White paper

Breaking through the data ceiling:

How content intelligence unlocks agentic AI and accelerates business value

IRON MOUNTAIN®

Contents

03/ Executive summary

05/ The productivity tax silently draining the ROI from AI

11/ Governance as a multiplier

12/ Business impact: the accelerator advantage

13/ Conclusion

Executive summary

Every enterprise IT leader has heard the same promise: artificial intelligence will transform decision-making, unlock productivity, and deliver revenue growth. Yet for most enterprises, high-value AI deployment remains stalled in the pilot phase. The reason is rarely the large language model (LLM)—more often than not, it's the fragmented content feeding it.

New primary research from Iron Mountain tracks how enterprise leaders are modernizing fragmented data estates to build content intelligence into their operations. Benchmarking organizations across three distinct archetypes—**Custodians**, **Connectors**, and **Accelerators**—the study exposes a widening divide: market leaders (Accelerators) are pulling ahead of their peers (Custodians and Connectors), leveraging a more mature content intelligence foundation to unlock faster decision speed, AI readiness, and measurable ROI.

- **Custodians (low content intelligence maturity - 17% of organizations surveyed)**: This group is characterized by manual processes, siloed data, and legacy technology. Because their data remains fragmented, their estate creates AI agent blind spots that stall adoption at the pilot stage.
- **Connectors (moderate content intelligence maturity - 75% of organizations surveyed)**: For these organizations, data and technology modernization is underway, but plagued by AI pilot fatigue, they are experiencing some gains or business impact. They can support basic AI use cases, but disconnected data, systems, and governance prevent scaling autonomous agents enterprise-wide.
- **Accelerators (high content intelligence maturity - 8% of organizations surveyed)**: These are our market leaders who are furthest ahead in AI innovation, leveraging higher rates of modern cloud architecture and automated workflows to deliver enterprise insights and measurable business performance. By building an AI-ready content foundation, they have established an active trust layer required to scale specialized AI agents.

What is content intelligence?

For most enterprises, meaningful business information remains trapped inside siloed, disconnected legacy content. With years, or even decades of records locked inside non-digitized or unstructured files, businesses are stalled in activating that information for advantage.

Content intelligence is an organization's capability—powered by AI-native technology—to discover, interpret, and activate information trapped in dark data like contracts, invoices, emails, and physical archives. The right approach and foundation turn previously static files into structured, governed, AI-ready assets that fuel models, specialized agents, and automated workflows.

How does it drive ROI?

Content intelligence maturity is associated with stronger business outcomes. Our survey reveals that Accelerators (leaders with more mature content intelligence) are 53% more likely to report improved AI readiness to scale beyond pilots and 36% more likely to report increased revenue growth than their global peers.

The divide between these groups is structural. **Accelerators** are not just managing content better; they are re-architecting how information is governed and fed directly into critical business workflows.

By investing in a consolidated information layer, **Accelerators** are breaking through the enterprise data ceiling to scale AI successfully and deliver higher revenue growth than their counterparts.

Percent of respondents that:	Custodians	Connectors	Accelerators
Say they have the ability to scale AI beyond pilots	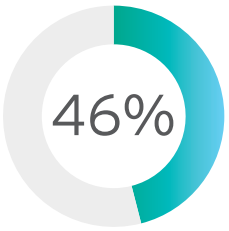	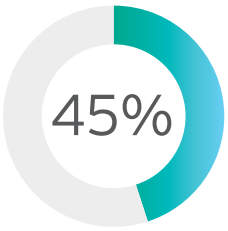	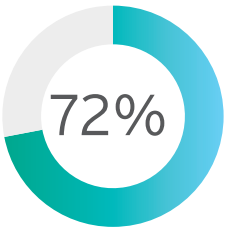 (+53%*)
Say their organization has experienced increased revenue growth over the last 12 months	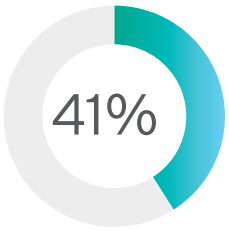	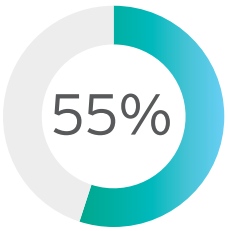	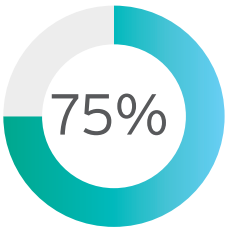 (+36%*)

This paper details several key differences between these groups, exposing how data fragmentation levies a tax on enterprise capital, resources, and AI investments. It also reveals how **Accelerators** are breaking free from legacy drag to unlock higher performance and value.

The research for this paper was conducted as a global study of 400 IT decision makers which included both qualitative and quantitative interviews. See page 15 for more details on the research methodology and demographics.

* Reflects the difference versus the global baseline established across all global respondents.

