility and automated risk recommendation across cloud and on-premises environments.

We use DSPM to find and classify data, enforce zero trust access and protect data with encryption and DLP.

Regulations have already made most of the data security architecture decisions for us. Most of these considerations used to sit with legal, but is now a key factor in pretty much every security conversation we have.

We use a layered, AI-driven approach combining DSPM, SIEM, IAM, and cloud-native security tools to protect data.

Our organization enforces strict access controls, continuous monitoring, and regular audits to protect data. We use endpoint security, firewalls, and are evaluating DSPM and AI-driven threat detection tools.

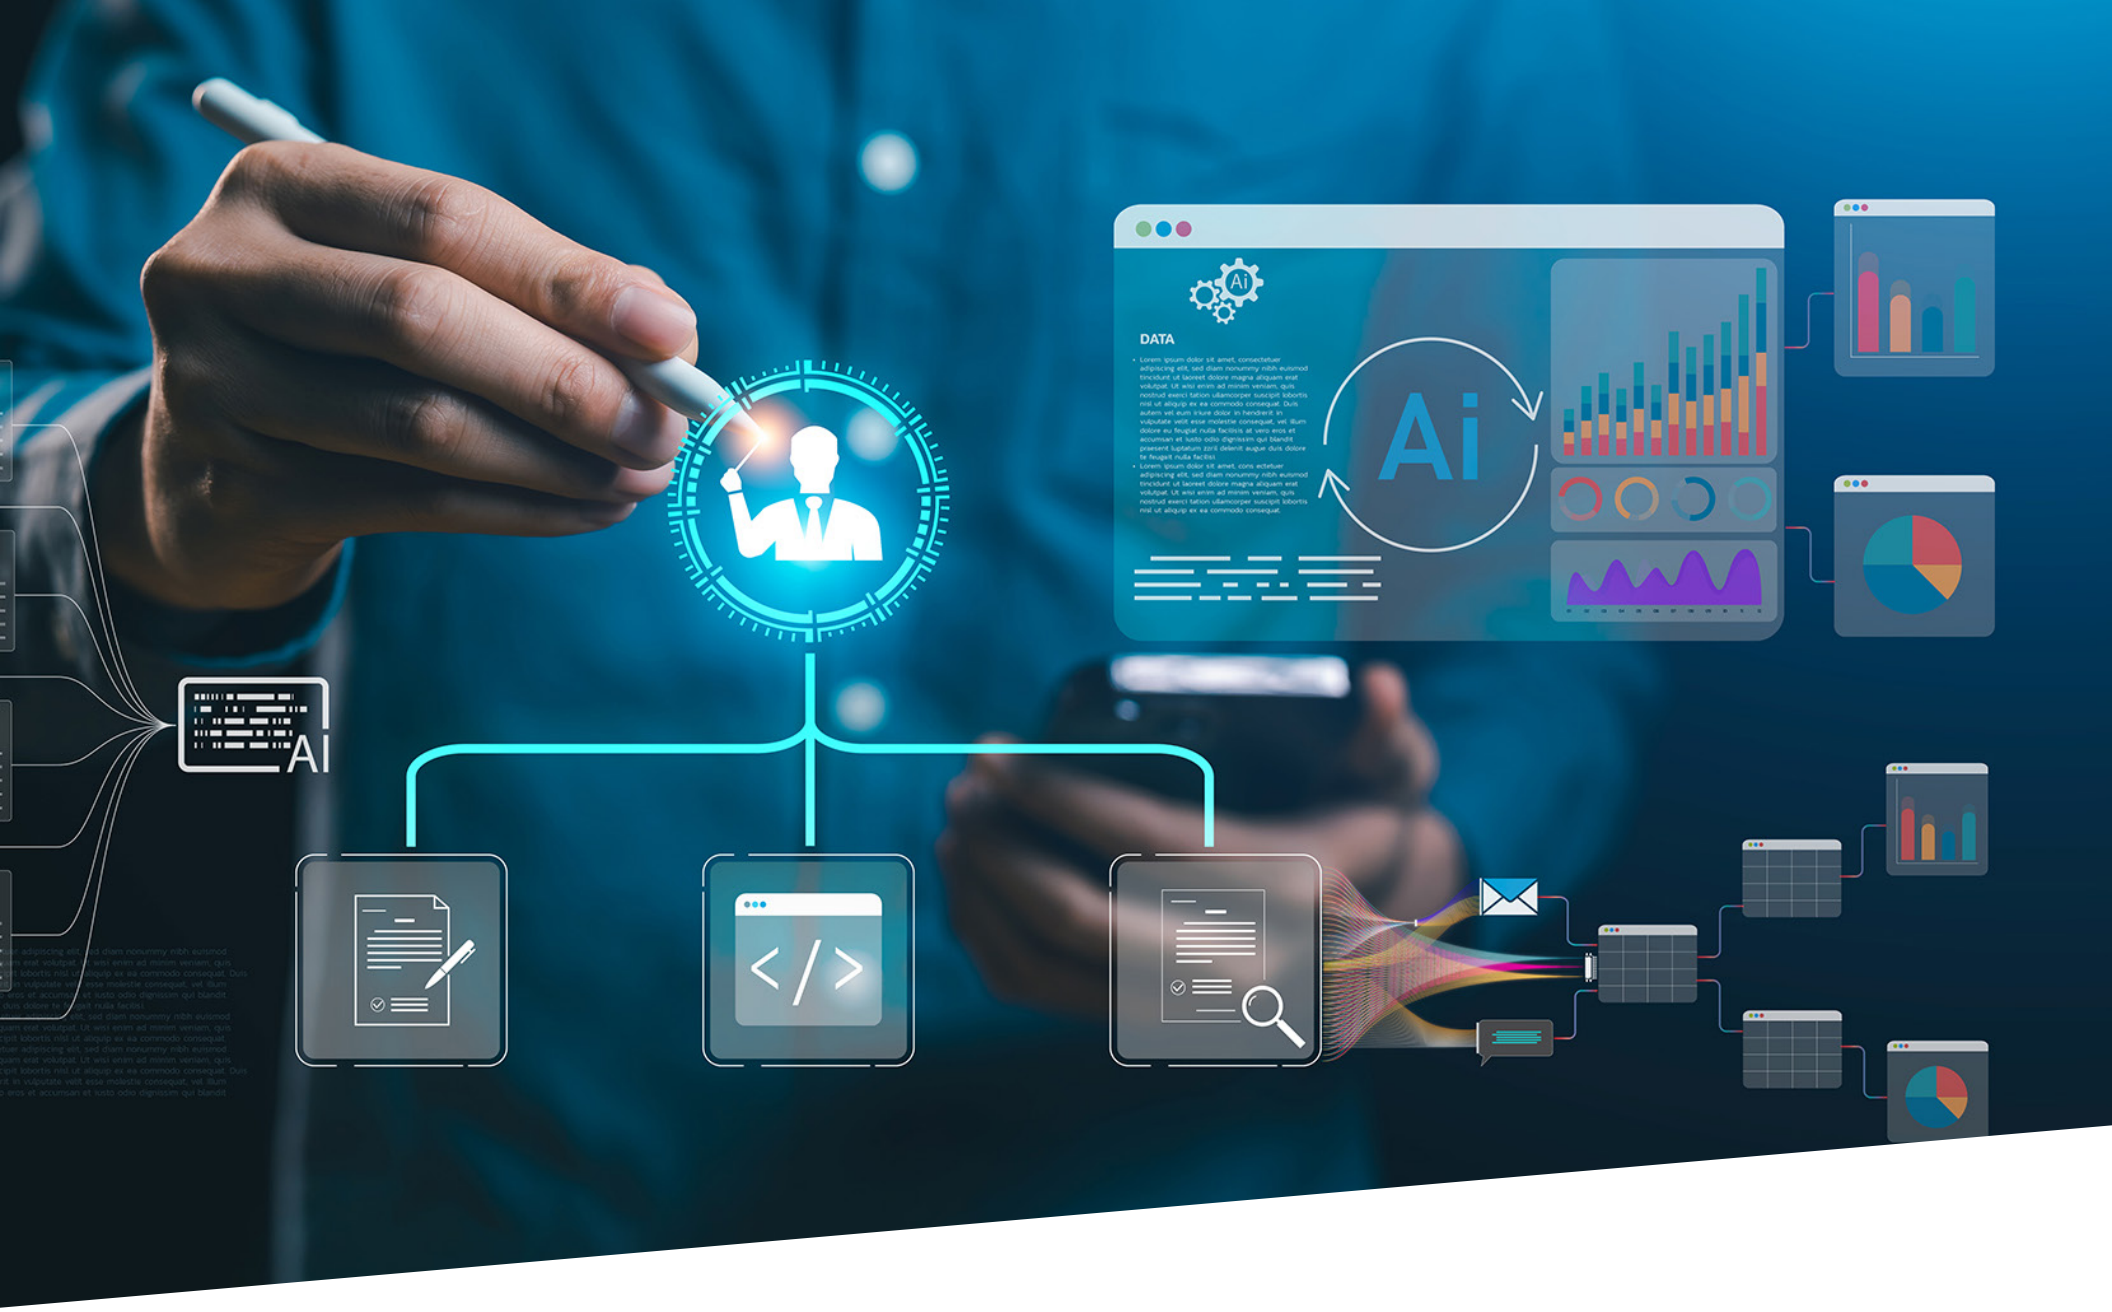

The AI-Driven Transformation of DSPM

Data security posture management was built for a specific data environment. When the category emerged, the problem it addressed was bounded and familiar: discover sensitive data in structured stores, classify it, govern access, and produce the compliance documentation security and audit functions required. The underlying assumption was that data was largely static between transactions – residing in known locations, accessed by identifiable users, and governed through established policy frameworks. That assumption no longer holds. The deployment of artificial intelligence across the enterprise changed the fundamental nature of enterprise data – who generates it, how quickly it accumulates, where it travels, and what governing it requires. This transformation is not a future trajectory to plan. The survey data from this research establishes it as the present operating reality, visible across multiple independent findings that together define a category in active transition.

The most direct evidence of this shift is in what organizations say they are buying DSPM to accomplish. When asked which outcomes matter most when evaluating a data protection solution, securing AI data flows – including prompts and training datasets – was selected by 64.4% of respondents, placing it first among all measured outcomes. This surpasses preventing sensitive data exfiltration at 56%, which was historically the central DSPM value proposition. For the first time in research covering this category, an AI-specific outcome displaced exfiltration prevention as the market's primary purchasing motivation.

The same appears in capability priorities. AI-specific data visibility – covering training datasets, model inputs, and outputs – is rated the most important DSPM capability by 57.8% of respondents, ranking above automated classification (52.9%), data lineage tracking (48.9%), and sensitive data discovery at rest (47.1%). These represent the foundational architecture of traditional DSPM. The market has not abandoned them, but it has placed AI data governance explicitly ahead of all of them.

