A Shifting Paradigm: Governance in the Age of AI

Cisco 2026 Data and Privacy Benchmark Study

Contents

Executive summary

The **Cisco 2026 Data and Privacy Benchmark Study** sheds light on the evolution of privacy programs within organizations around the world.

This report shows investments in privacy are delivering enterprise-wide improvements and returns. Almost all organizations (99%) report at least one tangible benefit from their privacy initiatives, topped by faster innovation, improved operational efficiency, and greater customer loyalty.

At the same time, artificial intelligence (AI) is driving a surge in demand for data to train new technologies, putting pressure on existing privacy and data governance frameworks. And, as the regulatory environment continues to evolve in the wake of this dynamic, organizations must be ready to adapt their privacy programs to be both competitive and compliant.

90% of organizations say their privacy programs have expanded because of AI, and 43% report that privacy spending has increased over the past year. In the next two years, 93% plan to allocate more resources to at least one area of privacy and data governance to manage the growing complexity of AI systems and expectations of customers, clients, and regulators.

This ninth edition of the study reinforces a theme consistent with the **Cisco AI Readiness Index**: AI ambition continues to outpace readiness. The data privacy foundations built in the GDPR era are proving essential, but no longer sufficient, as AI introduces new expectations around data use, accountability, transparency, and contractual clarity. Organizations largely recognize what AI implementation demands; the work now lies in extending established privacy frameworks to meet the scale, speed, and complexity of AI.

Global operations are also under pressure. Eighty-five percent of organizations say data localization adds cost, complexity, and risk to cross-border service delivery. Results show global organizations are more likely to see significant operational challenges arising from localization compared to single-market players (88% vs. 79%). Global companies prefer technology partners that match their footprint: 82% believe global-scale providers are better at managing cross-border data flows. The assumption that locally stored data is inherently more secure is gradually eroding, falling from 90% in 2025 to 86% in 2026.

Operationally, AI has amplified long-standing data quality and management challenges. Sixty-five percent of organizations struggle to access relevant, high-quality data efficiently, while 77% identify intellectual property (IP) protection of AI datasets as a top concern. Many (70%) acknowledge risk exposure from the use of proprietary or customer data in AI training – evidence that data governance discipline must evolve in step with innovation.

Vendor relationships are another focal point for trust in AI. Eighty-one percent of organizations say their GenAI providers are transparent about data practices. However, only 55% require contractual terms that define data ownership, responsibility, and liability, underscoring the need for stronger accountability frameworks as dependency on third-party AI partnerships expands.

Through these findings, one conclusion is unmistakable: data management – both personal and non-personal data – has become a critical part of digital trust and organizational readiness. As AI reshapes every industry, leading organizations must evolve data governance programs to be enablers of responsible innovation. They are moving beyond compliance to further build transparency, fairness, accountability, privacy, security, and reliability into every decision about data.

The **Cisco 2026 Data and Privacy Benchmark Study** captures an enterprise landscape in the midst of change. Privacy, AI, data governance, and cybersecurity are no longer operating as separate disciplines, but are increasingly intertwined around a shared requirement: trust. The organizations best positioned to succeed are not those reacting piecemeal to shifting rules or technologies, but those deliberately embedding trust into how decisions are made, data is governed, and AI is deployed. In doing so, they are turning trust from a defensive necessity into a durable driver of innovation and growth.

Key findings

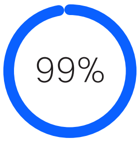

Privacy Delivers Tangible Value: 99% of organizations report measurable benefits from their privacy investments, with ‘enhanced agility and innovation’ now leading as the top outcomes.

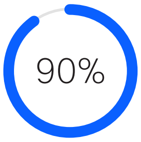

AI Reshapes Privacy’s Scope: 90% of organizations report their privacy programs have expanded due to AI, and 43% increased privacy spending in the past year.

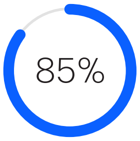

Localization Raises Operational Challenges: 85% of organizations say data localization adds cost, complexity, and risk to cross-border service delivery.

AI Governance Gap: Only 12% of organizations describe their AI governance committees as mature and proactive.

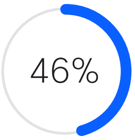

Transparency Drives Trust: 46% of organizations identify ‘clear communication about data use’ as the most effective action to build customer confidence, outranking ‘compliance or breach protection’.

Number of respondents

Brazil	358
China	459
France	411
Germany	437
India	362
Italy	421
Japan	445
Mexico	390
Saudi Arabia	355
Spain	469
United Kingdom	371
United States of America	746

Number of respondents

Educational Services	265
Financial Services	683
Healthcare	321
Hospitality	291
Government	97
Manufacturing	808
Professional Services	457
Public Administration	102
Retail	606
Technology	892
Transportation	227
Utilities / Energy	232
Wholesale Trade	242

*Totals may not sum to 100% due to rounding.

Methodology

United Kingdom
Japan
China
Germany
France
India
United States of America
Italy
Mexico
Spain
Saudi Arabia
Brazil

The **Cisco 2026 Data and Privacy Benchmark Study** offers a comprehensive analysis and view on the implications of the increasingly complex privacy landscape. This study is based on a survey of over 5,200 IT, technology and security professionals with data privacy responsibilities based in 12 markets across the Americas, Asia and Europe: Brazil, China, France, Germany, India, Italy, Japan, Mexico, Saudi Arabia, Spain, the United Kingdom, and United States.

Respondents represent 26 industries, including Educational Services, Energy and Utilities, Financial Services, Healthcare, Hospitality, Government, Manufacturing, Professional Services, Retail, Technology, Transportation, and Wholesale Trade.