The productivity tax silently draining the ROI from AI

While executives typically look for technical barriers to AI, a bigger bottleneck is often hidden in plain sight; it is the operational friction—such as manual wrangling, missing metadata, and legacy search hurdles—caused by a disconnected content estate. AI pilots that begin with promise often hit a wall when legacy systems fail to deliver clean, trusted, accessible data. As a result, internal teams get trapped in manual data work, stalling project momentum.

Across all IT leaders surveyed, **84% agree** that highly skilled IT and data science personnel are burning hours acting as manual data wranglers instead of executing strategic work.

This disparity sharpens as you focus more deeply across the groups studied.

Why this kills agentic AI scale:

Manual data wrangling isn't just a waste of expensive technical hours—it prevents enterprises from safely scaling autonomous AI agents. Deploying specialized AI agents on fragmented, dark data is an unacceptable operational gamble. Without clean, governed context, multi-step workflows become unreliable, prone to hallucinations, and vulnerable to permission leaks.

Percent who agree skilled IT/data personnel are acting as data wranglers

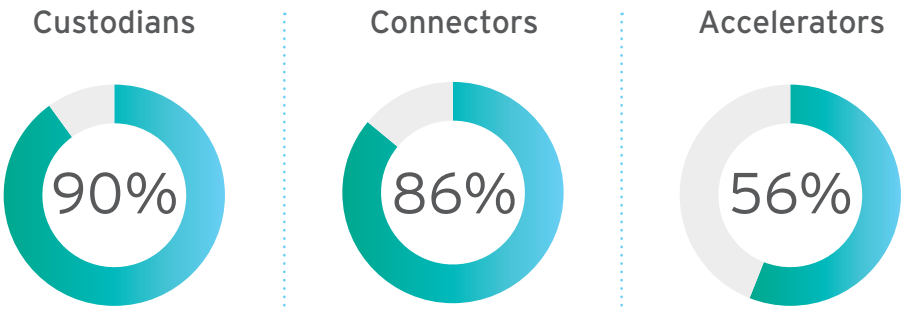

* Reflects the difference versus the global baseline established across all global respondents.

The productivity tax isn't contained to IT and data teams. When IT and data scientists spend project cycles prepping documents, extracting metadata, normalizing formats, and manually validating sources, the economics of enterprise AI collapse. Strategic modernization is often deferred, extending the resource drain to line of business (LOB) employees.

Average of LOB time lost to manual data verification and search hurdles

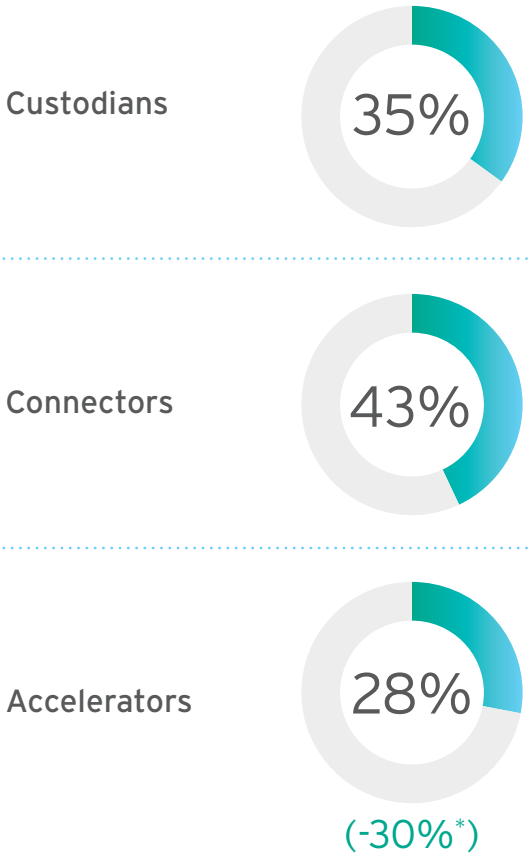

The difference between these groups isn't ambition or focus—it is the maturity of their data estate, and how advanced they are in modernizing their technology and workflows.

"Metadata is missing from many documents. Old contracts and invoices often lack the creation date or author. Searching often takes longer than doing the work itself."

—Senior IT leader, Healthcare and Life Sciences

* Reflects the difference versus the global baseline established across all global respondents.

The cost of waiting on actionable data

For an enterprise, decision speed is a competitive advantage. When business processes rely on insights buried in unstructured content—from core business documents like contracts and claims to audio, video, and media assets—getting trusted data to humans becomes a challenge. By transforming unstructured content into AI-ready data, organizations can trust AI agents to execute end-to-end workflows with minimal human intervention, driving faster decisions for the business.