Which outcomes matter most when evaluating a data protection solution?

Which capabilities would you prioritize in a DSPM solution?

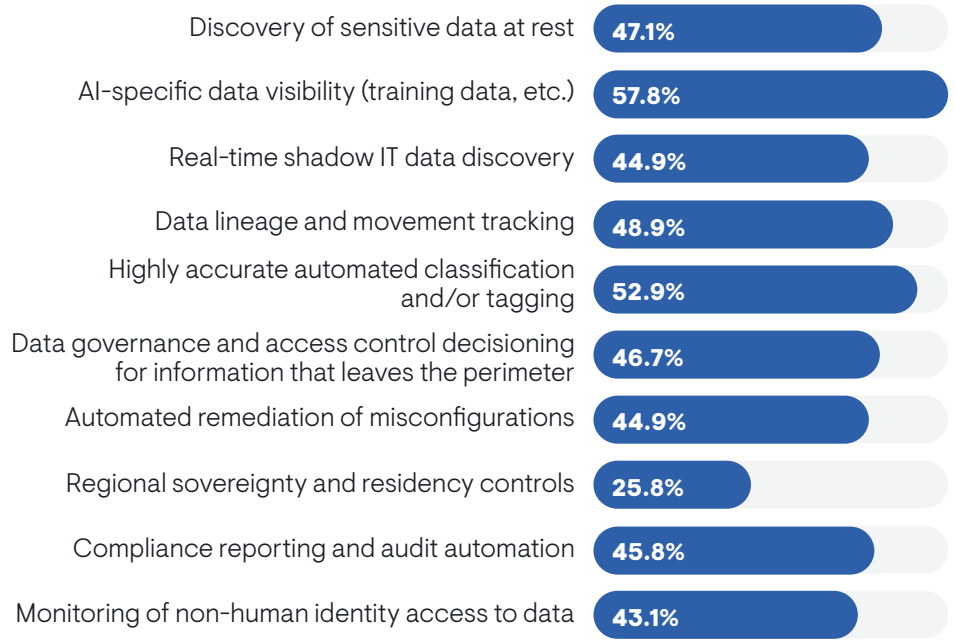

When organizations were asked how important it is for DSPM platforms to provide insight into how others or other processes are using AI Data, including prompt histories and model retention behaviors, it showed the depth of this consensus. Ninety-six percent rated this as either critical or important, with 49.3% selecting critical – the strongest available preference. No other single DSPM capability in the research achieved this level of agreement across the respondent population.

AI and ML workloads are already the leading cause of storage performance degradation, cited by 37.3% of respondents – above all other factors, including client count growth and operational complexity. More specifically, data science teams moving data for AI training have become the leading cause of sensitive data appearing in unmanaged locations at 40.9%, overtaking developer testing environments (25.8%) and misconfigured backup processes (16%). AI workloads are now the primary driver of shadow data creation in most enterprise environments – not a contributing factor, but the leading one.

Despite the clarity of this market signal, organizational readiness to govern AI-generated data lags significantly behind the pace of deployment. Only 37.3% of organizations report having a fully implemented AI governance framework. The majority (56.4%) describe their framework as only partially implemented, and 5.8% are still in development. Most organizations are deploying AI workloads against an incomplete governance foundation, generating data that is classified, traced, and enforced against only partially, or not at all.

How important is it for DSPM to provide insight into AI data usage (prompts/retention)?

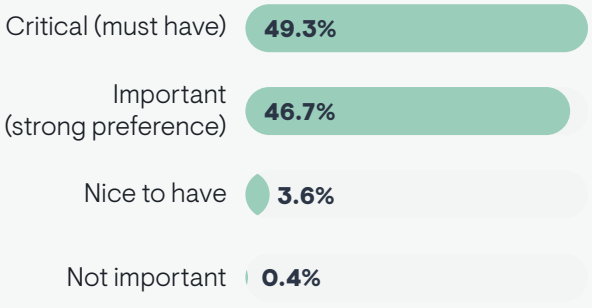

What is the most common reason for sensitive data appearing in unmanaged locations?

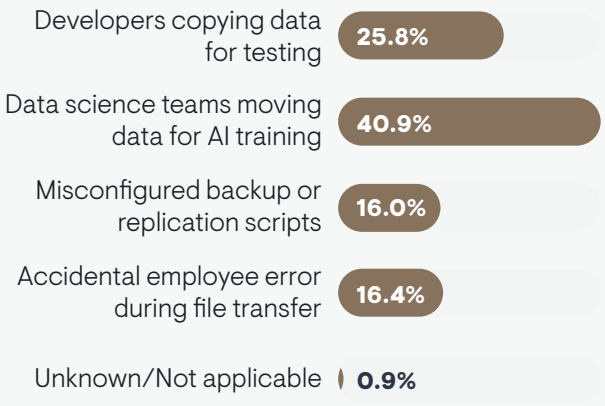

Does your organization have a formal AI governance framework in place?

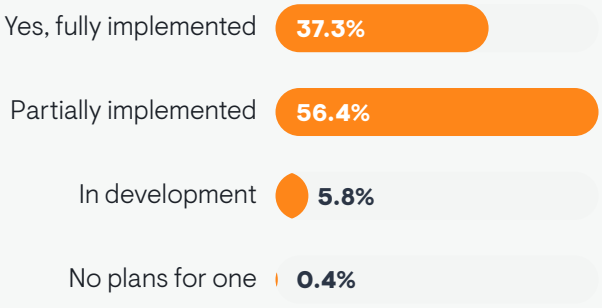

The ownership of AI governance compounds this gap. Responsibility is distributed across IT infrastructure teams (30.2%), security and CISO offices (29.8%), the CDO function (20%), and shared responsibility committees (18.7%), with no single function holding majority accountability. When an AI system generates a sensitive output, moves training data across a store boundary, or exposes prompt histories through an unsecured API, there is no clearly accountable team to detect it, classify it, or remediate it. The absence of a defined owner is not an organizational oversight; it reflects the genuinely cross-functional nature of AI data risk, which sits at the intersection of infrastructure, security, compliance, and business operations.