This sample size is nearly double that of previous editions of Cisco privacy research, which could result in minor variances when comparing year over year. With a global sample of this scale, results are statistically significant to within approximately +1% at the 95% confidence level. Differences of 2% or more year-on-year represent a true change.

The research was conducted through a double-blind online survey in September 2025 and the analysis by an independent third party, Sandpiper Research & Insights.

Privacy pays:
Measurable ROI strengthens year on year

In recent years, the business case for data privacy has shifted from theoretical to undeniable. In the earliest iterations of Cisco's privacy research that began in 2018, most organizations were just beginning to observe the tangible financial and reputational returns from their privacy programs.

Privacy was still often framed as a safeguard – important, but primarily reactive. By 2025, that perception had transformed. Companies not only maintained their privacy budgets through periods of economic uncertainty; many have reported steady or increased investment as privacy began proving its worth as a strategic business imperative.

Now, in 2026, the data speaks for itself. Ninety-nine percent of organizations report measurable benefits from their privacy investments. Nearly all respondents credit their privacy programs with delivering concrete

% of organizations experiencing benefits from privacy-related investments

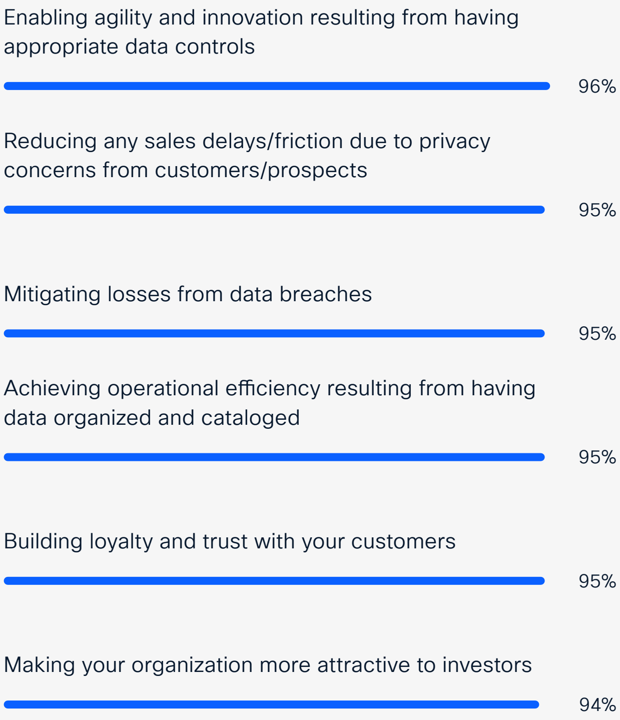

business outcomes: 96% report that enhanced data controls have unlocked agility and innovation, 95% say privacy initiatives have built stronger customer trust and loyalty, and another 95% have gained operational efficiencies through better data organization and management.

The trend line is clear: organizations with mature privacy programs outperform those that treat compliance as a check-the-box exercise. Privacy maturity has therefore become a reliable indicator of broader business performance, correlating with improved operational resilience, customer retention, and speed to market.

Even amid the rising complexity of AI, privacy continues to pay off. The organizations at the forefront recognize that privacy is not a cost to absorb, but a capability to leverage. By embedding privacy principles into how data is collected, governed, and used across AI systems, they are creating the conditions for sustainable digital transformation. In this new era, data governance isn't slowing innovation – it is what makes innovation possible. This trend is exemplified in the study findings with 'unlocking agility and innovation' rising to become the most-touted benefit of privacy investments – surpassing 'loyalty and trust' in 2025.

% of respondents using each metric to measure the value of privacy investments

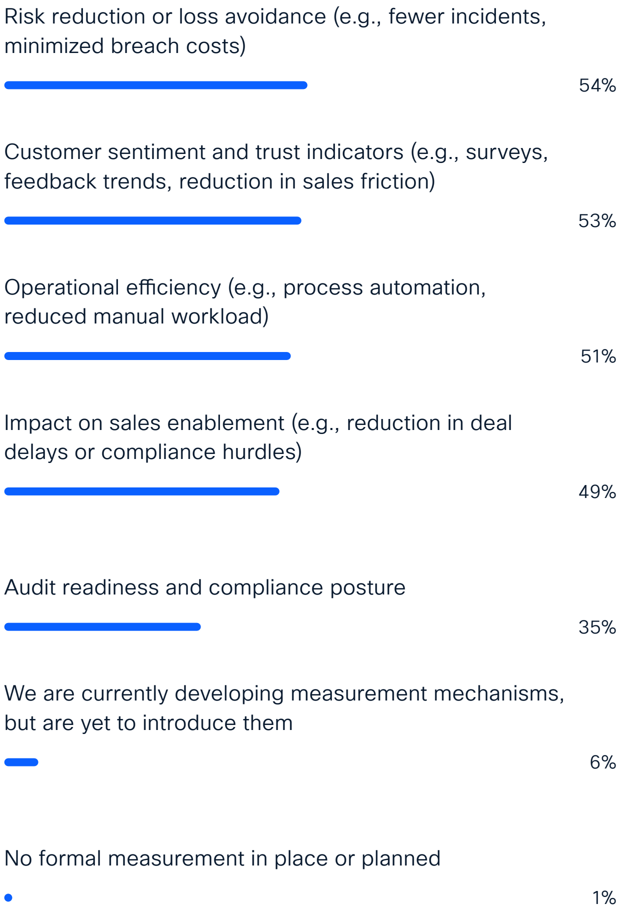

The shifting paradigm: Privacy evolves with AI

As AI reshapes how organizations collect, use, and secure data, the boundaries of privacy and data governance are rapidly expanding.

What began as a compliance conversation has evolved into an enterprise-wide governance challenge, one that demands coordination across legal, operational, technical, public policy, and ethical domains.