> "We face issues with data scattered around and not integrated for better tracking and management. Our biggest challenge has been dealing with data scattered across different systems and teams. It slows down the use of AI."
>
> —IT Leader, Banking

Percent of respondents that:

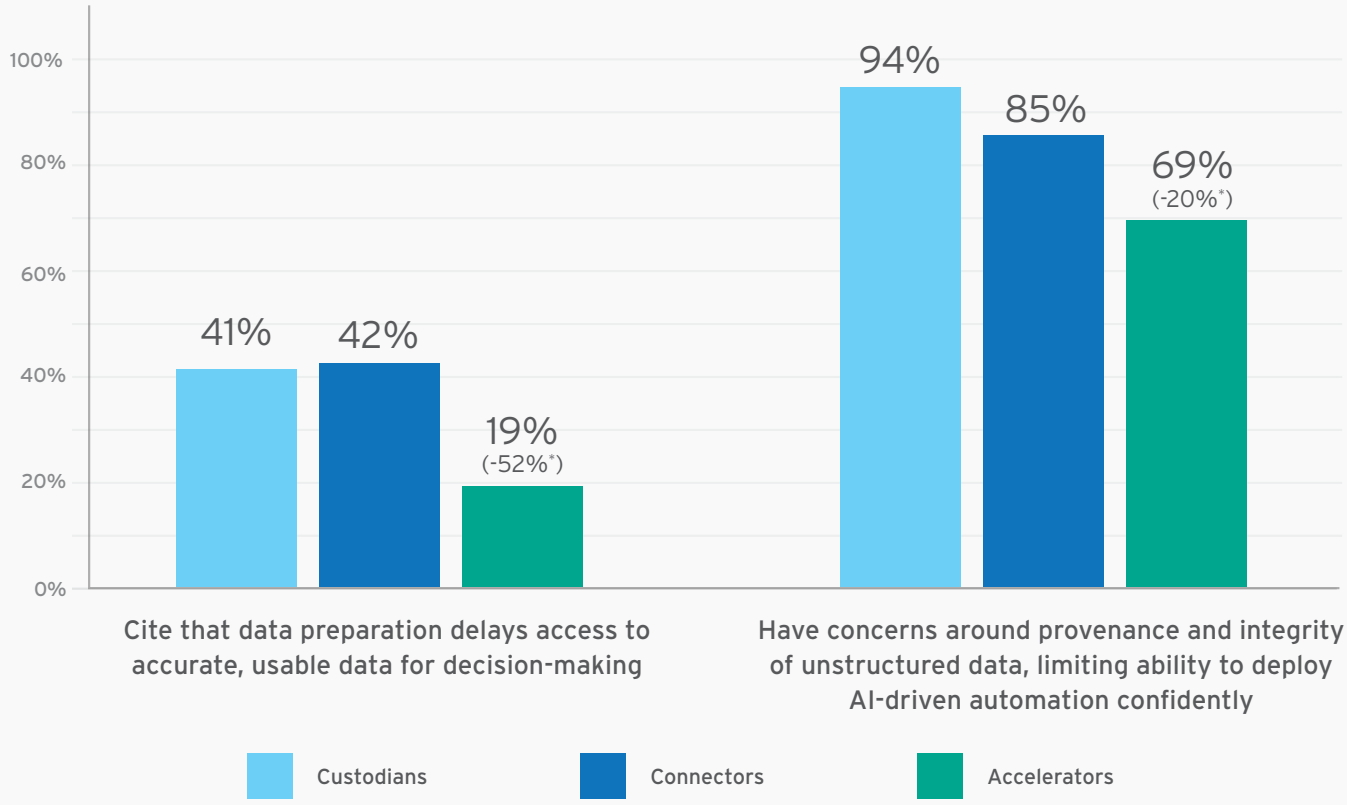

Custodians and **Connectors** are paying for this in real time. Both groups are experiencing longer delays in accessing accurate, usable data for decision-making and greater concerns about the provenance and integrity of their data, limiting their ability to deploy AI-driven automation.

Accelerators are breaking through. They are making decisions with greater speed and scaling AI with conviction, backed by content and IT estates engineered for modern intelligence.

* Reflects the difference versus the global baseline established across all global respondents.

The burden of legacy data and technology

Our research reveals a deeper divide: lower-maturity organizations are bleeding more than lost productivity maintaining unmanaged data. They are spending capital maintaining legacy content systems which cannibalize the budget needed to invest in innovation and advanced AI-enabled platforms. Beyond the financial drain, these legacy systems create AI agent blind spots—isolating unstructured content from the agentic layer required to transform enterprise operations.