The implication for technology selection is direct: this governance gap cannot be addressed by extending existing cloud security tools or applying generic data classification frameworks to new data categories. AI-generated data – prompts, completions, vector embeddings, model retention artifacts, and agentic workflow outputs – requires classification intelligence built specifically to understand the contextual relationships and sensitivity characteristics unique to these types. Tools architected for structured financial records or personal data fields do not have native visibility into these categories, and the risk they create remains invisible to platforms not designed to surface it. Organizations evaluating their AI governance readiness should assess their DSPM investments specifically against this capability requirement, treating AI data type coverage as a first-tier selection criterion.

The technical requirements data shows why traditional classification cannot close this gap. 43.6% prioritize custom machine-learning-based classifiers over prebuilt regulatory templates (18.2%) — but this isn't a rejection of deterministic matching. The highest tier of data discovery marries ML context-awareness with Exact Data Matching (EDM), since each covers a failure mode the other cannot.

AI-generated data — prompts, completions, embeddings — requires ML to understand contextual relationships between an output and its source sensitivity. But inference alone introduces probabilistic uncertainty unacceptable when a specific value must be confirmed with certainty, such as an account number in an agent's reasoning trace. EDM closes that gap with mathematical precision rather than inference.

This dual-engine approach secures complex AI completions with zero-trust accuracy: ML catches novel patterns, EDM confirms specific protected values. Rule-based systems weren't built for this complexity, and ML alone leaves organizations exposed to false negatives at scale.

What is your organization's preferred approach to implementing and consuming an SD-WAN or SASE solution?

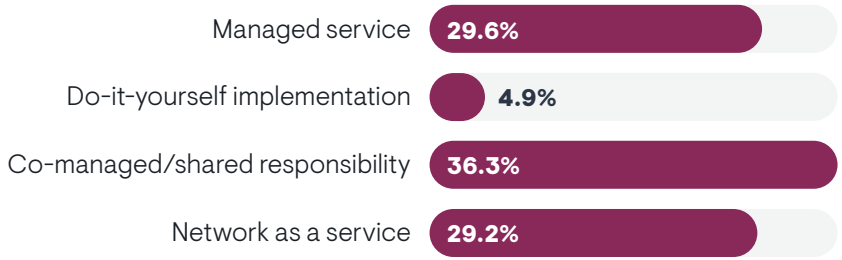

Which of the following best describes your organization's preferred approach to achieving a SASE architecture that integrates SD-WAN and cloud-based security (e.g., security service edge)?

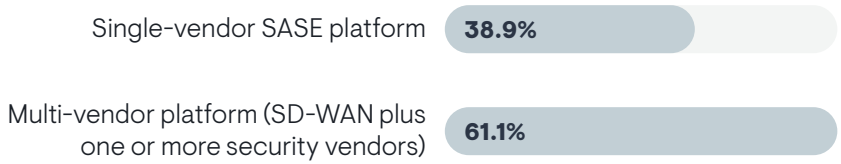

One dimension of the AI data security challenge where the gap between concern and deployed capability is particularly pronounced is prompt injection – attacks that manipulate AI system behavior by embedding malicious instructions within model inputs. Concern about prompt injection attacks is cited by 76.5% of respondents, placing it among the highest AI risk concerns in the research. Yet, the rate of specifically deployed solutions designed to detect and mitigate injection attacks targeting production AI systems remains substantially lower than that concern level. This asymmetry – high recognized risk, limited deployed mitigation – is a consistent pattern across AI security categories and reflects both the novelty of the threat vector and the absence, until recently, of purpose-built solutions positioned to address it. For security leaders evaluating DSPM platforms in 2026, a prompt injections feature as part of an AI security solution should be treated as an immediate selection criterion, not a capability to revisit after initial deployment.

The regulatory environment reinforces the urgency of closing this capability gap before AI deployments scale further. Emerging AI-specific regulations are cited by 57.8% of respondents as a driver for stronger data governance – the highest figure of any framework in the research, above GDPR (43.6%), HIPAA (43.1%), PCI-DSS (44%), and CCPA (44.4%). Many of these regulations require documentation of training data provenance and AI system input and output logging that cannot be reconstructed retroactively. Organizations evaluating DSPM investments in 2026 should treat AI data visibility, training dataset governance, and prompt lineage tracking as core selection criteria – not advanced features to assess in a second evaluation phase. For those operating under GDPR, the EU AI Act, or national AI regulatory frameworks, these requirements carry enforcement weight. AI-specific security capabilities – including specialized posture discovery, prompt lineage documentation, and injection attack mitigation – have moved from discretionary enhancements to compliance-required elements of a mature data security program.

How concerned are you about prompt injection attacks targeting your AI systems?

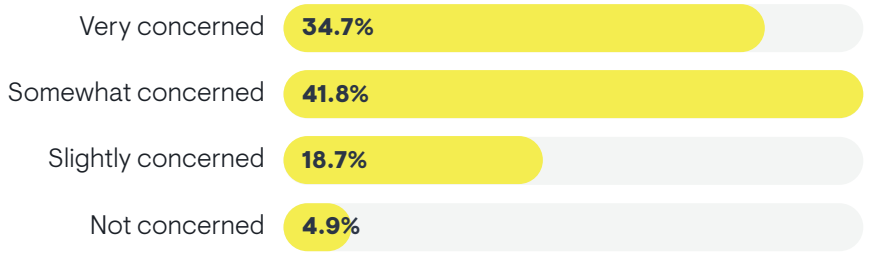

Which regulations are driving your need for stronger data governance?

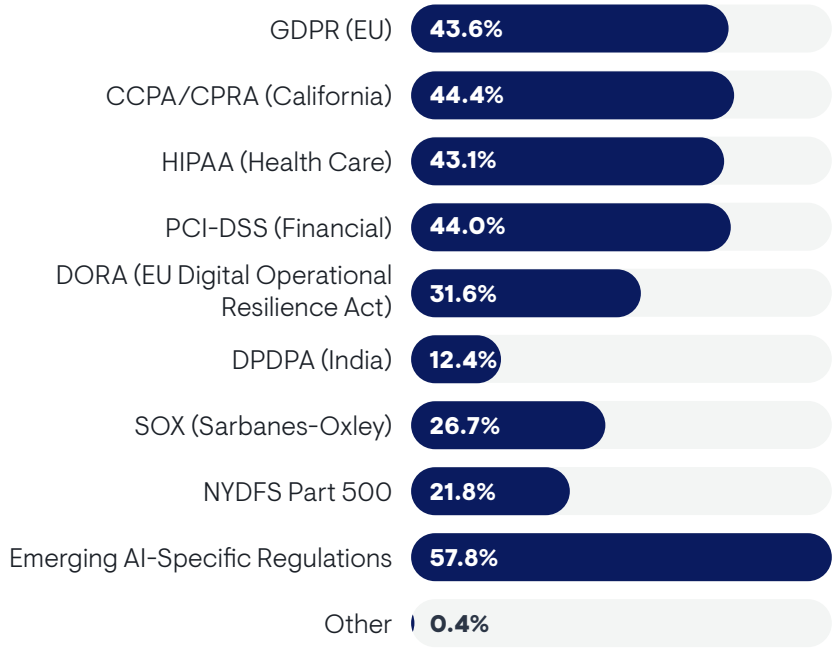

Data Sovereignty – The Emerging Need Gap

One of the most structurally revealing findings in this research is a gap that exists not between what organizations want and what vendors offer, but between how seriously organizations treat sovereignty at the operational level and how rarely they have connected DSPM to sovereignty enforcement. Data sovereignty – the set of requirements governing where data physically resides, who controls the encryption infrastructure protecting it, how its movement across jurisdictions is governed, and how compliance with residency obligations can be demonstrated – registers as a top-tier operational concern in storage and infrastructure data while appearing at the bottom of the DSPM capability priority list. This is the sovereignty gap the market has not yet solved.

When asked which constraints most influence their storage architecture decisions, 58.2% of respondents identify encryption key ownership and localized control as their top concern – the single highest-rated constraint in the research. Data residency, the requirement that data be physically stored within a defined jurisdiction, follows at 55.1%. Auditability and regional compliance reporting rank third at 48.4%, sovereignty-related metadata tagging at 43.6%, cross-border data transfer limitations at 41.3%, and jurisdiction-based access restrictions at 31.1%. These organizations are making active infrastructure decisions – cloud region selection, key management architecture, replication policy configuration – based directly on sovereignty requirements.