In 2026, nine in ten organizations (90%) report that their privacy programs have broadened in scope specifically because of AI, signaling a growing recognition that strong data governance underpins effective and trustworthy AI deployment. Investment is following this shift. Forty-three percent of organizations increased privacy spending over the past year, and an even larger share (93%) plan to direct additional resources to at least one area of privacy and data governance over the next two years.

% of organizations reporting expanded privacy scope due to the rise of AI

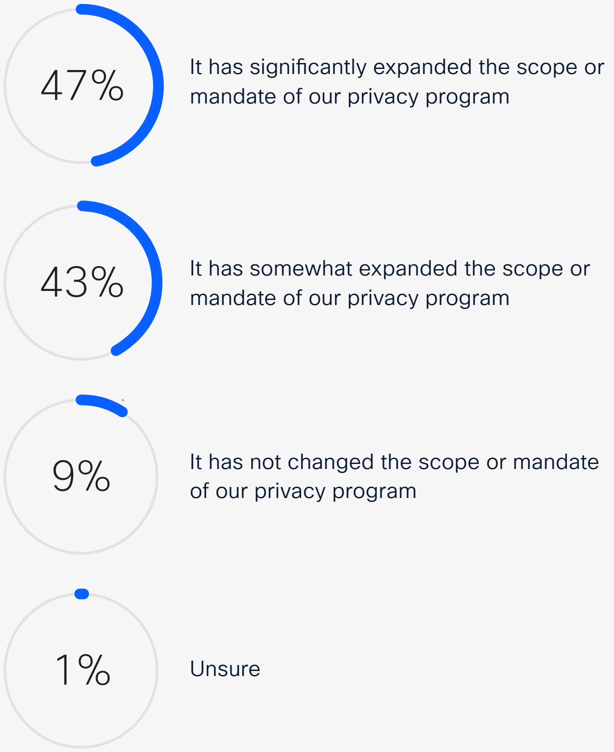

Areas of focus for greater spend are evenly spread across automation of compliance tools, using AI to monitor adherence to privacy and data governance requirements, and compliance with new and changing regulations.

Overall, 38% of companies globally spent $5 million or more on privacy in the past 12 months – an increase from 14% reported in early 2025. These figures underscore a decisive shift from reactive compliance to proactive capability-building.

And while investments are increasing across the board, the approach to governance is still evolving. Three in four organizations have a dedicated AI governance committee, and yet, only 12% describe them as mature and proactive.

This finding reveals an inflection point familiar from the **Cisco AI Readiness Index 2025** – a moment when aspiration outstrips infrastructure. Many organizations recognize the importance of governance but are still working to operationalize it across complex, fast-moving digital ecosystems.

Many organizations have expanded privacy programs and budgets, yet governance structures remain fragmented, and often IT or security-led, leaving gaps in cross-functional representation and operational maturity. The result is a hybrid landscape in which aspiration outpaces infrastructure.

When "by design" principles are embedded into AI governance, business is accelerated and trust strengthened. Respondents see 'achieving corporate values' (85%), and 'improved product quality' (85%) as the top benefits of AI governance followed by 'regulatory readiness' (84%), and 'enhancing employee relations' (83%). Privacy teams are central to AI and data governance, not peripheral.

% of organizations spending $5 million or more on privacy

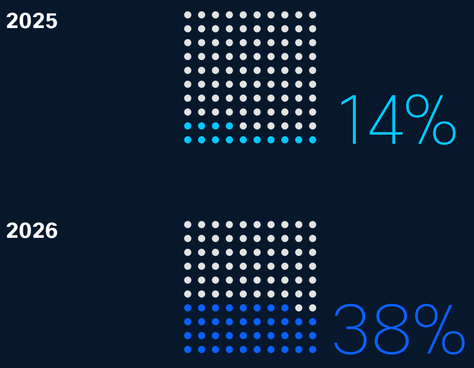

% of governance committees with representation from each function

IT/Technology	57%
Cybersecurity	42%
Legal, Risk, Compliance	35%
CEO-office	29%
IT	28%
Operations	26%
Human Resources	21%
Finance	21%
Engineering	16%
Product	8%

Top benefits of AI governance

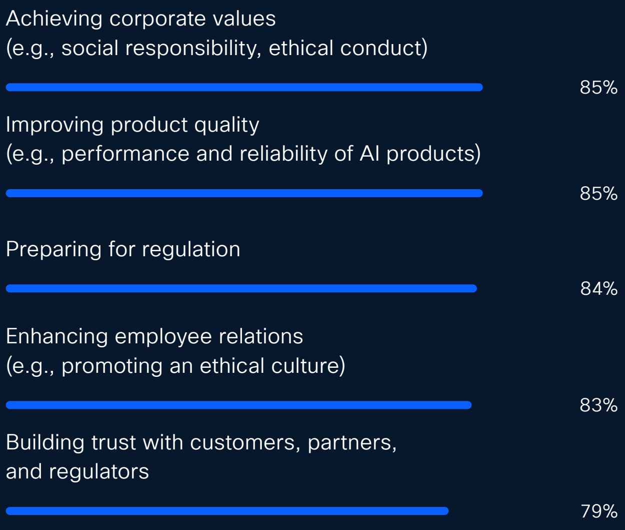

Beyond compliance:
Reframing data governance as a growth driver

Transparency has emerged as the single most powerful driver of customer trust. Almost half of organizations (46%) rank 'providing clear information about how data is collected and used' as the most effective action to build customer confidence, far ahead of 'demonstrating compliance with privacy laws' (18%) or 'avoiding data breaches' (14%). Ninety-seven percent report their transparency efforts with customers are proving effective, showing firms not only prioritize clarity but also see it deliver trust in practice.