> "The biggest challenge is compatibility between legacy systems and new technology solutions. Many of the systems we use were developed more than ten years ago and were not designed to handle the large volume of unstructured data we generate today. Migrating this data, without disrupting essential public services, is extremely complex and risky."
>
> —Senior IT leader, Government sector

Average estimated proportion of enterprise data across all organizations surveyed that is:

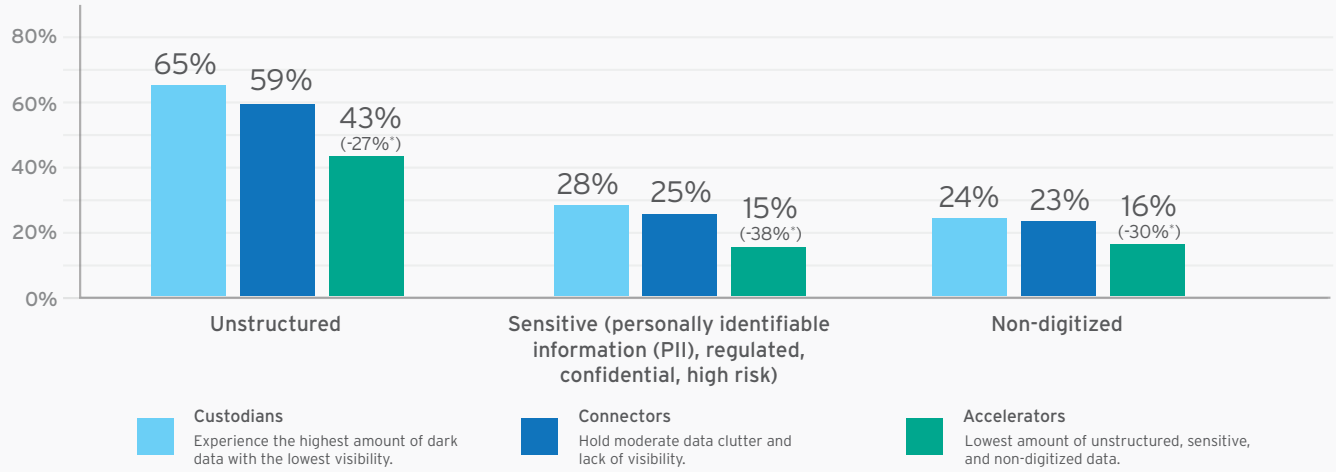

High performers are also prioritizing the transition from legacy ECM platforms and on-premises network drives to modern, AI-powered tooling.

Percent of respondents who say their organization is:

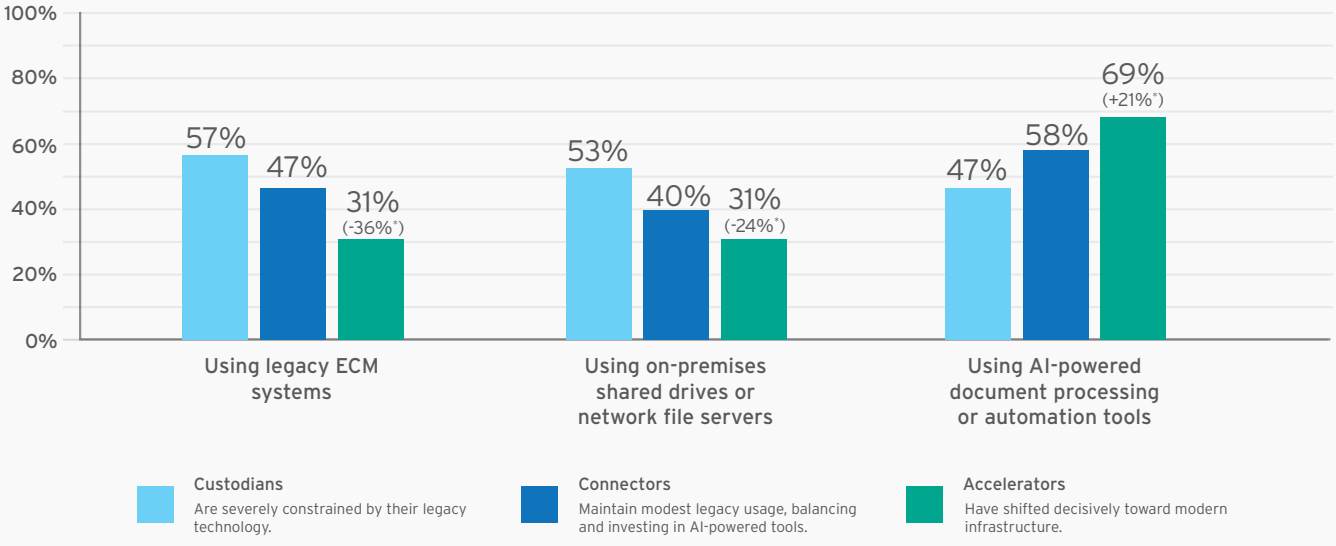

* Reflects the difference versus the global baseline established across all global respondents.

The impact on IT budgets and ROI

Spending on maintaining legacy systems used to store, manage, and access unstructured data places a higher drain on IT budgets, impacting each enterprise group's ability to fund modernization and scale AI initiatives.