The encryption key ownership finding deserves particular attention. Storing data in a compliant geographic location while having the keys that protect it managed by an entity in a different jurisdiction may not satisfy the sovereignty requirements of GDPR, the EU AI Act, or national data localization frameworks. Physical location and cryptographic control are distinct sovereignty dimensions and organizations recognize this distinction. A sovereignty program that addresses only data residency without addressing key ownership is incomplete by the standards that 58.2% of the market is already applying.

Which digital sovereignty constraints most influence your storage architecture?

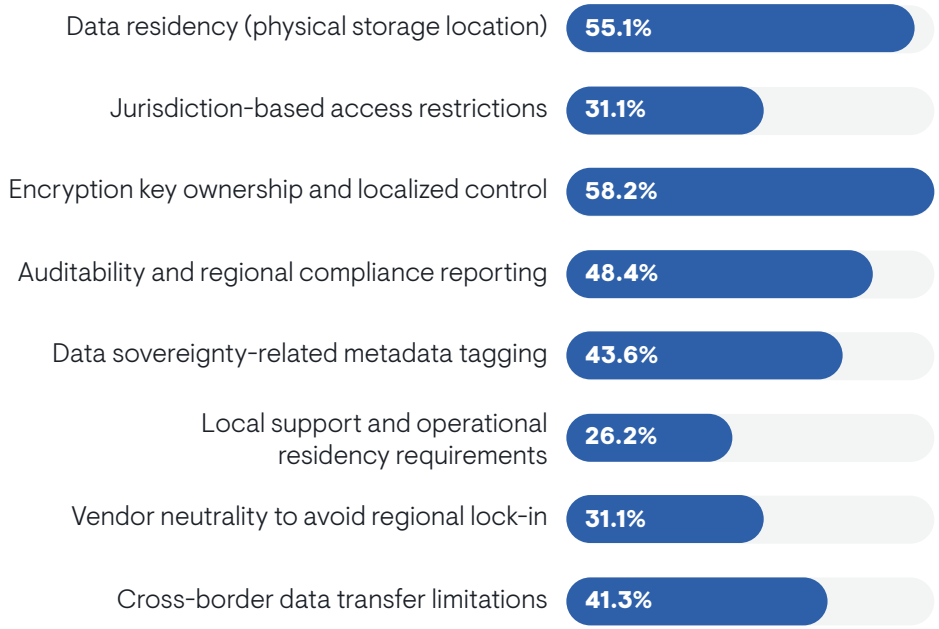

Against this operational intensity, the DSPM capability priority data reveals a stark disconnect. Regional sovereignty and residency controls rank as the lowest-prioritized DSPM capability in the research at 25.8% – thirty-two percentage points below AI-specific data visibility (57.8%) and below every other capability in the survey. This does not indicate that organizations care less about sovereignty than the infrastructure data suggests. It indicates that most organizations have not yet formed a clear conceptual connection between DSPM tools and sovereignty enforcement. Sovereignty is currently managed through infrastructure controls, legal frameworks, and manual operational processes, not through integrated DSPM capabilities.

This conceptual separation is the primary reason sovereignty is underrepresented in DSPM purchase criteria, and it is a gap that will not scale as data environments grow more complex and AI workloads proliferate across regional boundaries. As organizations close this conceptual gap and begin evaluating DSPM through a sovereignty lens, the scanning architecture of the platform will emerge as a primary selection criterion: not a technical specification to review in a final procurement phase, but a compliance requirement to validate before any other capability assessment begins.

Which capabilities would you prioritize in a DSPM solution?

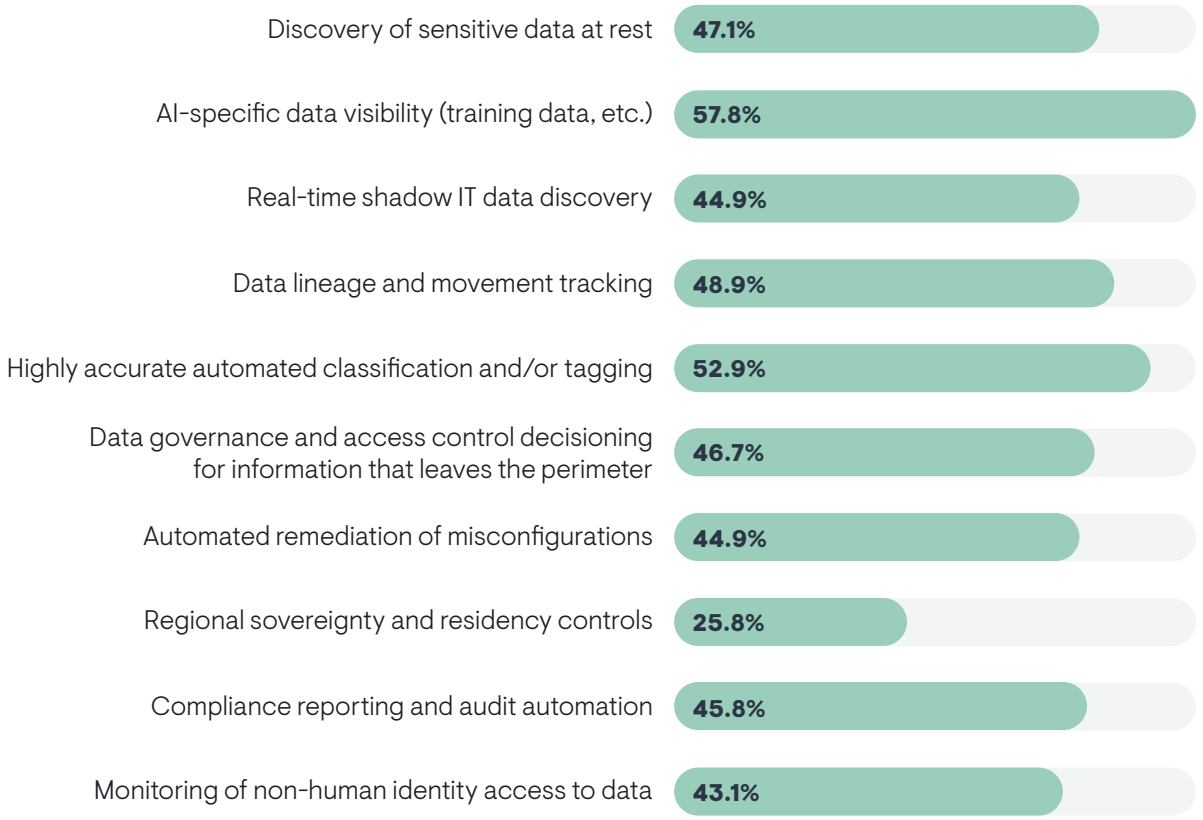

The vendor requirements data provides the clearest evidence of what sovereignty-capable DSPM must deliver operationally. When asked how vendors should support data sovereignty, 49.8% of respondents require that all scanning be performed locally within the relevant region – not as a preference, but as a structural requirement. An additional 30.2% require that metadata may leave the region, but raw data must remain in place. Only 2.7% accept sovereignty managed through manual tagging alone. DSPM platforms that operate through centralized hybrid-based scanning engines, moving data to a central analysis environment for classification, are architecturally incompatible with the requirements of nearly half the enterprise market. For these organizations, the scanning architecture is a regulatory boundary.

This requirement places specific and non-negotiable architectural constraints on DSPM vendors. Platforms that perform classification through centralized analysis pipelines – regardless of how their sovereignty capabilities are described in marketing materials – are structurally incompatible with the requirements of the nearly 80% of organizations that either prohibit raw data from leaving the region entirely or require fully in-region scanning operations. Satisfying these requirements demands that vendors maintain regional infrastructure capable of executing classification locally, hold the compliance certifications required by the data protection frameworks in force in each region they serve, and ensure that neither raw data nor sensitive metadata traverses geopolitical boundaries during processing without explicit organizational authorization and documented legal basis. These are architectural requirements that must be verified through technical validation during the procurement process – not accepted on the basis of general compliance claims or certifications that address data storage without addressing scanning and processing operations.

How should a vendor support data sovereignty requirements?