Customer demand for transparency has risen sharply, with 85% reporting increases over the past three years. Firms are responding with more varied channels – over half (55%) now offer interactive dashboards that let users view or control their data in real time, and 50% embed transparency directly into contracts so customers and partners can clearly understand how AI systems use their information.

Top actions to build customer trust in data handling (% of organizations selecting each)

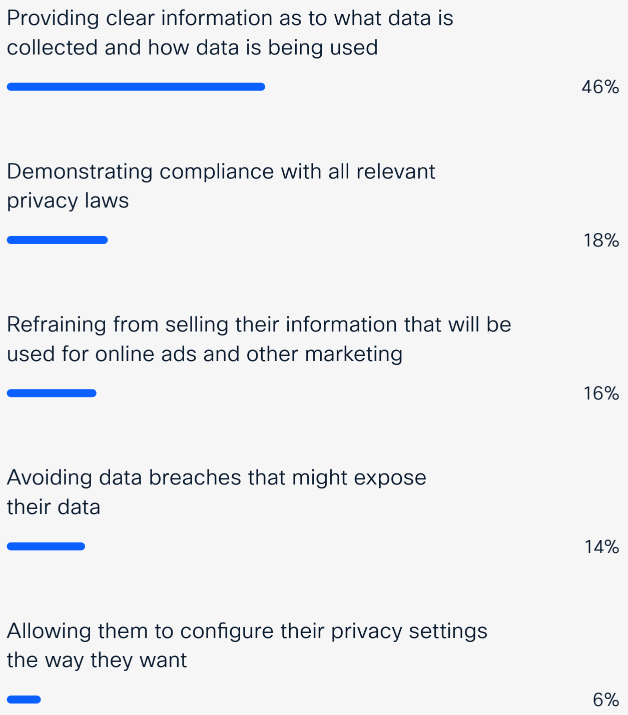

These approaches are not just improving perceptions; they are driving measurable outcomes. Organizations report that customers who understand a company's data-use policies are significantly more willing to share their information, especially for AI-driven services. In fact, 87% of respondents agree that strong privacy laws make their customers more comfortable engaging with AI applications. What was once seen as a regulatory burden is now recognized as a competitive advantage: compliance and trust are not opposing forces but mutually reinforcing goals.

The shift marks a fundamental turning point in the role of governance within digital organizations. As AI systems grow more complex and autonomous, customers increasingly expect accountability from the organizations deploying them. The most forward-looking organizations understand this and are embedding transparency into every layer of their governance model.

By treating trust as a measurable business outcome, these organizations are reframing data governance as a growth enabler.

"Transparent communication about how data is collected, processed, and safeguarded isn't simply a best practice; it's a business imperative and differentiator that defines market leaders."

Harvey Jang
Cisco Vice President, Deputy General Counsel, and Chief Privacy Officer

This, in turn, also supports an environment where data compliance measures are still largely welcomed rather than feared. Aside from 'the avoidance of fines and penalties,' around three-quarters (72%) of respondents generally assess the impacts of compliance with data privacy laws as an overall positive.

Estimated benefit from privacy-related investments over the past year

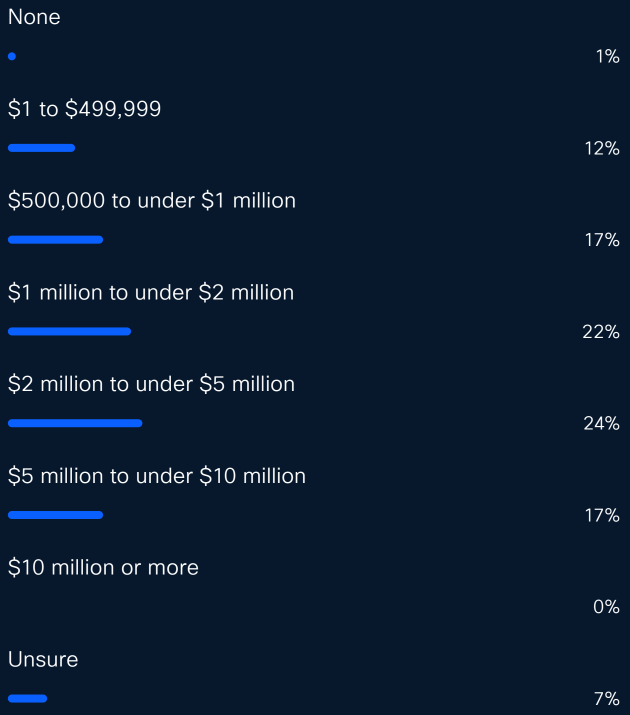

Organizational benefits realized through compliance with data privacy laws

Classifying risk	75%
Identifying business purpose for processing	72%
Identifying legal basis for processing	72%
Cataloging data	70%

Global pressure points:
Localization costs rise as trust globalizes

As organizations scale their digital operations across borders, data localization has emerged as one of the most significant — and complex — privacy challenges.

Since this report's inception, Cisco's research has examined how national data protection laws have shaped business models and operating costs. In 2026, data localization requirements have become increasingly common across markets. As their adoption grows, organizations are grappling with the practical challenge of operationalizing these requirements at scale, while continuing to balance regulatory expectations, infrastructure decisions, and cross-border trust.