Average current IT budget spent maintaining legacy systems for storing, managing, and accessing unstructured data

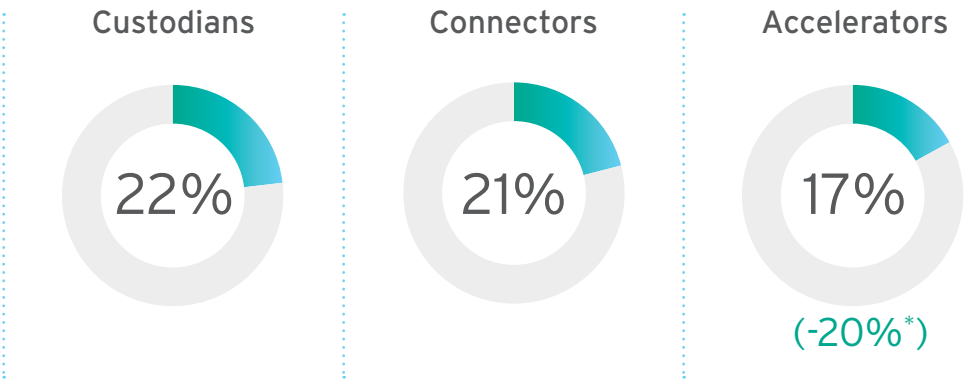

Rather than sinking capital into maintaining static repositories, **Accelerators** are free to reallocate IT budgets to modern automation tools. This shift addresses one of the most persistent bottlenecks in enterprise AI: scalable data extraction.

Rigid, legacy extraction tools collapse under the volume, velocity, and variety of information modern enterprises require to run agentic-powered workflows. **Accelerators**, however, are further ahead in building a unified informational layer. Deploying AI-driven ingestion engines that automatically classify, extract, and structure unstructured content allows Accelerators to overcome data extraction limits and accelerate AI deployment.

Percent of respondents that cite:

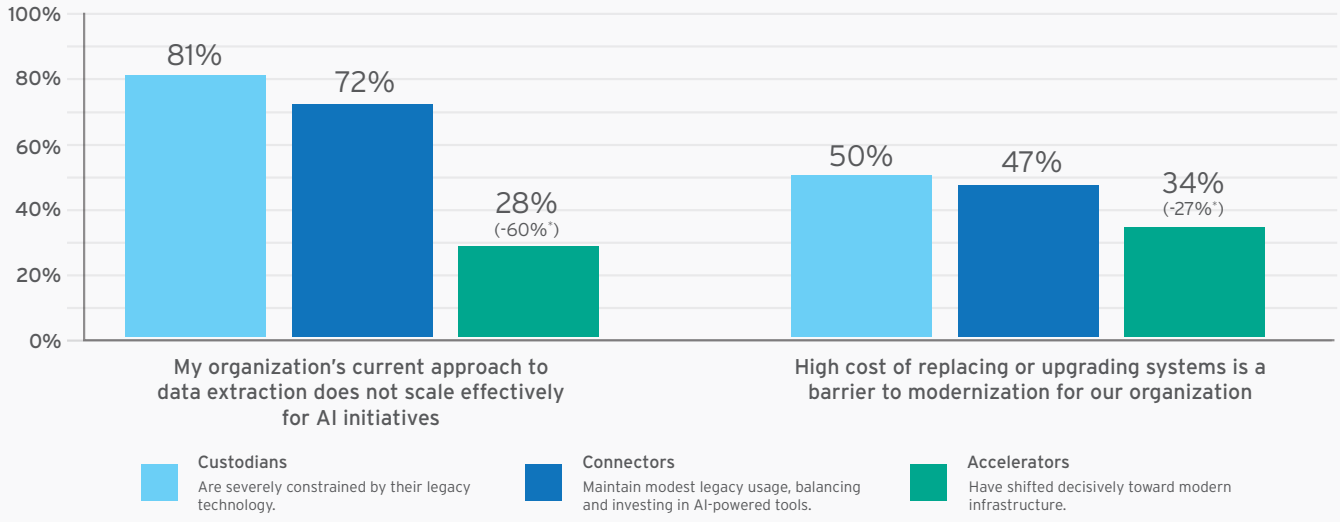

* Reflects the difference versus the global baseline established across all global respondents.

Accelerators likely have not inherited cleaner data and technology environments. Instead, they made a deliberate trade-off, reallocating investments from legacy maintenance to a unified content intelligence layer that serves both current operations and future AI initiatives.

Not satisfied with the lead they have already built, **Accelerators** are ahead in demonstrating ROI from current investments and have higher levels of future investment planned in the coming year.