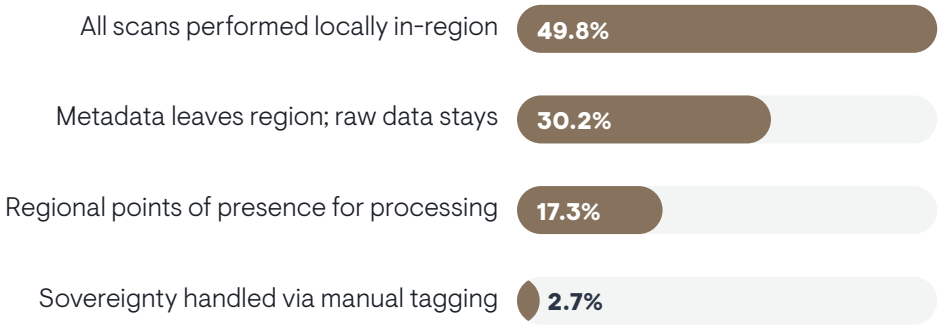

The multi-cloud environment that characterizes most enterprise data estates makes sovereignty governance significantly more demanding. As organizations standardize on S3-compatible object storage for AI training and RAG workloads, many are deploying a unified delivery layer — an "S3 front door" — in front of distributed object stores to provide a consistent global endpoint, uniform security policy, and traffic management across clouds and on-prem environments. When asked how hybrid and multi-cloud adoption changed their storage requirements, 45.3% of respondents cite the need for consistent security policies across all cloud environments as their leading concern, and 28% cite the need for a unified global endpoint — both describing requirements of that front-door layer. Each major cloud provider implements data residency controls, encryption key management, and cross-region access restrictions differently. Organizations operating across two or more providers must maintain equivalent sovereignty postures across architectures that do not natively align, using provider-native controls that produce consistent posture within a single platform but not across the full estate. DSPM that can abstract these differences — applying consistent classification, policy, and compliance reporting across heterogeneous cloud environments — closes a gap that provider-native tools structurally cannot.

The intersection of sovereignty requirements with AI workloads represents an emerging risk that the research captures through its shadow data finding. Data science teams moving data for AI training are the leading cause of sensitive data in unmanaged locations at 40.9%. The sovereignty implication is direct: when a data science team extracts records from a regionally restricted store for AI training and processes that data in a compute environment in a different jurisdiction, the organization may have violated data residency requirements through a legitimate operational workflow that was never evaluated against sovereignty constraints. Real-time enforcement of data residency rules is identified as a requirement by 46.7% of respondents — confirming that auditability alone is not sufficient. Organizations need governance that prevents violations at the point of data movement, and that requires sovereignty-aware classification in place before AI training pipelines scale.

How has hybrid or multi-cloud adoption changed your S3 front door requirements?

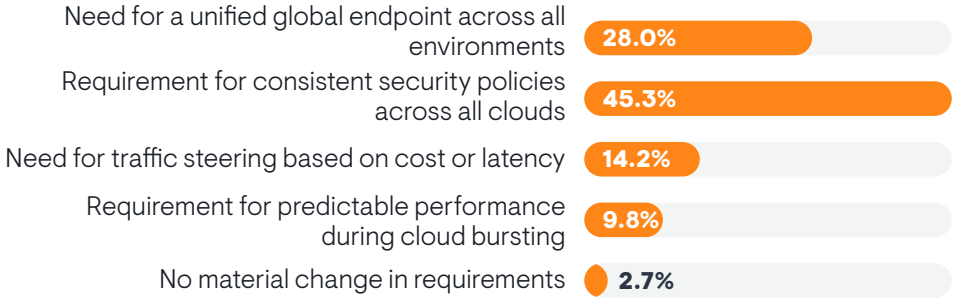

What is the most common reason for sensitive data appearing in unmanaged locations?

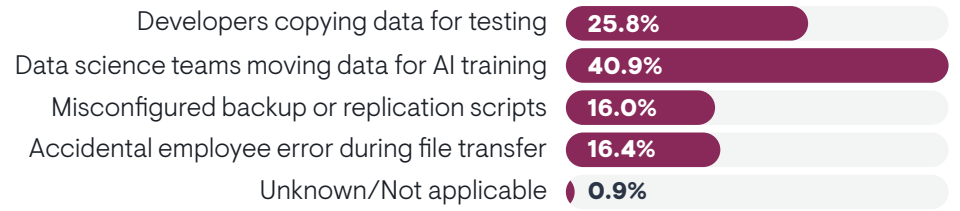

Which use cases require real-time (inline) enforcement rather than just auditing?

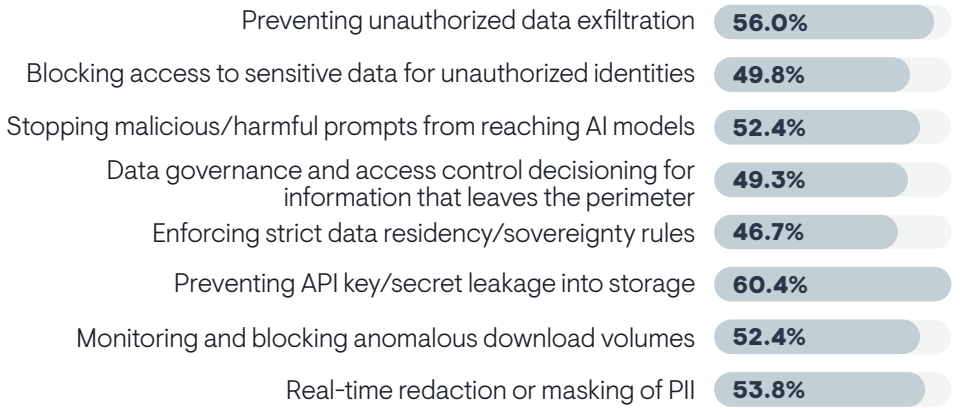

Automated Remediation and the Enforcement Paradox

Enterprise demand for automated data security enforcement is broad, consistent, and growing. When asked whether automated remediation could reduce their manual security workload, 85.7% of respondents expect meaningful reduction. The question is not whether organizations want automation – they clearly do – but how they want it designed, where they want it applied, and what failure modes they are working to avoid. The nuance in the research data on this topic is as important as the headline figure.

The automation demand disaggregates into two meaningfully different positions. Thirty-three percent expect significant workload reduction exceeding fifty percent – a population that has either implemented effective automated remediation or is confident in its potential. The larger cohort, at 52.4%, expects moderate reduction in the twenty to fifty percent range. This is not skepticism: only 1.3% explicitly distrust automated remediation. It is a deliberate hedge that reflects either prior experience with the operational failure modes of automated enforcement, or a considered preference for human oversight at key decision points, or both.

The ROI framing matters for how security leaders communicate automated remediation investments to executive sponsors. An 85.7% expectation of meaningful workload reduction is not simply a projection about security improvement; it is a statement about operational efficiency with direct financial consequences. Security analysts redirected away from routine classification, alert triage, and manual policy enforcement can be reallocated to higher-value activities: threat investigation, governance framework development, exception handling, and the cross-team collaboration that effective data security programs require. For organizations in which security operations headcount is constrained, automated remediation is not a quality-of-life improvement; it is the mechanism by which existing teams extend their effective capacity. Executive sponsors evaluating DSPM investments should be presented with this efficiency dimension alongside the risk reduction case – the two together make a materially stronger justification than either does in isolation.

To what extent can automated remediation reduce your manual security workload?

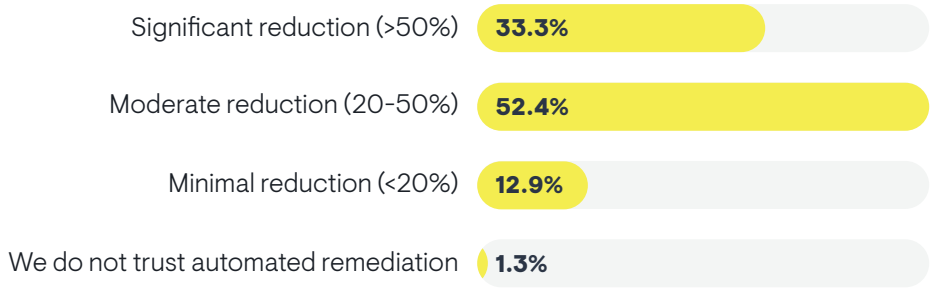

The zero trust integration data provides the clearest articulation of how organizations actually want automated enforcement to function. When asked how DSPM should best integrate into a zero trust architecture, the most common response (37.3%) was that DSPM should deliver metadata to a downstream platform for policy decisioning. The second most common (36%) was that DSPM should provide real-time context for firewall and proxy policies. Only 8% of respondents want DSPM to act as the policy decision point for all data access.