Eighty-five percent of organizations say data localization adds cost, complexity, and risk to cross-border service delivery; similarly, 84% report it creates significant operational challenges. The effects are especially acute for multinational companies:

% of companies experiencing specific operational challenges due to localization

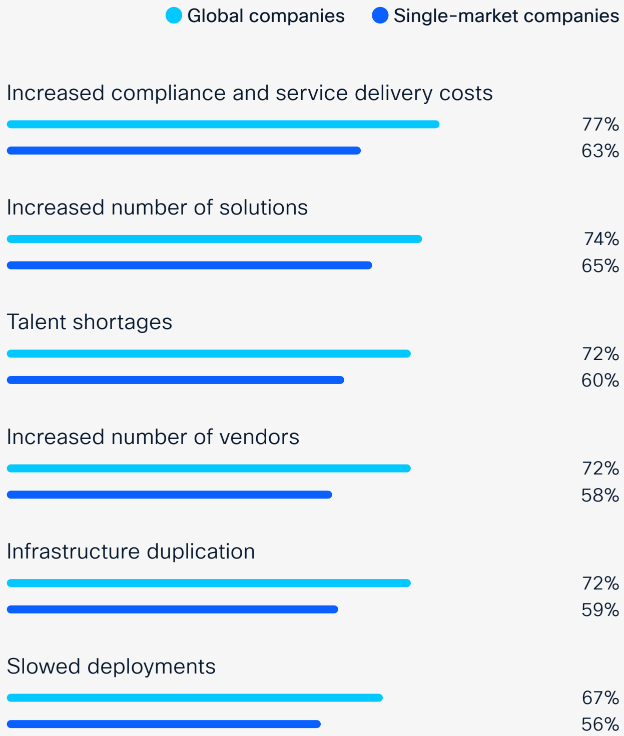

89% state that operating across multiple jurisdictions requires customized compliance strategies to address the inconsistencies in national privacy laws compared with 79% of companies operating in a single market. For global businesses operating in multiple markets, the result is not just higher costs but also slower service delivery, infrastructure duplication, and a growing risk of fragmentation in data management.

Expanding localization requirements tend to have a disproportionate impact on multinational organizations. Managing compliance across multiple jurisdictions can lead to slower service rollouts, more complex technology environments, and increased demands on scarce technical talent. At the same time, these requirements introduce greater fragmentation into data operations, challenging organizations to preserve consistency and scalability while continuing to deliver reliable, global services.

As AI adoption accelerates, this complexity has intensified. Seventy-eight percent of organizations now report increased costs linked to localization and data sovereignty because of AI developments. Seventy-seven percent say data localization requirements limit their organization's ability to offer seamless 24/7 services across markets, and the same percentage also say sustainability ambitions are hindered by localized infrastructure. Seventy-six percent predict these costs will rise even further in the coming year. Meanwhile, 81% already report a heightened demand for localization due to the rise of generative and agentic AI models that rely on massive, distributed datasets.

Amid this complexity, organizations are also refining how they assess data security in a multi-jurisdictional environment. While many continue to view local storage as an important safeguard, perceptions are becoming more nuanced. In 2025, 90% of respondents associated local data storage with greater security; in 2026, that figure has edged down to 86%. At the same time, confidence in global-scale providers is growing. Eighty-two percent of multinational organizations now believe that large, globally integrated providers are well positioned to manage and secure data flows across jurisdictions.

Organizations must be able to accommodate requests for localized services where necessary, but more broadly, progress will depend on continued public and private sector dialogue on regulatory frameworks that preserves safety and security while balancing flexibility and innovation.

% of organizations agreeing that data free flow with trust can boost economic growth and respect national sovereignty

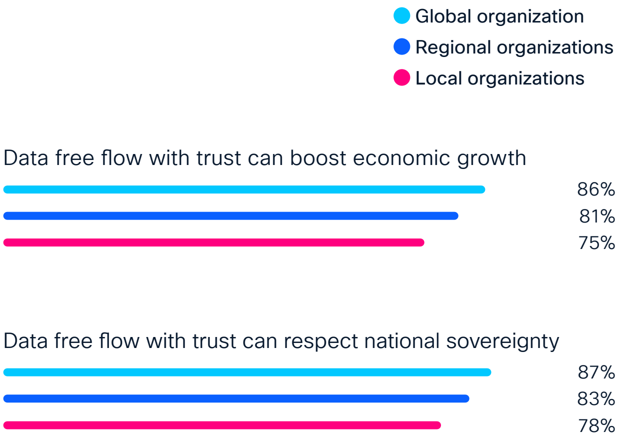

Many organizations are now advocating for global standards under the principle of Data Free Flow with Trust – a model that respects national sovereignty while promoting secure, interoperable data exchange. Eighty-three percent of respondents agree that harmonized global standards for data protection would reduce the need for country-specific localization laws, making compliance more efficient and sustainable. Initiatives, such as the Global Cross Border Privacy Rules Forum, operationalize these principles into a concrete action plan, turning Data Free Flow with Trust from aspiration into practice.

These shifts point to a broader reimagining of data governance at scale – one that prioritizes globally credible standards, measurable accountability, and integrated enterprise risk management as the foundation for responsible data practices and sustained digital trust.

For forward-looking organizations, this is the moment to rethink global data strategies, shifting from reactive compliance toward proactive alignment. Those that invest early in harmonized governance frameworks will not only reduce operational drag but also position themselves as trusted global players in the AI economy.

Data discipline under pressure:
Quality, tagging, and IP protection

As AI becomes embedded across the enterprise, the center of gravity in privacy has shifted. The question is no longer only how data is protected, but whether it is sufficiently understood, structured, and governed to support increasingly complex AI systems.

% of organizations with data tagging systems in place

Yes	No	Currently in development
66%	16%	18%

Sophistication level of tagging where in place

Comprehensive tagging system	51%
Limited tagging	33%
Customer-identified tagging	10%
Tagging in development	5%
Manual or ad hoc tagging	1%

Nearly seven in ten organizations report ongoing difficulty accessing relevant, high quality data efficiently, often citing the cost and effort of data preparation as a barrier to scaling AI initiatives. At the same time, concerns around intellectual property protection are intensifying. More than three quarters of organizations identify IP protection of AI datasets as a top governance challenge, reflecting the growing value – and sensitivity – of the data feeding AI systems.

