White paper

Building enterprise trust for back office

IRON MOUNTAIN®

Before organizations can safely deploy and scale agentic AI, they need a firm grasp on what their AI systems are doing, why they are doing it, what information they use, and who is accountable. This level of **enterprise trust** can only be achieved through **compliance by design**—weaving governance directly into back-office operations.

Contents

03/ Barriers to realizing agentic AI benefits

04/ Compliance cannot be a checkpoint

05/ Compliance by design supports back-office operations and decision intelligence

06/ Next steps: building the enterprise trust layer

Barriers to realizing agentic AI benefits

Enterprise automation is transitioning from rigid, task-based automation to dynamic agentic systems capable of autonomous decision-making and cross-workflow execution. Autonomous systems can now access data, invoke tools, and execute workflows. A finance agent can investigate an invoice discrepancy. A procurement agent can evaluate supplier documentation. A facilities agent can interpret maintenance records and service-level requirements to prioritize intervention.

The opportunities are substantial, but so are the risks. Without enterprise trust, organizations may lack the confidence and control needed to allow AI agents to make business decisions. In fact, McKinsey's 2026 AI Trust Maturity Survey found that nearly two-thirds of leaders cite security and risk as key obstacles to scaling agentic AI.

In related research, the Everest Group reported that the gap between data governance and AI governance is "where audit, operational, and reputational risks accumulate." The risk profile grows considerably when AI agents must make decisions using unstructured and structured data.

Across shared services, **four core friction points** undermine enterprise trust.

1. **Fragmentation** occurs when information spans physical archives, legacy repositories, systems of record, and data clouds. Agents can access data without knowing what's authoritative, causing validation delays. Processing cycle time slows, increasing cost per transaction, delaying cash conversion, and ultimately causing margin leakage. Establishing enterprise trust requires connecting agents to the right information by establishing unified access to authoritative physical and digital evidence before reasoning begins.

2. **Context gaps** open as unstructured documents lack the metadata and relationships AI agents need to reason correctly. Agents cannot reliably connect a document to the right customer, supplier, or transaction—lowering throughput and increasing cost-to-serve. Overcoming this friction requires supplying the right context, enriching unstructured assets with domain logic, obligations, and relationship metadata so agents can accurately interpret intent and business state.

3. **Governance gaps** exist around documents with unclear or varying permissions, retention rules, ownership, consent, and jurisdictional constraints. Agents default to manual escalation or halt entirely, driving up legal review costs and delaying contract cycle times. Addressing this friction demands embedding the right controls—integrating policy-as-code, dynamic permissioning, and human-in-the-loop (HITL) review directly into the execution flow rather than relying on manual checkpoints

4. **Explainability gaps** leave regulators and boards unable to validate AI-driven outcomes. Without traceable evidence trails, an organization cannot show how an agent reached a decision or confidently deploy that agent in key workflows—limiting scalability and reducing return on AI investment. Overcoming this barrier requires traceable outcomes, with immutable audit trails and lineage records that prove compliance to regulators and leadership.

After-the-fact compliance review is no longer sufficient for back-office operations. To overcome the friction points, organizations need to know—and show—why a decision was made, which evidence and rules informed it, and who or what is accountable.

With governance built into workflows, compliance by design becomes the foundation for enterprise trust.

Compliance cannot be a checkpoint

> By 2027, 40% of enterprises will demote or decommission autonomous AI agents because governance gaps are identified only after production incidents occur.
>
> —Gartner

Are AI agents ready to navigate the complexities of multi-region shared services centers? These organizations already struggle with cross-functional data discrepancies and inconsistent metadata across enterprise resource planning (ERP) systems, procurement platforms, and legal repositories.

Agents must adhere to localized privacy controls such as the General Data Protection Regulation (GDPR), along with policies for cross-border data transfers, data sovereignty requirements, and regional retention mandates. Frameworks like the EU AI Act and localized cloud-storage requirements expect organizations to document data use, maintain governance and traceability controls, and help ensure AI workflows do not move or expose regulated data outside approved environments.

Imagine an AI finance agent investigating an invoice exception. It finds a contractual provision allowing a disputed charge and approves the invoice. Was the contract current? Was the agent permitted to access it? Did the provision apply in the relevant country? Was the supplier's legal entity correctly matched?

With incomplete governance, an agent could reason correctly and still produce the wrong outcome —triggering an incorrect payment, processing an invalid charge, or missing a duplicate invoice. When these errors leak into production, the business pays a double penalty: immediate financial losses through unearned payouts and missed early-payment discounts, followed by delayed processing across the broader workflow as panicked teams freeze automation to manually re-examine flagged transactions. Operations slow when manual, post-hoc reviews and "audit fatigue" set in—impacting velocity, inflating compliance management costs, and exposing the organization to regulatory penalties.

Compliance by design supports back-office operations and decision intelligence

In a recent survey of business and IT leaders, Deloitte reports that only 21% say their organizations have mature governance for agentic AI. The numbers make the case for compliance by design—building policies, regulatory schedules, and data rules into the processing layer rather than retrofitting controls.