This is a coherent and sophisticated market preference. Organizations want DSPM to classify data, score risk, and surface enforcement signals, and they want that intelligence to inform decisions made by their existing security infrastructure rather than creating a new autonomous enforcement layer. The preference reflects operational experience: concentrating enforcement authority in a single system that lacks full business workflow context generates the specific failure modes that erode trust in automated controls and lead business teams to find workarounds that undermine the security program entirely.

How should DSPM best integrate into a zero trust architecture?

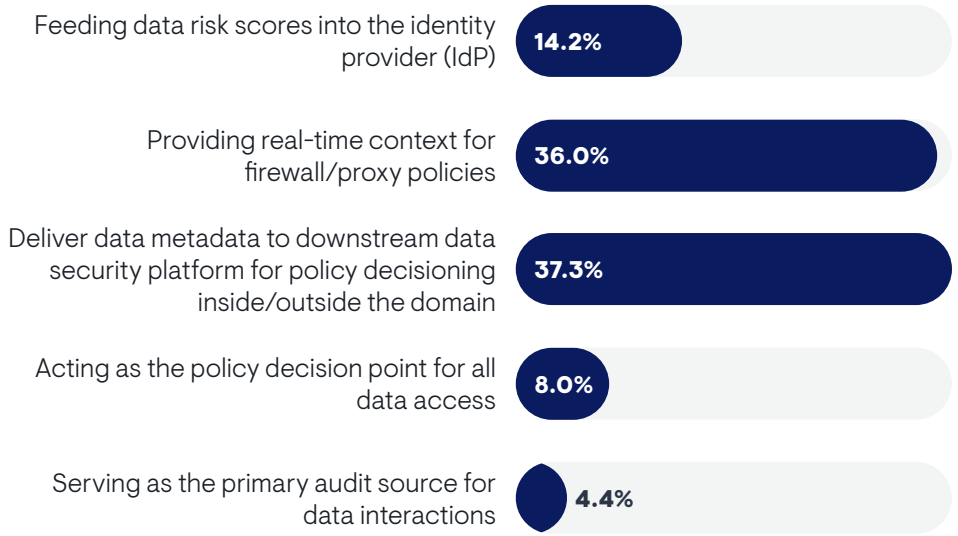

There is a specific organizational dynamic that is the primary source of remediation friction in enterprise environments. When asked how DSPM implementation affects the roles of security operations versus data governance teams, 44.9% report that DSPM forces deeper collaboration on a shared platform – a productive outcome. However, 32% report that DSPM implementation shifts data security responsibility entirely toward security operations teams. This 32% represents the friction population.

When security operations teams apply automated remediation to data they do not own — a contract file that travels by email to a business partner each quarter, a training dataset a data science team has prepared over months, a customer record set an analytics team needs to close a time-sensitive commitment — they are making enforcement decisions without the business workflow context needed to distinguish a legitimate data movement from a policy violation. The result is not only operational disruption. It is the erosion of organizational trust in automated controls, because business teams discover that critical workflows can be interrupted without warning or recourse by security systems that do not understand the work they are attempting to protect.

The research points toward a graduated remediation model as the approach best aligned with both the security requirements and the operational realities of enterprise environments. Rather than binary block-or-allow enforcement, a graduated model calibrates response intensity to the risk level of the detected event. Anomalous events surface as alerts to data owners and security teams with no automated action, but with the context needed for rapid human review. Confirmed higher-risk patterns trigger temporary access restriction with expedited exception workflows for cases in which business continuity requires fast resolution. Persistent data protection controls apply to confirmed sensitive data, protecting it in place without blocking the workflows that depend on it.

How does DSPM implementation affect the roles of SecOps vs. data governance?

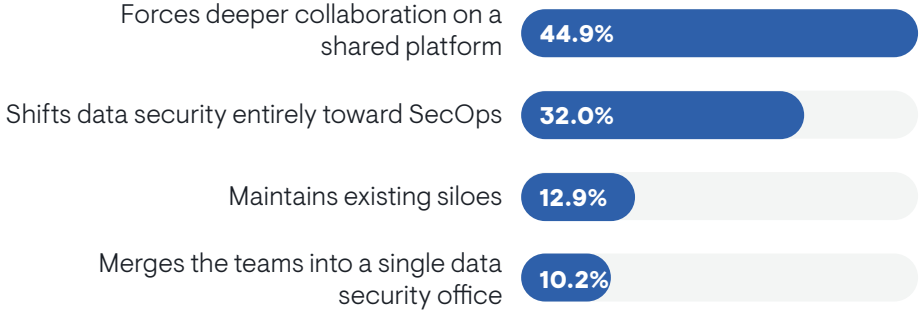

The graduated model is also better suited to the correlated risk intelligence that effective DSPM generates. When a platform detects not an isolated misconfiguration or a single over-permissioned identity, but the convergence of multiple individually manageable conditions — sensitive data, over-permissioned access rights, and a policy gap that together constitute an outsized exposure — the risk level of that convergence exceeds the sum of its parts.

The broader strategic context of the remediation discussion is shaped by a market debate this research captures in near-even division. When asked whether DSPM will eventually replace traditional data loss prevention, 47.1% believe it represents the natural evolution of the DLP category, while an almost equal 44.4% believe the two approaches are complementary rather than substitutive. This near-even split reflects a market at a genuine inflection point. DLP and DSPM have historically addressed adjacent problems with different architectures: DLP focused on data in motion, enforced at network and endpoint layers, while DSPM focused on data at rest, governed through agentless discovery and posture assessment.

The convergence pressure is real – the sensitive data that DSPM classifies needs inline enforcement, and the DLP policies enforcing in-motion controls need the classification context that DSPM provides. Organizations that resolve this through integrated platforms, where classification intelligence directly informs in-motion enforcement, will close the gap between current posture and the enforcement capability the market is converging toward.

For organizations with established DLP programs, this framing matters practically: DSPM does not render existing data protection investments obsolete. It extends their value. The sensitivity intelligence that DSPM builds – context about what data exists, where it resides, which identities access it, and how it flows – is precisely the classification input that DLP enforcement engines require to operate with accuracy and low disruption. Organizations that integrate DSPM classification with existing DLP enforcement infrastructure find their DLP programs becoming more precise and less prone to false-positive interruption, not competing with a parallel enforcement layer.

Do you believe DSPM will eventually replace traditional data loss prevention (DLP)?

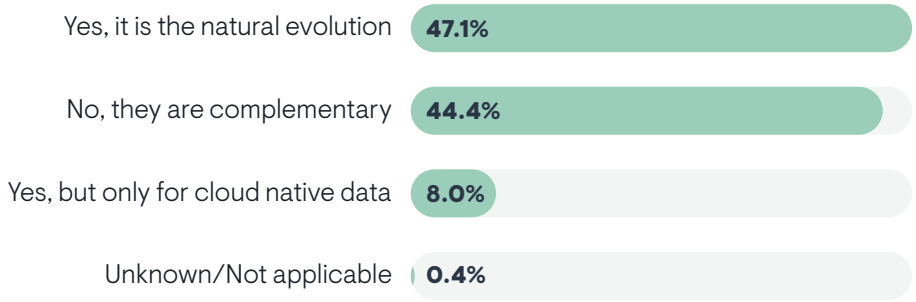

Conclusion

The enterprise data security landscape is under pressure from multiple directions, and the common thread running through each pressure point is the same. Artificial intelligence arrived in the enterprise before the governance infrastructure needed to manage it was fully built, and the consequences of that sequencing are visible across every dimension of this analysis.

The market is already clear about where data security priorities have moved. Protecting AI data flows has overtaken exfiltration prevention as the primary purchase outcome. AI-specific regulations have displaced frameworks that defined the compliance agenda for nearly a decade. Non-human identities – the agents, pipelines, and service accounts that now constitute the majority of data access activity in most environments – have become the highest-rated AI risk concern, yet remain among the least governed. The sovereignty requirements organizations treat as operational imperatives at the infrastructure level have not yet been connected, in most cases, to the DSPM programs responsible for enforcing data policy. The demand for automated remediation, while broad and growing, reflects a nuanced expectation: organizations want protection that follows data intelligently, not enforcement that disrupts the business workflows it is intended to secure.