Data tagging and classification remain critical enablers, but maturity is uneven. While two thirds of organizations report having a formal tagging system in place, only half describe their approach as comprehensive or automated. For many, tagging remains partial, manual, or fragmented, creating blind spots that complicate data use, governance, model oversight, and risk management.

These challenges underscore a broader reality: as AI systems rely more heavily on enterprise data, long standing gaps in data discipline become more visible, and more consequential.

Rising GenAI and agentic demand:
Reshaping data expectations

Compounding these foundational challenges is a sharp increase in data demand driven by generative and agentic AI. More than 80% of organizations report rising demand for both personal and non-personal data as AI use cases expand beyond experimentation into core business functions.

The types of data being used reflect this expansion. Organizations draw from a wide range of sources, including systems generated data, anonymized and aggregated customer information, telemetry data, synthetic datasets, and data acquired specifically for model training. This breadth reflects the diversity of GenAI applications, as well as the pressure teams face to source data that is timely, relevant, and fit for purpose.

At the same time, the growth in demand is outpacing clarity. Teams responsible for sourcing AI training data cite data quality as the leading obstacle,

Data types used in developing or using GenAI

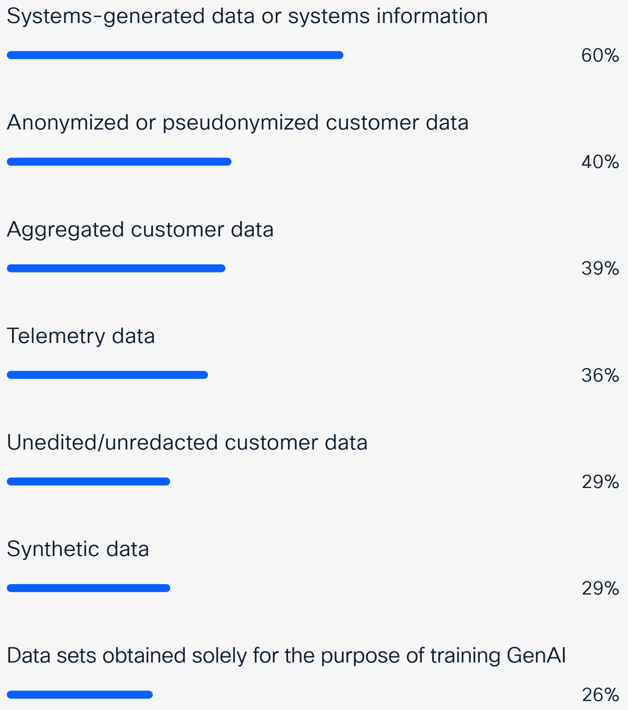

followed closely by unclear ownership and stewardship responsibilities. Localization requirements further add complexity, particularly for organizations operating across multiple jurisdictions.

The result is not a lack of intent, but an operational challenge: securing the right data, with the right permissions, under the right governance conditions, at the speed AI development requires.

% of organizations with outright bans on AI usage and limits on data entry into GenAI tools

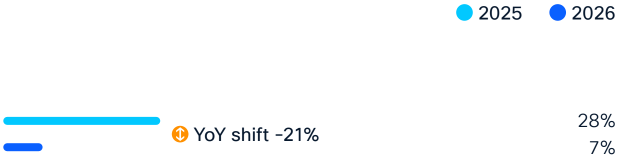

> **“Strong governance doesn’t slow innovation – it accelerates it, by creating clarity, accountability, and trust in how AI systems learn and operate.”**
>
> **Jen Yokoyama**
> Cisco Senior Vice President and
> Deputy General Counsel

From data sourcing to data use: Governance shifts closer to practice

As data demand accelerates, attention increasingly turns from sourcing to use. Many organizations acknowledge that proprietary business data and customer data are being incorporated into AI development pipelines, sometimes without full redaction or transformation. This reality reflects the practical pressures of building and deploying AI systems, rather than a deliberate disregard for policy.

It also helps explain why governance models are evolving. Between 2025 and 2026, organizations moved decisively away from blunt restrictions. Outright bans on generative AI use declined sharply, as did rigid limits on what data could be entered into AI tools. Experience appears to have shown that blanket prohibitions are difficult to enforce and often misaligned with how AI is used day to day.

Instead, governance is shifting closer to the point of interaction. Organizations are placing greater emphasis on user awareness, technical safeguards, and contextual controls that guide how data is handled in practice. Despite this shift, gaps remain. Most organizations still cite a lack of formal oversight, exposure of sensitive data, and insufficient privacy controls as top risks associated with GenAI usage, underscoring the need for governance approaches that match real world behavior.

Agentic AI expands the governance horizon

Emerging agentic AI systems raise the stakes further. While familiarity with agentic AI is high, active deployment remains limited. Even so, organizations recognize that autonomous and semi-autonomous systems expand the surface area for data access and decision making, moving governance beyond individual prompts to continuous workflows.

To manage this transition, organizations are adapting existing safeguards, including human validation, shared human AI responsibility models, pre action approvals, escalation thresholds, audits, and override mechanisms. Nearly half are extending current AI governance frameworks to agentic use cases, while others are beginning to build dedicated oversight structures.

This evolution reinforces a central lesson from the findings: as AI systems become more capable and autonomous, data privacy and governance do not become less relevant. They become more operational, more integrated, and more central to AI readiness.

Across data quality, sourcing, usage, and autonomy, the trajectory is clear. Effective AI deployment increasingly depends on disciplined, transparent, and well governed data practices that can scale alongside innovation.