Percent of respondents that say:

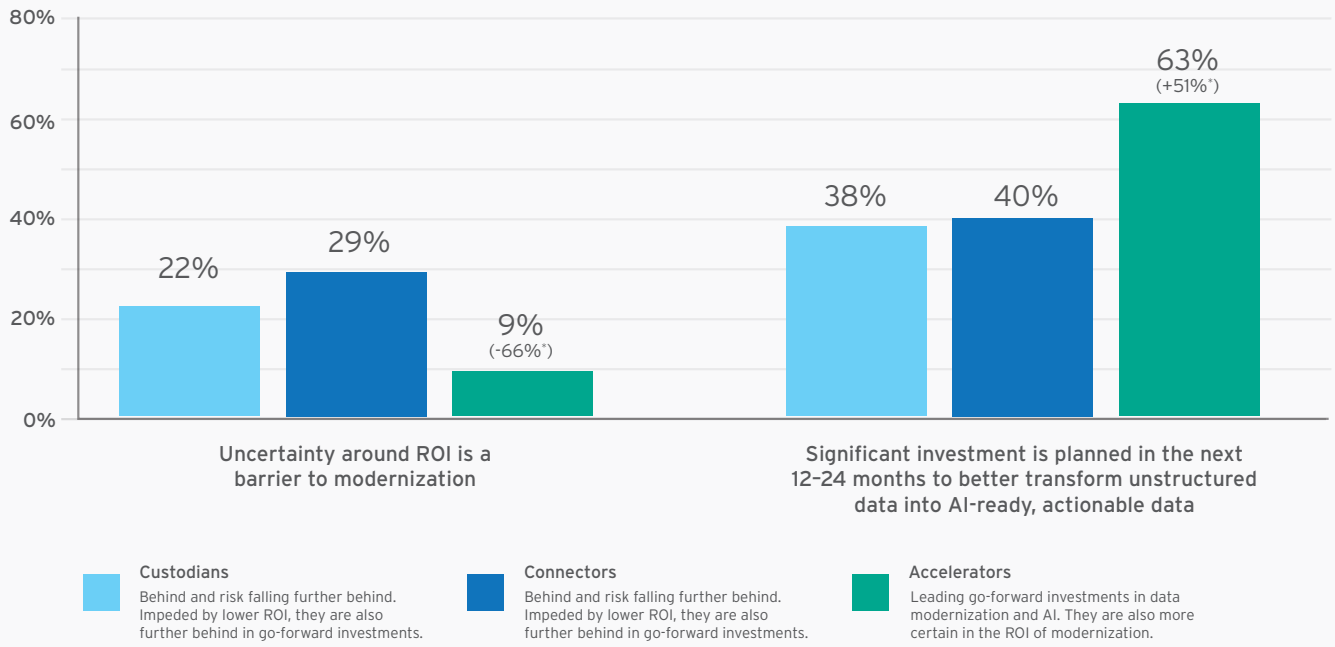

The takeaway for leadership: Legacy modernization and AI investment are not competing budget lines—for **Accelerators**, they're the same budget line. Organizations that treat a modern content intelligence foundation as a prerequisite for AI, rather than a parallel workstream, free up budget locked in maintenance and build confidence to invest further.

* Reflects the difference versus the global baseline established across all global respondents.

Governance as a multiplier

Data governance functions as a powerful multiplier across the enterprise, but its positive or negative impact depends entirely on an organization's maturity. Outdated governance approaches exacerbate the need for manual, downstream checkpoints—slowing AI deployments, driving up resource costs, and most importantly exposing the organization to compliance drift. By modernizing enterprise governance and including advanced AI solutions with automated ingestion-level information tagging, businesses lower risk while empowering teams to scale enterprise AI more safely and seamlessly.

This difference shows up in our research:

Percent of respondents that agree: Without stronger data governance, our organization cannot fully trust or operationalize AI outputs

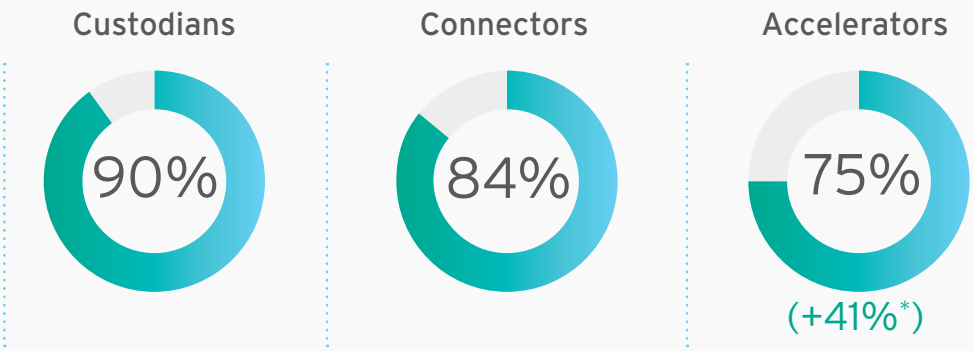

For lower-maturity **Custodians**, governance operates as a negative multiplier. With manual, downstream reviews and reactive policy checks, it acts as an operational brake that stalls AI initiatives, inflates labor costs, and opens organizations up to compliance problems.

Connectors occupy a middle ground, having digitized basic compliance workflows, yet they still struggle with disconnected enforcement policies that lead to localized bottlenecks and inconsistent risk controls.

