IDC PERSPECTIVE

Preparing Data for Safe AI Enablement

Jennifer Glenn Ryan O'Leary

EXECUTIVE SNAPSHOT

The success of AI initiatives hinges on the quality, governance, and protection of enterprise data, making trusted data a foundational requirement. To ensure data is AI-ready, organizations must establish a secure foundation for data that is free from sensitive information and non-valuable data. Maintaining trusted data is an ongoing process to ensure clear usage policies and enforcement. The benefits of AI can be greatly outweighed by its risks of AI if safeguards are ineffective. Effective AI deployments also require prompt monitoring of AI outputs and rapid response to errors or hallucinations.

Technology buyers should prioritize data governance that includes data loss prevention (DLP) and data access governance (DAG) policies to AI agents, privacy-enhancing technologies, and foster cross-functional collaboration among data stakeholders. End users and suppliers should both embrace the idea of "least privileged context" when it comes to data seeding their AI activities. Ultimately, establishing a foundation of trusted, well-governed data is a prerequisite for scaling autonomous AI agents and minimizing organizational risk.

Key takeaways

- The success of AI initiatives is entirely dependent on the quality, governance, and protection of enterprise data, making trusted data a foundational requirement for effective and secure AI deployment.
- Organizations must actively reduce data volume, eliminate redundant, obsolete, or trivial (ROT) data, and rigorously protect sensitive and confidential information to minimize risk and ensure data quality for AI models.
- Organizations should embrace "least privileged context." Embracing data masking principles to allow AI and agents to access only the essential information within data sets needed to enable workflows.
- Maintaining trusted, AI-ready data is a continuous process that requires ongoing oversight, clear access controls, detailed change logging, data lineage mapping, and robust obfuscation techniques to safeguard intellectual property and sensitive information.

- Demonstrating trust in AI output demands prompt monitoring, rapid response to errors or hallucinations, and comprehensive governance frameworks that extend to AI agents, ensuring transparency, compliance, and resilience across the AI life cycle.

Recommended actions

- Implement robust data governance frameworks that encompass data discovery, classification, and data security posture management (DSPM) to ensure visibility and control over sensitive and confidential data across hybrid environments, supporting effective AI life-cycle management.
- Extend DLP and DAG policies to explicitly include AI agents as identities, ensuring comprehensive coverage and minimizing security gaps as AI adoption increases.
- Investigate data masking and synthetic data to enable "least privileged context."
- Establish cross-functional governance processes involving data security, management, and resilience teams to facilitate collaboration, shared objectives, and unified oversight, thereby enabling trusted AI data at enterprise scale.
- Prioritize the establishment of trusted, well-governed data foundations before scaling autonomous AI agents, as unverified or poorly protected data significantly increases organizational risk and undermines AI outcomes.
- Maintain backup and recovery systems that are air-gapped from production environments to ensure resilience against rogue agents and limit their potential blast radius.

SITUATION OVERVIEW

AI is quickly moving from experimentation to active implementation, requiring demonstrable results. Organizations are under pressure to use AI to augment daily operations to improve productivity, accelerate innovation, and offset staffing limitations. AI tools are positioned as the mechanism for accomplishing all three objectives simultaneously. Currently, AI tools are supplementing sales and finance teams, supporting product development, streamlining customer support, and, in some cases, assisting with cybersecurity operations. Half of the respondents to IDC's *Future Enterprise Resiliency and Spending Survey*, April 2026 indicated that AI agents are in production across multiple business areas, with another 27% indicating production in a single area. Further, 76.3% of respondents expect to have at least 10 agents deployed by the end of 2026, with 35.8% expecting to deploy at least 50 agents. In many organizations, AI is no longer a side initiative; it is a strategic imperative for digital business (see Figure 1).

FIGURE 1

Expected Deployment of Agents by the End of 2026

Q. How many agents do you expect to have deployed by the end of 2026?

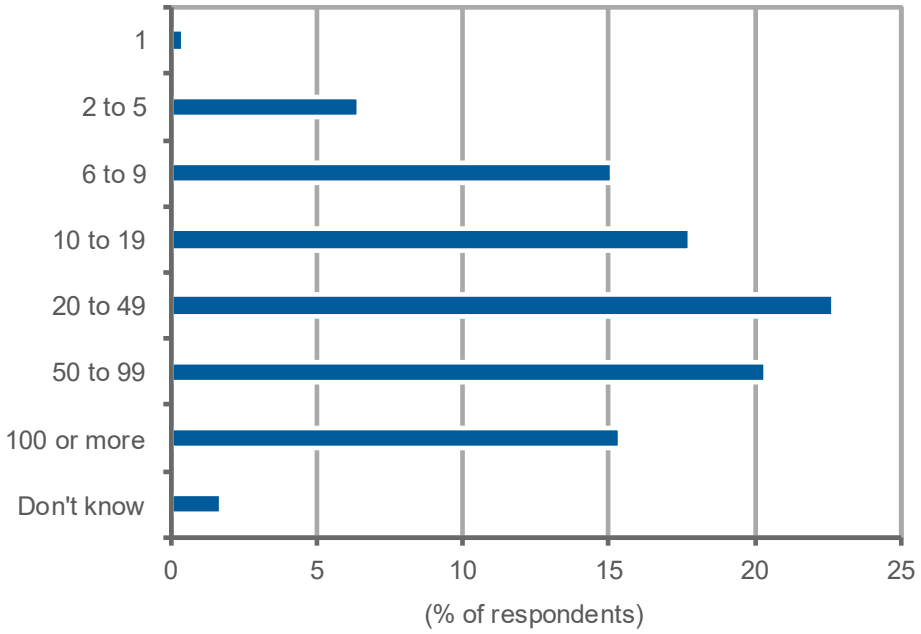

n = 1,009

Source: IDC's *Future Enterprise Resiliency and Spending Survey* Wave 2, March 2026