Vendor governance:
Transparency and accountability lead to trust

As AI adoption accelerates, organizations are no longer operating in isolation. They depend on a growing ecosystem of technology vendors, cloud providers, and AI tool developers and deployers to power their digital transformation.

This shift has created a new frontier in AI and data governance – one where vendor transparency and contractual clarity are as critical to trust as internal controls. This study reveals that while confidence in AI providers is rising, formal mechanisms to safeguard that trust are still catching up.

Trust in AI vendors is strong: 81% of organizations say that generative AI providers have been transparent about how their tools use data, and 81% report that these vendors have clearly explained how their systems operate. Transparency, once a differentiator, has quickly become a baseline expectation. Most organizations now expect AI providers to articulate data-handling practices, clarify model behavior, and disclose how customer or proprietary information contributes to model training or outputs. This visibility is essential in a landscape where AI decision-making must be explainable not only to regulators, but to customers, partners, and employees as well.

Yet while confidence is high, formal accountability lags behind. Only about half (55%) of organizations require clear contractual terms outlining data ownership, usage rights, and IP parameters when working with AI vendors. Organizations may trust their vendors, but without formal guardrails, that trust is fragile. In a climate of rising regulatory scrutiny, informal assurances are no longer enough.

Forward-looking organizations are addressing this through a new discipline of responsible vendor governance. Nearly three-quarters (73%) now conduct active verification and ongoing monitoring to ensure third-party tools align with emerging AI regulations and responsible AI principles such as transparency, fairness, accountability, privacy, security, and reliability.

Meanwhile, 96% of respondents say that external, independent third-party privacy certifications play an important role in vendor selection – a sign that the market is rewarding verifiable privacy and compliance credentials.

The growing reliance on external AI providers also reinforces the need for greater contractual precision. Organizations are increasingly negotiating data usage limits, opting for configurable tools that allow them to control how their data contributes to model learning. Seventy-nine percent report that their generative AI providers are willing to adjust contract terms or tool configurations to limit data exposure – an encouraging sign that the ecosystem is maturing toward partnership and shared accountability.

The organizations leading in this space treat vendor transparency as an extension of their own governance frameworks. They recognize that accountability doesn't end at the enterprise boundary; it extends throughout the supply chain.

In the AI-driven economy, every partner, platform, and provider contributes to a shared trust equation. Companies that embed transparency into their procurement, contracting, and vendor management processes are not only protecting themselves from risk, they are also signaling to customers and regulators alike that they operate with integrity. As AI ecosystems grow more complex, this blend of confidence and accountability will become the new currency of digital trust.

% of organizations reporting that GenAI providers are willing to negotiate contracts or tool configurations to limit how data is used

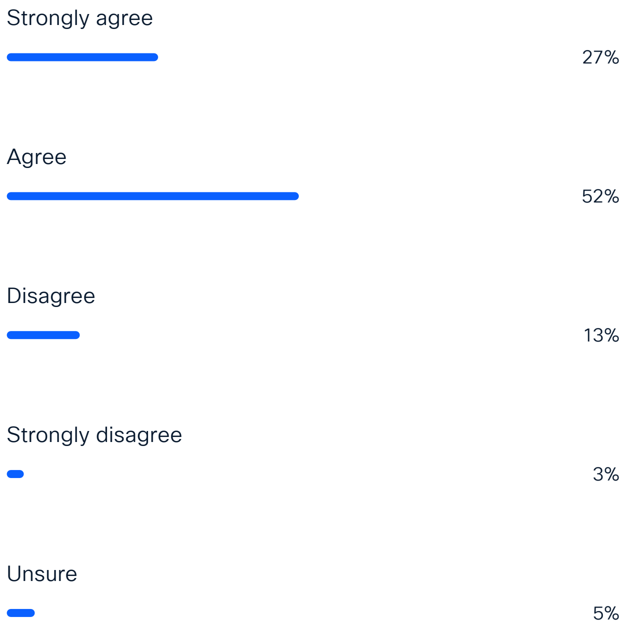

"What's emerging is a new model of trust by design – one that extends beyond technical controls into the legal and ethical frameworks that define digital relationships."

Dev Stahlkopf
Cisco Executive Vice President and
Chief Legal Officer

The convergence era:
Privacy, AI, and the governance horizon

Across nearly ten years of Cisco research, one message is clear: privacy, data governance, and AI accountability are converging into a single framework for digital trust. As AI becomes embedded across every operational layer, privacy leaders now find themselves at the center of a new governance ecosystem – one that unites compliance, security, ethics, and innovation under a broad umbrella of data responsibility.

This shift reflects the growing complexity of AI itself. The same technologies that drive personalization, automation, and insight also introduce new risks: bias in algorithms, opaque data usage, and the potential exposure of proprietary or sensitive information. To meet this challenge, organizations are expanding the scope of privacy and data management from protecting data to governing intelligence. Privacy officers are now collaborating with AI ethicists, cybersecurity experts, and data scientists to ensure that every dataset, model, and system operates within clearly defined ethical and regulatory boundaries.

The **Cisco 2026 Data and Privacy Benchmark Study** illustrates just how rapidly this integration is unfolding. Ninety percent of organizations report that their privacy programs have expanded because of AI, and nearly all see tangible benefits from these investments. Strengthened data controls and governance frameworks are not only reducing risk, they are also enabling cleaner, higher-quality data that directly improves AI performance. Data governance, once viewed as a safeguard, is now recognized as a strategic enabler of innovation and trust.

AI is also upgrading privacy operations. Respondents rank the top benefits of AI's integration into privacy functions as regulatory compliance (41%), automation of workflows (25%), and data discovery and classification (15%). This automation frees privacy teams from manual oversight and improves accuracy and consistency at scale. As a result, privacy is becoming more predictive, proactive, and embedded into the digital core of the enterprise.