Individually, each of these findings carries clear implications for how organizations should evaluate, deploy, and mature their data security posture management programs. Collectively, they point to a more foundational conclusion – one that extends beyond the boundaries of this research.

The deployment of agentic AI systems without an established DSPM program in place is not a manageable risk. Autonomous agents retrieve, process, generate, and act on sensitive data continuously and at machine speed. Prompts, completions, retrieval logs, generated artifacts, and agent memory stores begin accumulating from the moment a system goes live. Without classification frameworks, lineage tracking, and policy enforcement already operational, that data estate becomes ungovernable at the pace it is created. The gap does not stabilize; it compounds with every agent interaction, and retroactive governance cannot close what real-time governance failed to capture.

The organizations best positioned to benefit from agentic AI are those that have already established the data security infrastructure to govern what those systems will generate. DSPM is not a parallel workstream to AI deployment. It is a prerequisite for deploying AI responsibly.

EMA Perspective

The findings in this research paint a consistent and urgent picture: enterprise data security has entered a period of structural transformation, and the tools, frameworks, and governance models that defined the previous era are no longer adequate for the environment in which organizations are operating today. Artificial intelligence has not simply created new data types to protect; it has changed the fundamental logic of what data security must accomplish, how it must be architected, and which capabilities must be treated as non-negotiable.

Two patterns in the research deserve particular emphasis:

- The gap between recognized risk and deployed capability is wide and widening. Organizations understand the threats – 96% prioritize AI data insight, 76.4% are concerned about prompt injection, and 81.8% expect specialized AI agent security to be a 2026 priority. Yet, governance frameworks are only fully implemented in 37.3% of organizations and purpose-built mitigation for the highest-concern categories remains substantially under-deployed. Security programs built on awareness without capability are not security programs; they are risk registers.
- The sovereignty finding reveals a structural blind spot. Organizations are making consequential infrastructure decisions – cloud region selection, key management architecture, replication policy – based directly on sovereignty constraints, while simultaneously ranking sovereignty at the bottom of their DSPM capability priorities. This disconnect will not hold as AI workloads push sensitive data across regional boundaries at scale. The scanning architecture of the DSPM platform is a compliance boundary, not a deployment preference, and the market has not yet absorbed that reality.

Several recommendations stand out for those evaluating DSPM solutions in the age of AI:

- Treat AI data type coverage as a first-tier selection criterion. Evaluate DSPM platforms specifically against their ability to classify prompts, completions, embeddings, model retention artifacts, and agentic workflow outputs. Generic classification frameworks are not sufficient for these data types.
- Close the prompt injection gap now, not in a future evaluation phase. About 76% of organizations recognize the risk; the deployment rate of purpose-built mitigation does not reflect that concern. Organizations operating AI systems in production should treat injection defense as an immediate operational requirement.
- Validate scanning architecture before evaluating any other DSPM capability. For organizations subject to GDPR, the EU AI Act, or national data residency frameworks, the question of whether a platform performs classification locally within the data's jurisdictional boundary is a compliance prerequisite, not a technical footnote. Evaluate it first.
- Assign a defined owner for AI data governance, not a committee. Distributed responsibility across IT, security, CDO, and shared committees produces the outcome the research documents: no one is accountable when an AI system generates a sensitive output, moves restricted data, or exposes prompt histories through an unsecured endpoint. Accountability must be singular and explicit.
- Build the ROI case for automated remediation around operational efficiency, not just risk reduction. Almost 86% of security leaders expect meaningful workload reduction from remediation automation. Executive sponsors respond to efficiency arguments. The combined case – risk reduction plus analyst capacity recaptured – is materially stronger than the security case alone.

Having your data security process implemented and refined is really the only way possible to deal with the onslaught of AI-generated data. Those organizations that deploy agentic solutions before understanding the risks associated with the mountains of data created by AI through prompt creation, analysis, and revision will very likely not be able to recover in a meaningful way.

Research Methodologies and Demographics

You indicated that your department is IT-related. Which of the following BEST describes your specific role?

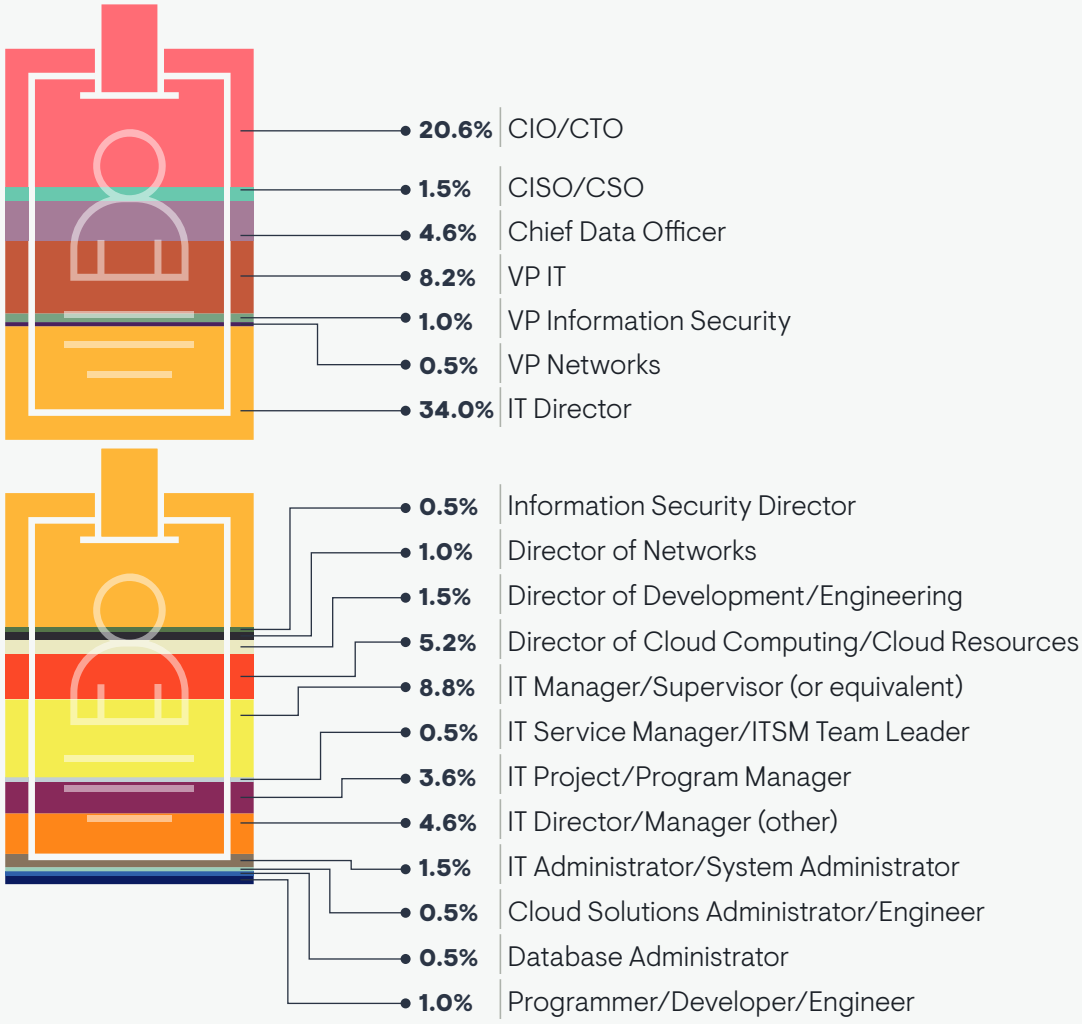

In total, how many employees are currently working in your organization?