AI tools are nothing without data. The success (or failure) of AI initiatives is entirely reliant on the quality, availability, governance, and protection of enterprise data. Data is what fuels models, trains algorithms, informs decision-making, and generates outputs that may influence customers, partners, and regulators. This means that trustworthy organizational data must be available to the right users or applications whenever and wherever it is needed.

What is trusted data? Trusted data is valuable, relevant, and available. It also needs to be carefully curated, securely managed, and well-governed to ensure sensitive data is not exposed or exfiltrated.

Preparing the trusted data

Having the right data is foundational for trusted AI. Getting the right data often means *removing* what is no longer necessary to minimize risks of data exposure or loss.

Reduce volume: The increasing volume of data is not a new problem. The advent of AI tools has exposed two problems. First, users and applications are holding onto more data,

in some cases to train personal AI tools. Second, AI itself generates a massive amount of data. Excess data is expensive to store, obscures liabilities, and dilutes value. For AI specifically, bloated data pipelines degrade model performance, increase inference costs, and expand the blast radius of any breach.

Limit ROT: ROT data creates multiple problems for AI. In addition to increasing the attack surface, ROT data can undermine data quality. Multiple versions of data may produce biased results. Older data may include outdated assets that are no longer available and should not be used in AI output. An example of this might be insurance policies that are no longer available being offered as options for users to purchase. Eliminating ROT is not just a storage optimization — it is a data quality and security imperative.

Remove or mask sensitive and confidential data: IDC defines sensitive data as information included under regulatory rules, such as personally identifiable information (PII) or protected health information (PHI). Confidential data is defined as information that could cause material harm to an organization's operations if accessed or stolen, such as company secrets or intellectual property. Again, management of sensitive data is not a new challenge. Most organizations have been actively working for years to identify and monitor this type of data to adhere to compliance and privacy requirements. However, AI adds a new layer of complexity: this data must be protected from ingestion or dissemination into AI models and must not be accessible to unauthorized AI agents.

The other wrinkle is that many documents and pieces of data contain lots of PII, but also other information that is necessary for large language models (LLMs) to have access to. Least privilege access is the principle of granting every user, agent, and system only the minimum permissions necessary to perform its intended function and nothing more. This limits an attacker's ability to move laterally across systems if a single account or agent is compromised. This is no longer practicable for organizations that want to leverage data to seed their AI. According to IDC's *Future Enterprise Resiliency and Spending Survey*, April 2026, 49.5% of respondents are putting customer data into their AI instances. Under "least privileged access," none of these data sources would be available to the models.

This is where data masking and synthetic data can augment data protection. Masking sensitive information that LLMs and agents cannot access, while allowing context to drive insights. This is what IDC calls "least privileged context." Enabling models to use only what they need limits risk but enables insights and outputs of AI to not only be actionable but also trusted.

Ensure appropriate protection is in place: Data is the lifeblood of AI. When sensitive data must be used, it needs to be protected not only in production but also at rest. This includes encryption, backup, and recovery. It is impossible to predict what autonomous AI agents will do to data or how it will be used, so it needs to be encrypted and backed up.

Maintaining trusted data

AI is dynamic. Models evolve, accumulate new training data, and become significant sources of intellectual property. Agents learn new access patterns. The data ecosystem that feeds and trains AI is constantly changing. This means the risks and threats to this data are changing, too. Maintaining data trust is a continuous project.

Maintaining AI-ready data requires the same practices used to prepare it — volume reduction, ROT elimination, sensitive data protection — applied on an ongoing basis. It also requires continuous oversight of how data is used and by whom. This includes tracking every identity that touches the data.

Create dedicated source and training data

Curating data into dedicated sources for AI tools establishes provenance. As a model grows and learns, knowing where the data originated provides a clear foundation for reference if the output goes awry. Curated and dedicated sources offer a tighter management surface. This means controlled data collection that can be overseen and managed by human experts, fewer access points to control, and cleaner audit trails.

Clear access controls and use policies

Knowing who (or what) accesses data — from the source or training layer, through runtime, to the outcome — is necessary for auditing and governance. It also offers the level of transparency needed to build trust with users. Clear policies for role-based access and usage, as well as granular enforcement, are the keys to balancing data security with availability.

