database
TRENDS AND APPLICATIONS
SUPER GUIDES

Redefining Trust: Data Security and Governance Strategies for the AI Era

As AI becomes central to enterprise strategy, the importance of trusted data continues to grow. Organizations are facing increasing pressure to secure sensitive information, ensure compliance, and maintain data integrity across complex and distributed environments. At the same time, AI introduces new challenges that require more dynamic and scalable approaches to governance.

This DBTA Super Guide highlights the strategies and technologies organizations are using to strengthen data security and governance in the AI era. By embedding governance into the data lifecycle and leveraging automation and intelligence, organizations can reduce risk while enabling faster, more confident innovation.

SPONSOR

INDUSTRY SNAPSHOT

Redefining Trust: Data Security and Governance Strategies for the AI Era

1. Agentic AI—State of the Industry

The continuing rise of AI in all its forms—machine learning, generative, and agentic—puts pressure on organizations and data managers to deliver data sources that are secure, timely, contextual, and well-aligned with the business. AI relies on large datasets for training, inference, and operation, making governance policies' security essential to assuring the trustworthiness of data. Increasing AI regulation, growing adoption, and greater risks require more comprehensive approaches to data security and governance.

- **Most organizations don't have AI governance formulated and do not feel secure about AI for mission-critical applications.** Only 27% responding to a survey by the Cloud Security Alliance (CSA) indicate they are confident that they can secure AI in core business operations. Only 26% have comprehensive AI governance policies in place, which jumps to 44% for larger enterprises (CSA, "The State of AI Security and Governance," 2025).

- **AI and AI governance have become foundational to security work.** Almost all (99%) of security processes now use AI, a recent survey shows. Half have formalized, active AI policies in place. Another 42% are making progress. AI is seen as highly effective for many critical tasks. However, concerns about security and compliance are the biggest barriers to effective automation, which limits teams' ability to scale impact (Tines, "Voice of Security," 2026).

- **Tellingly, most organizations that experienced security incidents lack formal governance policies.** Close to two-thirds (63%) of organizations reporting a data breach within the past year did not have a formal AI governance policy in place, an IBM study found. In addition, 97% of AI-related security breaches involved AI systems that lacked proper access controls. Even when organizations have a policy, less than half have an approval process for AI deployments, and 61% lack AI governance technologies. Among organizations that have governance policies in place, only a minority (34%) perform regular audits for unsanctioned AI (IBM, "Cost of a Data Breach Report," 2025).

- **Few mature governance models exist.** While agentic AI usage is poised to rise sharply in the next 2 years, oversight is lagging: Only 1 in 5 companies in a January 2026 Deloitte survey has a mature model for governance of autonomous AI agents ("The State of AI in the Enterprise").

ADOPTION OF GOVERNANCE GUIDELINES

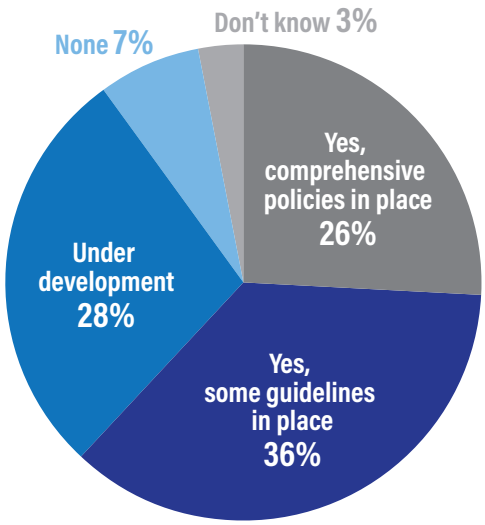

Source: Cloud Security Alliance

2. AI Data Security and Governance—Top Challenges and Opportunities

OPPORTUNITIES:

While strong governance and security policies will help ensure AI is delivered with the lowest levels of risk possible, AI itself also acts as an engine to ensure greater data security.

- **AI governance is the "maturity multiplier" driving responsible adoption.** According to CSA's research, organizations with formal AI governance are significantly more advanced in their AI journeys. They are two times more likely to adopt agentic AI, three times more likely to train staff on AI security tools, and two times more confident in protecting AI systems.
- **AI is driving down the cost of data breaches.** The global average cost of a data breach dropped to $4.44 million from $4.88 million in 2024, representing a 9% decrease and a return to 2023 cost levels, according to IBM's data breach report. "Faster identification and containment of breaches—much of it from organizations' own security and security service teams, with help from AI and automation—drove this decline."
- **AI enables greater intelligence and less manual intervention in governance and security processes.** AI systems can rapidly sift through data to spot anomalies or security issues—tasks that previous took days and weeks to complete. This will lower compliance reporting costs and provide for real-time protection against security threats.