The benchmark findings show that 95% of organizations report greater customer loyalty and trust as a direct outcome of privacy investment – a critical advantage in markets where transparency is increasingly demanded. Companies that clearly communicate how data is collected, used, and protected are strengthening customer loyalty, attracting privacy-conscious partners, and building more sustainable regulatory relationships. Transparency has become both a trust signal and a differentiator in AI adoption.

This convergence is not only structural, but cultural. Privacy, data, and AI governance are no longer confined to specialized teams – they are becoming shared organizational values.

Organizations that embrace this convergence are already demonstrating measurable performance advantages. Those with strong privacy and governance frameworks report higher agility, faster AI deployment, and greater success in scaling innovation responsibly. By embedding governance early, they create the resilience to adapt to new regulations, drive ethical AI use, and differentiate in increasingly competitive markets.

As AI continues to evolve from predictive systems to autonomous and agentic models, the governance horizon will expand again. The next phase will see the full integration of privacy, AI accountability, and cybersecurity into a unified model of trust-driven governance. Cisco's research across privacy, AI readiness, and cybersecurity consistently points to the same conclusion: value follows trust. The organizations that understand privacy and data management not as the barriers to progress, but as the foundation that enables it, will lead in the era of responsible innovation.

Conclusion and recommendations

Organizations that treat privacy, data governance, and AI accountability as interconnected pillars are better positioned to scale AI safely, manage regulatory complexity, and differentiate themselves in a market where transparency is now a competitive advantage.

Cisco's broader research into AI readiness and cybersecurity echoes this same theme: readiness is the true differentiator. Just as AI Pacesetters use disciplined investment and governance to turn ambition into measurable value, privacy leaders are using governance to turn compliance into confidence. They are proving that data governance is not a barrier to innovation; rather, it is the infrastructure that makes innovation possible.

Organizations that thrive will be those that build data governance into the DNA of their AI strategy. They will use transparency as a lever for trust, governance as a framework for scalability, and accountability as a source of competitive strength. In a world defined by intelligent systems and interconnected data, privacy is no longer a constraint on progress – it's the key to it.

Recommendations for organizations navigating privacy and data governance in the AI era:

1 **Prioritize data understanding and transparency**
Embracing a deep understanding of your data – its origin, classification, and usage – coupled with transparent communication about data practices, is fundamental for future-proofing your organization. This proactive approach not only builds stronger customer trust and loyalty but also positions your business to adapt swiftly and effectively evolving privacy, AI, and data regulations, turning compliance into a strategic advantage.

2 **Invest in robust data infrastructure**
A foundational investment in data infrastructure and comprehensive governance frameworks is non-negotiable for success in the AI era. This includes establishing consistency in data collection, format, labeling, and overall architecture. Infrastructure discipline ensures the high-quality, accessible data necessary to fuel AI innovation responsibly, mitigate IP protection concerns, and unlock operational efficiencies.

3 **Strategically evaluate data localization and infrastructure choices**
While data localization can address specific regulatory requirements, organizations must thoughtfully and deliberately assess the implications of local or on-premises data storage. Consider the potential risks associated with local providers, the scalability of infrastructure, and the true security benefits versus the costs and complexities of fragmentation.

4 **Establish a single, empowered AI governance body**
Beyond technical infrastructure, an organizational measure like a dedicated, cross-functional AI governance body is critical. It ensures that ethical considerations, regulatory readiness, and responsible AI principles and data handling are embedded into every stage of development and deployment, translating aspiration into actionable governance.

5 **Empower your workforce with training and safeguards**
Recognizing the human factor as a significant element in data security and responsible AI usage, invest in comprehensive workforce training and implement point-of-use safeguards. By fostering a culture of accountability and providing tools to prevent risky data exposure, organizations can mitigate potential data leaks and ensure that responsible innovation is a shared organizational value.

Meeting our customers’ standard of trust

Organizations have always required security to protect assets, help manage risk, and build customer confidence and loyalty. Privacy is a critical element of customer trust in today’s complex business environment. As customers set their standards of trust, Cisco continues to listen, learn, and evolve to meet those standards, prioritizing trustworthiness, transparency, and accountability throughout our holistic approach.

In addition to this report, Cisco also publishes Offer Disclosures for its major products and services, enabling anyone interested to understand what data is used, who has access to it, and how long it is retained.

Our Responsible AI Principles and Framework show how these principles and practices form our broad AI governance framework. And the Cisco Purpose Report and Cisco Purpose Reporting Hub offer relevant resources related to how we prioritize trustworthiness, transparency, and accountability in our environmental, social, and governance (ESG) initiatives.

All of this and more are available on the Cisco Trust Center.

For additional information about our privacy research, email the Cisco Privacy Center of Excellence at ask_privacy@cisco.com.

About the cybersecurity report series

Over the past decade, Cisco has published a wealth of security and threat intelligence information for security and privacy professionals interested in the state of global cybersecurity and privacy. These comprehensive reports have provided detailed accounts of threat landscapes and their effects on organizations, as well as best practices to defend against the adverse impacts of data breaches and privacy concerns.

In our approach to thought leadership and to help enable stronger cybersecurity and privacy globally, Cisco publishes a series of research-based, data-driven studies. We have expanded the topics to include varying reports for security and privacy professionals with distinct interests.

Calling on the depth and breadth of expertise from threat researchers and innovators in the security industry, the reports in each series offered annually include the **Cisco Talos Year in Review**, **State of AI Security**, **Cisco AI Readiness Index**, and **Cisco Cybersecurity Index**, with others published throughout the year.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**