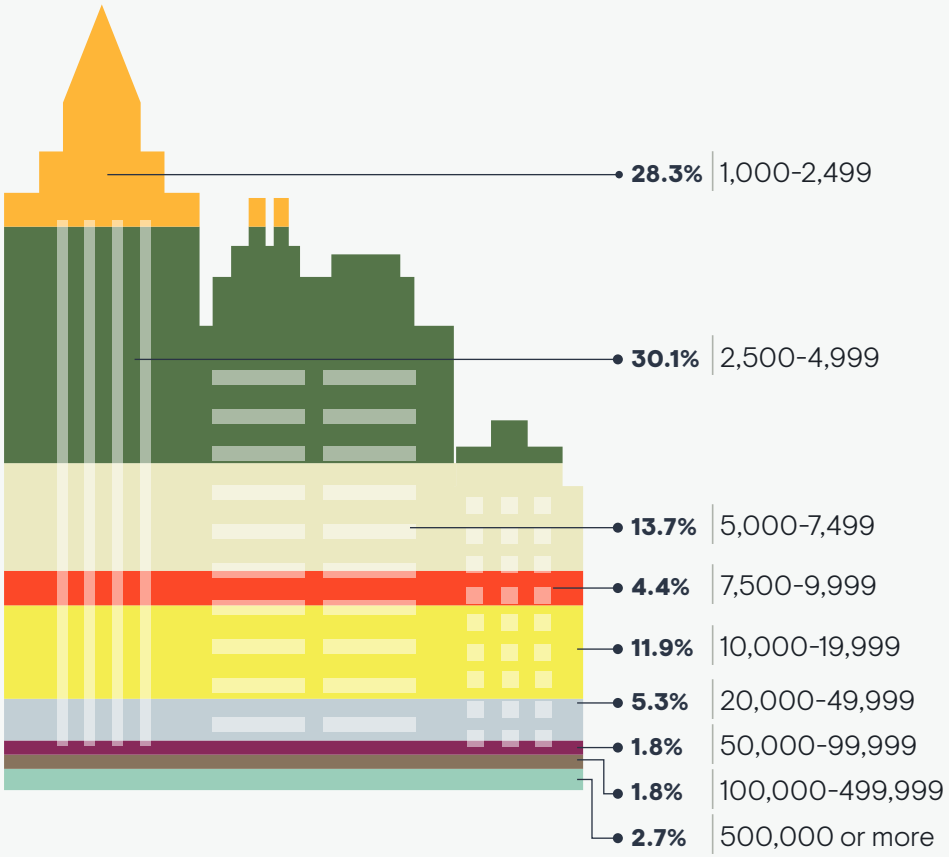

What is your organization's annual sales revenue?

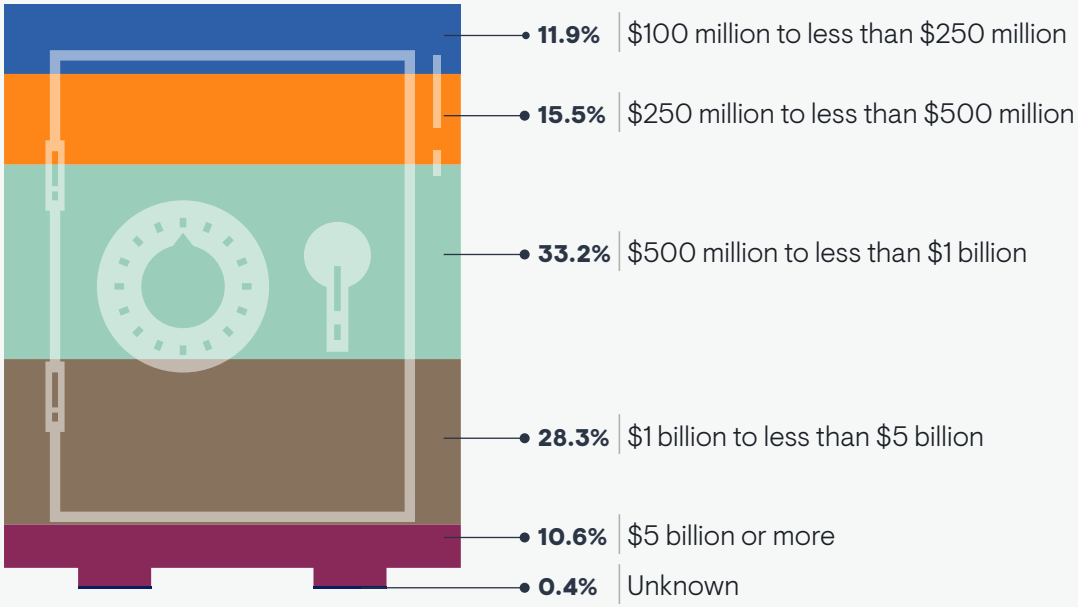

Which of the following best describes your organization's primary industry?

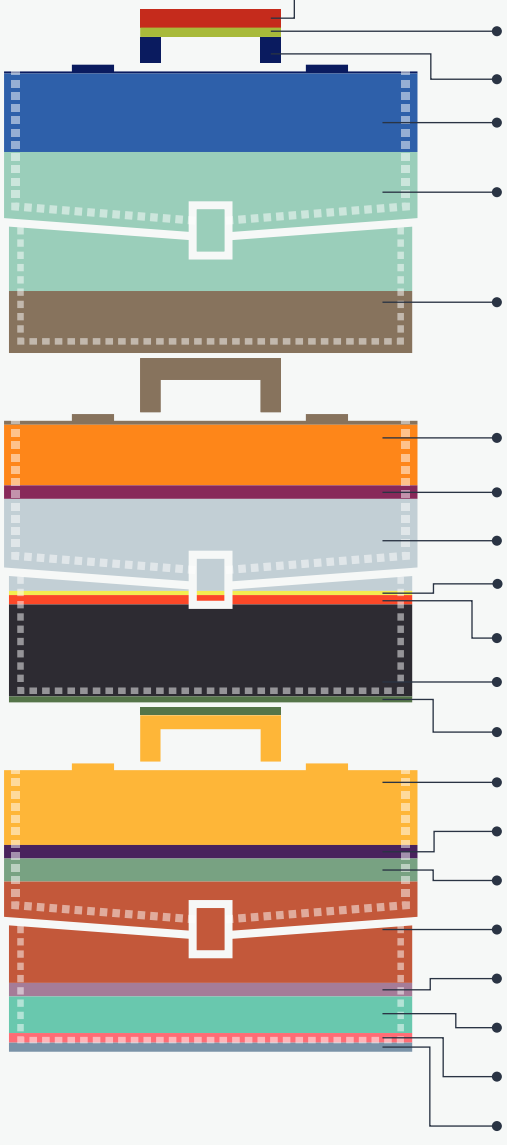

- **1.8%** Aerospace/Defense
- **0.9%** Automotive
- **3.5%** Business Services/Consulting
- **7.5%** Computer/Technology Hardware (devices, chip, computer/networking hardware)
- **13.3%** Computer/Technology Software (mobile app, consumer, custom, web-based)
- **12.8%** Computer/Technology Services (IaaS, SaaS, MSP, MSSP, cloud provider)
- **5.8%** Computer/Technology: Other
- **1.3%** Ecommerce
- **8.8%** Finance/Financial Services/Banking
- **0.4%** Gaming/Digital Entertainment
- **0.9%** Government (federal, state & local)
- **8.8%** Health Care/Medical/Pharmaceutical
- **1.8%** Insurance
- **12.4%** Manufacturing
- **1.3%** Oil/Gas/Chemicals
- **2.2%** Professional Services (non-technical)
- **9.7%** Retail/Wholesale/Distribution
- **1.3%** Telecommunications
- **3.5%** Transportation/Airlines/Trucking/Rail
- **0.9%** Utilities/Energy
- **0.9%** Other

EMA
IT & DATA MANAGEMENT RESEARCH, INDUSTRY ANALYSIS & CONSULTING

Christopher M. Steffen, CISSP, CISA
Vice President of Research
CSteffen@enterprisemanagement.com

Check out the CyberSecurity Awesomeness Podcast:
https://www.cybersecurityawesomeness.com

CYBERSECURITY
AWESOMENESS PODCAST

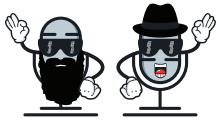

EMA

About Enterprise Management Associates, Inc.

Founded in 1996, Enterprise Management Associates (EMA) is a leading IT research and consulting firm dedicated to delivering actionable insights across the evolving technology landscape. Through independent research, market analysis, and vendor evaluations, we empower organizations to make well-informed technology decisions. Our team of analysts combines practical experience with a deep understanding of industry best practices and emerging vendor solutions to help clients achieve their strategic objectives. Learn more about EMA research, analysis, and consulting services at **www.enterprisemanagement.com** or follow EMA on **X** or **LinkedIn**.

2770 Arapahoe Road, Ste 132 - 544, Lafayette, CO 80026 +1 303.543.9500 **www.enterprisemanagement.com**

4743.062226