Continuous change logging and monitoring

In addition to access controls, a clear and detailed log of any changes to the data or model is essential for building trust and confidence in AI tools. When any source data is accessed, modified, or deleted, there should be a clear record of who made the change and when it was executed. Ideally, there would also be a reason provided for the change. This aids in detecting malicious activity, such as compromised credentials, and provides an auditable trail of information.

Lineage and data mapping

Ensuring continuous AI-ready data requires full knowledge of how data is used, where it is presented, and which applications use it. Data lineage and mapping answer the critical governance questions: Who created this data? Where did it come from? Where is it going? When was it last changed? This is essential for data and AI governance.

Data obfuscation

At the end of the day, protecting the data itself is the best way to limit exposure risk. Data used for model training, as well as data within the trained model, should be considered intellectual property and treated appropriately. Obfuscation techniques, including tokenization, masking, and anonymization, ensure that sensitive values can be used for AI processing without being exposed in plaintext. These controls are particularly important for AI systems that surface data to external users or partner-facing applications.

Demonstrate trusted output

When AI output is incorrect, or worse, shares sensitive or inappropriate information, end users will notice. Monitoring output is also a great way to demonstrate transparency and trust. Initial rollouts of AI tools will not be perfect, so responding quickly to make adjustments will be an important factor in the success of the project.

Prompt monitoring

Prompt monitoring helps detect attempts to extract sensitive information, bypass data governance policies, or manipulate model behavior through injection attacks. It also provides an audit trail that supports compliance and incident response. In an interesting twist, a recent IDC survey found that 52% of organizations are using AI tools to augment data security functions for user behavior analytics and prompt analysis (*IDC IT Data Management Quick Poll*, February 2026, [n = 107]).

Protect sensitive data from being exposed

Even in the best circumstances, sensitive data may slip through or may be inadvertently created. Monitoring both the prompt and the output can provide a clear context for where sources should be modified or access privileges adjusted.

Limit and respond quickly to hallucinations

When hallucinations happen, responding quickly to identify and fix the problem is paramount. Building and maintaining a foundation of AI-ready data should make this process faster and more accurate.

ADVICE FOR THE TECHNOLOGY BUYER

Technology buyers should evaluate the security of data environments against the full lifecycle of AI solutions — from initial preparation through ongoing maintenance, to output. The following guidelines will offer context and priority for investment and policy decisions.

- **Know where and what your data is.** DSPM tools offer visibility needed to discover sensitive and confidential data across complex hybrid environments. Effective

governance and AI outcomes are not achievable without knowing where the sensitivity of data is or where it lives.

- **Context is king.** Data security tools, such as discovery, classification, and DSPM, provide valuable information — such as provenance and lineage — about organizational data. This context is essential for granular policy enforcement that balances control with availability.
- **Extend DLP and data access governance to AI pipelines.** DLP and DAG solutions must be configured to account for AI agents as identities. Policy frameworks that cover only human users will leave significant gaps as agentic AI proliferates.
- **Build cross-functional governance processes.** AI is a true unifier for data stakeholders. IDC research shows that data security, data management, and even data resilience teams share similar objectives and challenges. Each group may have different success metrics, but shared information and collaboration about organizational data are required to deliver trusted AI data at enterprise scale.
- **Treat trusted data as a prerequisite for AI autonomy.** Organizations should establish trust in their data foundation before scaling agentic AI. Autonomous agents operating on unverified, ungoverned, or poorly protected data amplify risk exponentially. Trust must come before autonomy.
- **Deploy data masking technology to enable "least privileged context."**

LEARN MORE

Related research

- *RSAC 2026: Data Security Is a Top Priority for Agentic AI* (IDC #US54471026, April 2026)
- *IDC Survey Spotlight: Data Security and AI Workloads Fueling a Move to On-Premises Infrastructures* (IDC #US54353926, March 2026)
- *Strengthening Your Data Privacy Program: IDC's Future Enterprise Planning Guide* (IDC #US54244926, January 2026)
- *IDC's Data Security and Privacy Survey: The Integration of Data Security and Privacy Gets Stronger* (IDC #US53802025, September 2025)
- *Data Classification Taxonomy: It Is All About the Data* (IDC #US53526025, June 2025

Synopsis

This IDC Perspective emphasizes that the reliability and competitiveness of AI initiatives hinge on the quality, governance, and protection of enterprise data. Organizations must rigorously curate, secure, and continuously monitor data to ensure it is AI-ready, minimizing

risks and ensuring compliance. Robust governance frameworks, cross-functional collaboration, and proactive oversight are essential to maintain trusted data foundations, enabling effective, transparent, and resilient AI deployment across the enterprise.

"AI is only as trustworthy as the data it consumes," says Jennifer Glenn, research director for Information and Data Security. "Without secure, curated data, every AI-driven decision becomes a potential risk."

ABOUT IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology, IT benchmarking and sourcing, and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives. Founded in 1964, IDC is a wholly owned subsidiary of International Data Group (IDG, Inc.).

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com

Copyright notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, and web conference and conference event proceedings. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/about/worldwideoffices. Please contact IDC at customerservice@idc.com for information on additional copies, web rights, or applying the price of this document toward the purchase of an IDC service.

Copyright 2026 IDC. Reproduction is forbidden unless authorized. All rights reserved.