For **Accelerators**, governance serves as a positive multiplier. By embedding automated classification, encryption, and zero-trust security directly at the point of ingestion, these leaders ensure downstream AI platforms inherit clean, secure, and fully auditable data—making compliance an intrinsic property of the information itself. Accelerators also elevate HITL as a strategic best practice: reserving expert oversight for high-risk edge cases, policy validation, and model refinement. This shift turns governance into the fail-safe boundary that lets specialized AI agents take autonomous actions safely. By embedding compliance, access controls, and PII protections directly at the data layer, enterprise teams gain the structural confidence to deploy, scale, and trust AI agents at speed—without exposing the company to regulatory or security risks.

> "Aligning governance, security and accessibility across unstructured data has been our most significant challenge over the past year."
>
> —Senior IT leader, Healthcare

* Reflects the difference versus the global baseline established across all global respondents.

Business impact: the accelerator advantage

Tangible business metrics offer perhaps the most vital validation point for building content intelligence maturity. The structural differences between **Custodians**, **Connectors**, and **Accelerators** yield clear, measurable business outcomes where executives care most: top-line revenue growth, faster decision velocity, and productivity gains that create a compounding market advantage.

In addition to outpacing their peers in their ability to scale AI projects and experience revenue growth, **Accelerators** are leading in several additional business metrics.

With faster decision-making, improved productivity, and the ability to point high-performing internal teams at more strategic work, organizations are further along the path to realizing their overall AI ambitions.

Percent of respondents that cite their strategies as delivering:

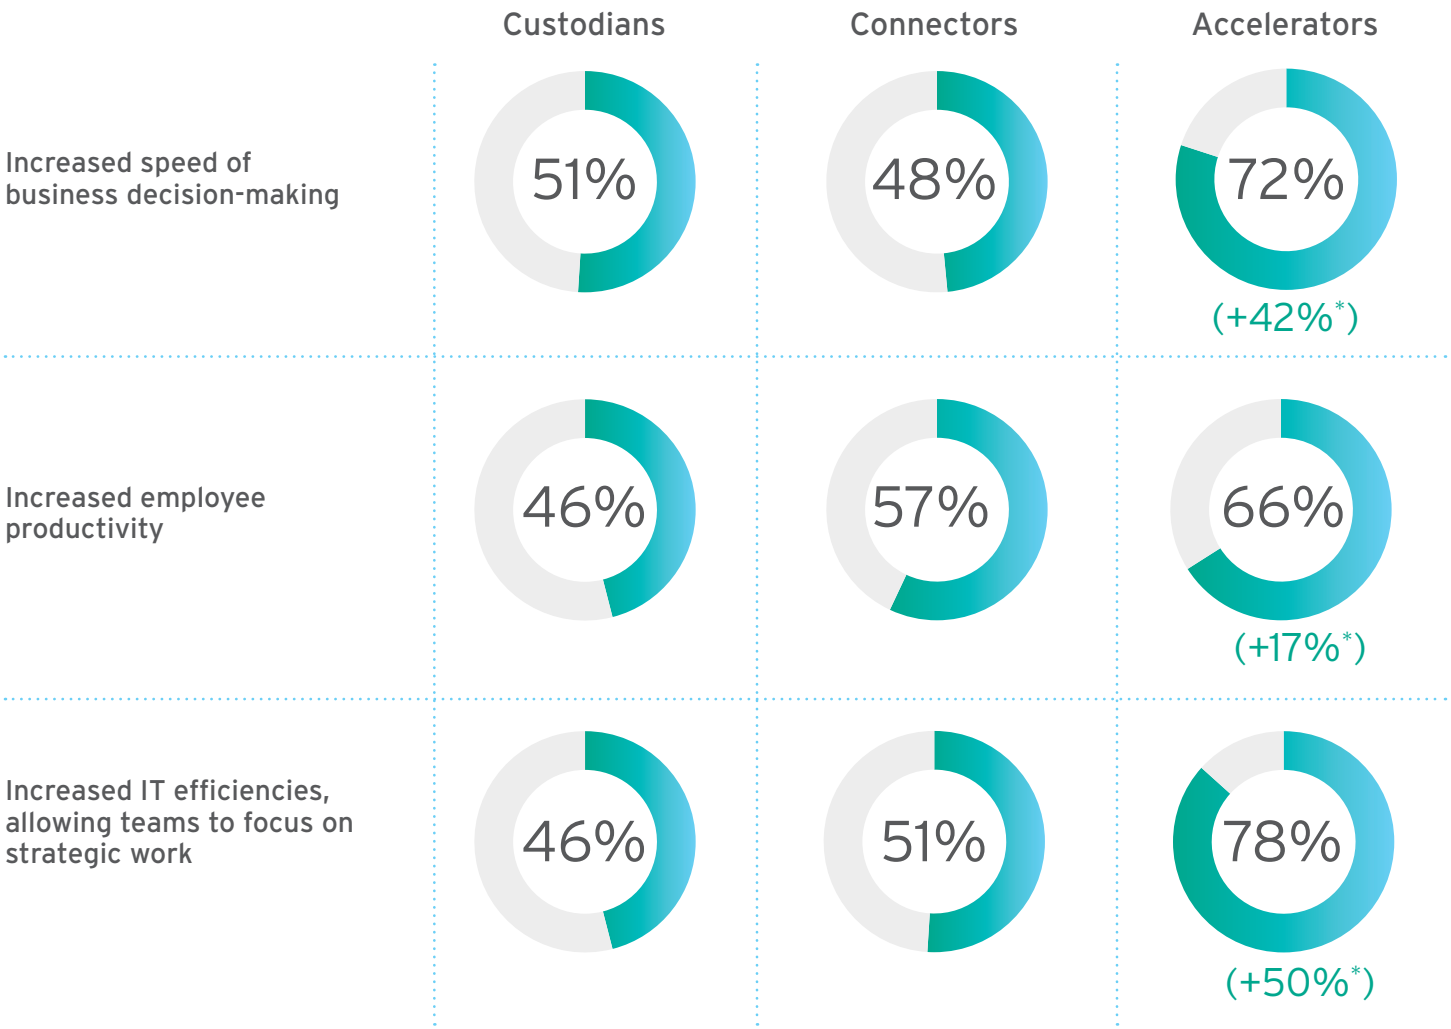

* Reflects the difference versus the global baseline established across all global respondents.

Conclusion

The enterprise AI race will not be won simply by purchasing more powerful models. The true differentiator is an organization's capability to unify, govern, and activate its unstructured content estate at scale.

Organizations that remain **Custodians** will continue to see their AI ambitions hindered by manual wrangling, security fears, high maintenance costs, and uncertain ROI. Conversely, **Accelerators** are proving that unifying content intelligence, automating governance, and reducing legacy debt creates a competitive advantage.

To break through the data ceiling and unlock the full promise of enterprise AI, technology leaders must first solve the unstructured content challenge.

Iron Mountain InSight DXP: The intelligent content platform built for enterprise trust

As our research reveals, closing the gap between dark data and AI-ready decision velocity requires more than basic automation—it demands an intelligent content platform rooted in enterprise trust. For regulated organizations, that trust comes from embedded security, policy-driven governance, and transparent AI decision-making. At its core, Iron Mountain InSight® DXP transforms passive content into active intelligence, giving organizations the power to deeply understand their dark data, support employees in day-to-day decisions, and give regulated enterprises the structural confidence to unleash specialized AI agents—accelerating decision velocity while maintaining rigorous risk and privacy controls.

InSight DXP accomplishes this through three core pillars:

1 Discover & trust (data readiness)

Converts digital and physical dark content into structured, high-fidelity datasets. By extracting context and metadata at the point of ingestion, InSight DXP helps mitigate agent blind spots and provides AI models with clean, accurate inputs.

2 Activate & govern (enterprise trust in action)

Applies zero-trust permissions, policy-driven safeguards, and persistent PII redaction directly at the data layer. By pairing these guardrails with transparent source citations and AI confidence scoring, InSight DXP verifies data accuracy and establishes safe boundaries for autonomous AI agents to execute workflows—reducing compliance and security risks.

3 Preserve & prove (agentic execution & ecosystem integration)

Delivers robust lineage tracking, legal holds, and auditability across processed content. InSight DXP enriches and governs dark data at the source, seamlessly feeding trusted, AI-ready data back into 300+ enterprise platforms to power downstream workflows across your existing ecosystem.

About the research

Iron Mountain commissioned independent market research specialist Vanson Bourne to conduct this study. The research included a survey of 400 IT decision-makers with knowledge of or responsibility for their organizations' content intelligence strategy. Respondents' organizations had at least 1,000 employees across the following countries: US, UK, Brazil, Germany, India, and Australia.

Participants represented a mix of public and private sectors, with strong focus in banking, financial services, and insurance; government and public sector (excluding healthcare); and healthcare and life sciences.

About Vanson Bourne

Vanson Bourne is an independent specialist in market research for the technology sector. Their reputation for robust and credible research-based analysis is founded upon rigorous research principles and their ability to seek the opinions of senior decision-makers across technical and business functions in all business sectors and all major markets. For more information, visit www.vansonbourne.com.

Ready to shatter your enterprise data ceiling?

Establish an active, governed content foundation that gives your organization the structural confidence to unleash specialized AI agents—accelerating decision velocity while maintaining rigorous risk and privacy controls.

Explore Iron Mountain InSight DXP

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 makes no guarantees to develop, deliver, or market any specific product, feature, functionality, or service. All future products, functionality, features, technical specifications, or service offerings, including their availability or release dates, are subject to change or withdrawal at any time without notice. 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. "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.