This approach moves back-office records from passive to active assets—de-risked for agentic AI execution and structured for the decision intelligence agents need to reason well, not just operate safely. Beyond safeguarding compliance, embedding these guardrails directly into the workflow delivers critical operational efficiencies. It reduces tedious manual reviews, accelerates processing speeds, and empowers finance teams to automate high-volume workflows at scale without sacrificing control. By continuously embedding compliance rules and policy updates into workflows, organizations stay audit-ready with a clear accountability record.

Behind the scenes, the enterprise trust layer's architecture includes built-in zero-copy system access, unified metadata, and policy-as-code functionality.

- **Zero-copy system access** allows data to be securely accessed from the enterprise system where it resides without creating additional copies that increase risk.
- **Unified metadata** provides a consistent way to classify information across systems so the organization can maintain governance wherever that information is used.
- **Policy-as-code** means compliance requirements are built into automated processes to apply controls consistently and automatically—rather than bolting them on after the fact.

Finally, HITL machine learning acts as a core trust mechanism for exception management and continuous model retraining. Qualified employees review ambiguous, high-risk decisions, and their corrections create an auditable feedback loop that improves model performance while preserving accountability and control.

Next steps: building the enterprise trust layer

As AI moves from assisting employees to executing complex work on behalf of the business, one clear imperative emerges: establish the right conditions so agents can be trusted to act at scale. Enterprise trust is the ultimate catalyst for organizational scalability; without embedded governance, scaling autonomous agents merely scales operational risk and compliance friction. By establishing a policy-aware trust layer, global back offices can confidently expand agentic decision intelligence across regions, business units, and high-volume workflows without increasing headcount or exposing the enterprise to regulatory penalties.

Iron Mountain InSight® DXP forms the core of that architecture, unifying back-office solutions and systems through intelligent, non-disruptive automation with built-in policy management. Partnering with organizations to operationalize compliance by design, Iron Mountain brings a 75-year track record of securely handling document-centric assets and complex processes for companies worldwide. The Iron Mountain InSight DXP content intelligence platform was recently recognized as runner-up for the Shared Services and Outsourcing Network's "Technology of the Year" award. Furthermore, The Forrester Wave recognized the InSight DXP platform for its agentic AI functionality, document labeling, and data privacy, and for pairing everyday automation with advanced AI agents and reliable human-in-the-loop controls to keep data highly accurate while protecting customer privacy.

Ready to scale your back-office operations safely? Iron Mountain experts offer no-obligation discovery and demo sessions.

Iron Mountain empowers your organization to achieve enterprise trust

Trust the content

- Clean up redundant, obsolete, trivial data and resolve metadata discrepancies
- Facilitate the continuous application of global retention policies and data privacy controls across connected repositories
- Empower C-suite leaders to safely scale agentic decision intelligence across a global enterprise

Govern the action

- Employ specialized back-office AI agents to operate within strict guardrails
- Leverage ERP connectors and human-in-the-loop controls to prevent unapproved actions
- Eliminate audit fatigue

Prove the outcome

- Generate immutable evidence packets, audit trails, and defensible disposition logs to prove regulatory compliance to external auditors and regulators across back-office operations
- Recover lost profit margins

Learn more about how Iron Mountain supports secure, scalable AI-driven workflows for the enterprise back office.

IRON MOUNTAIN®

800.899.IRON | ironmountain.com

About Iron Mountain

Iron Mountain Incorporated (NYSE: IRM) partners with more than 240,000 customers across 61 countries, including approximately 95% of the Fortune 1000. We serve as the trusted guardians of what matters most, helping unlock what's possible. By connecting physical, digital, and intelligent workflows, our lifecycle solutions help organizations realize enterprise value at scale. From complex challenges across information management and digital transformation to data security, data centers, and asset lifecycle management, we empower organizations to navigate an ever-changing landscape to prepare for tomorrow's possibilities. Our longstanding commitment to safety, security, sustainability, and continuous innovation defines everything we do. Visit www.ironmountain.com for more information.

© 2026 Iron Mountain, Incorporated and/or its affiliates ("Iron Mountain"). All rights reserved. Information herein is proprietary and confidential to Iron Mountain and/or its licensors, does not represent or imply an invitation or offer, and may not be used for competitive analysis or building a competitive product or otherwise reproduced without Iron Mountain's written permission. Iron Mountain does not provide a commitment to any regional or future availability and does not represent an affiliation with or endorsement by any other party. Iron Mountain shall not be liable for any direct, indirect, consequential, punitive, special, or incidental damages arising out of the use or inability to use the information, which is subject to change, provided AS-IS with no representations or warranties with respect to the accuracy or completeness of the information provided or fitness for a particular purpose. The Iron Mountain solutions are designed to provide tools and a technical framework to assist in meeting your compliance goals. Use of the solution does not guarantee compliance with any law or regulation. "Iron Mountain" is a registered trademark of Iron Mountain in the United States and other countries, and Iron Mountain, the Iron Mountain logo, and combinations thereof, and other marks marked by ® or TM are trademarks of Iron Mountain. All other trademarks may be trademarks of their respective owners.