CHALLENGES:

Moving forward with AI governance and security is not an overnight project. It requires a comprehensive approach with collaboration from across the enterprise.

- **AI security faces a maturity gap.** Even with clearer ownership emerging, organizations continue to face significant challenges in building the skills and risk understanding required to secure AI effectively. Organizations cite understanding AI risks (61%), skill gaps (53%), and lack of knowledge among current staff (52%) as the top hurdles to getting started with security for AI implementations (2025 CSA survey).
- **Zero-trust data governance will advance to meet the AI data challenge.** Within the next 2 years, 50% of organizations will adopt zero-trust data governance as unverified AI-generated data grows, Gartner predicts.
- **Shadow AI is creating security issues.** The swift rise of shadow AI has displaced security skills shortages as one of the top three costly breach factors tracked by the IBM data breach report. At least 20% of respondents said they suffered a breach due to security incidents involving shadow AI. For organizations with high levels of shadow AI, those breaches added $670,000 to the average breach price tag compared to those that had low—or no—levels of shadow AI. These incidents also resulted in more personal identifiable information (65%) and intellectual property (40%) data being compromised. "[T]hat data was most often stored across multiple environments, revealing just one unmonitored AI system can lead to widespread exposure."

3. AI Security and Governance—Views from the Experts

As organizations move AI into production, questions around security, governance, and trust are coming into sharper focus. Two industry experts offer their perspectives on what it takes to manage AI responsibly while continuing to innovate.

"The consolidated, custom-governed data classification, integration, cleansing, masking, and reformatting operations in IRI Voracity support the preparation of reliable, compliant model data. Even the logs from these jobs can help you measure, enforce, and sustain trust in AI, too."

—David Friedland, IRI SVP

"The organizations getting real value from AI aren't just experimenting with models, they're building disciplined data foundations. Governance, quality, and security aren't barriers to innovation, they're what make it scalable."

—Stephen Faig, Research Director Database Trends and Applications

Everyone is excited, and even enchanted, by the power of AI, and is looking to the technology to lift their teams and organizations to new levels of achievement. That's all fine, but AI will stop dead in its tracks—or worse, create chaos—without a solid foundation of governance and security.

—Joe McKendrick, Lead Analyst, Unisphere Research.

Creating Trust in AI Data

IRI Voracity
An Insatiable Appetite for Data

David Friedland, IRI SVP

As AI adoption accelerates, trust in the data behind it lags. As DBTA highlights in this guide, only 40% of organizations have invested in the governance, controls, and transparency needed to support trustworthy AI systems, and 87% report rising AI-related vulnerabilities. The IRI Voracity data management platform can help close this gap between AI ambition and the operational readiness required for trustworthy AI.

Enterprises with Voracity can discover, integrate, cleanse, secure, and wrangle data—all from a single interface and within a custom-governed runtime infrastructure. And as organizations are rethinking how they protect and manage data in increasingly diverse environments, and prepare it for AI models, they need end-to-end capabilities that foster and demonstrate trust in that data.

THE PROBLEM: FRAGMENTED TOOLS AND RISING AI-DRIVEN RISK

Today's data ecosystems are sprawling across cloud, hybrid, and multi-platform environments. This fragmentation creates blind spots around data lineage, governance, and AI-related threats.

Consider your:

- Disjointed data classification, integration, masking, quality, and reformatting tasks
- Limited visibility into data flows and transformations performed in separate tools
- Challenges to enforcing data security and compliance controls
- More brittle pipelines that hinder AI model building and AI adoption
- Difficulty proving data provenance, quality, and explainability

These challenges undermine the trust required for responsible AI and regulatory alignment.

THE SOLUTION: A UNIFIED PLATFORM FOR GOVERNED, EXPLAINABLE DATA

IRI Voracity addresses these challenges by consolidating data discovery, integration, migration, governance, and analytics into one platform. Its design aligns directly with the Super Guide's focus on zero-trust data access, automated governance, observability, and explainability.

KEY CAPABILITIES

AI-Aware Data Security & Testing

- Through built-in data classification and masking facilities like DarkShield, Voracity finds and masks sensitive data with encryption, redaction, pseudonymization, fabrication, and other anonymization functions for content-aware DLP and data privacy law compliance. Pipelining its data obfuscation or synthesis functions protects or replaces sensitive data before it reaches AI models or analytics.

AUTOMATED GOVERNANCE & METADATA MANAGEMENT

- Built-in metadata hubs, business glossaries, and policy-driven controls ensure consistent governance across Voracity jobs.

END-TO-END DATA LINEAGE & OBSERVABILITY

- Voracity describes and records data workflows, transformations, and dependencies. Its PII discovery dashboards, mapping diagrams and job scripts, machine-readable audit files, and log wrangling tools, enable transparency and auditability.

HIGH-PERFORMANCE DATA INTEGRATION & TRANSFORMATION

- Powered by IRI CoSort, Voracity accelerates ETL workloads and combines processing steps into one fast I/O pass.

EXPLAINABLE, TRUSTWORTHY DATA PIPELINES

- Through data profiling, cleansing, and lineage, Voracity users can demonstrate the provenance and reliability of the data feeding their AI systems

FASTER INSIGHTS, LOWER RISK, AND TRUSTED AI OUTCOMES

Voracity results:

- Up to 10x faster data processing than SQL and ETL tools
- Significant reductions in data risk across databases and files
- Consolidation of multiple tools
- Improved data quality and trustworthiness for AI and analytics
- Lower operational costs and simplified governance

By unifying data processing, protection, and proof, Voracity users are speeding information delivery, modernizing confidently, and maintaining the controls required for trustworthy AI.

For more information on preparing data for AI, see: https://www.iri.com/blog/business-intelligence/prepare-and-protect-data-for-ai-with-voracity/

ACTION CHECKLIST

Building AI Governance and Security Foundations

Building an AI governance and security foundation is an enterprise-scale endeavor intended to manage and ensure that AI efforts are secure and aligned with the business.

Here are the four key elements needed for delivering a well-governed and secure AI implementation, as outlined in the NIST Artificial Intelligence Risk Management Framework Playbook:

1. GOVERN. A CULTURE OF RISK MANAGEMENT IS CULTIVATED AND PRESENT.

- Be aware of legal and regulatory requirements.
- Integrate AI governance into existing organizational governance and risk controls, particularly as they relate to sensitive data.
- Create standards for experimental design, data quality, and model training, as well as model testing and validation processes. Provide descriptions of training data.
- Pursue monitoring, auditing, and regular reviews. This applies to information related to AI contact information, business justification, scope, usages, and potential risks and impacts.
- Develop policies and procedures related to AI system performance and trustworthiness, accounting for bias and security problems.

2. MAP. CONTEXT IS RECOGNIZED AND RISKS RELATED TO CONTEXT ARE IDENTIFIED.

- Establish teams that reflect a wide range of skills, competencies, and capabilities for AI efforts. Empower these teams to capture, learn, and engage the interdependencies of deployed AI systems and related terminologies and concepts from disciplines outside of AI practice such as law, sociology, and art.
- Define the tasks, purposes, and benefits of the AI system.
- Identify potentially more suitable non-AI or even nontechnology alternatives.
- Determine the downstream events, such as decision making, that may be impacted by AI performance.

3. MEASURE. IDENTIFIED RISKS ARE ASSESSED, ANALYZED, OR TRACKED.

- Focus on detecting, tracking, and measuring the known risks, errors, incidents, or negative impacts of AI.
- Identify appropriate testing procedures and metrics, particularly as they demonstrate AI trustworthiness.
- Define acceptable limits for system performance and include course-correction measures.
- Develop metrics to ensure transparency regarding system design, development, deployment, use, and evaluation. This includes ensuring transparency within the lines of communication with AI teams.
- Maintain records of errors, incidents, and negative impacts of AI systems.

4. MANAGE. RISKS ARE PRIORITIZED AND ACTED UPON BASED ON A PROJECTED IMPACT.

- Apply AI risk factors to existing data management, quality, and privacy controls.
- Make trustworthiness a top priority, and document AI system performance related to trustworthiness.

- Monitor risks and benefits throughout the AI system lifecycle.
- Evaluate AI performance and trustworthiness as applied to real-world use cases.
- Make sure AI systems with lower risk tolerances receive greater oversight and management resources. Re-calibrate as necessary.
- Apply regulatory, organizational, and professional standards to AI risk management.
- Develop mechanisms and procedures for capturing feedback about AI-based incidents.

Additional Resources

- For more information on preparing data for AI, see: https://www.iri.com/blog/business-intelligence/prepare-and-protect-data-for-ai-with-voracity/