Why Enterprise Information Access Is Broken

And How Lenovo Closes the Enterprise Knowledge Gap

Summary

Enterprise organizations have never had more information at their disposal. They have also rarely had a harder time using it. Data lives in dozens of disconnected systems. Policies sit in SharePoint folders nobody updates. Institutional knowledge lives in the heads of employees who may leave next quarter. And as organizations push AI into core workflows, this fractured foundation is becoming a material liability.

This dynamic is not new, but the stakes are higher than ever. Enterprises have spent years and real money on search platforms, AI assistants, and custom-built retrieval systems, and most of those investments have underdelivered. Not because the technology was wrong, but because the underlying knowledge infrastructure was not ready for it. Data fragmentation, weak governance, and the absence of trust mechanisms have quietly undermined initiative after initiative. Now, as agentic AI moves from pilot to production, the knowledge layer that supports those systems is becoming genuinely critical infrastructure. More than that, the gap between organizations that have it right and those that do not is starting to show.

Underlying it all is a lack of trust. Enterprise knowledge systems fail when employees stop believing the answers. A system that surfaces results without showing its work, or returns content the user suspects they should not be seeing, loses credibility fast. Citation grounding, access controls enforced at the retrieval layer, and a structured feedback mechanism form the architecture of trust. Without them, adoption data looks fine — until it suddenly doesn't.

Enterprise knowledge management is harder than most organizations give it credit for, and it is important to understand why. It is also useful to see how a solution tailored to address the problem like the Lenovo Knowledge Super Agent can provide a strong knowledge- system foundation.

MARKET OVERVIEW: THE CHALLENGE IS BIGGER THAN WE REALIZE

Most organizations know they have a knowledge problem. But Moor Insights & Strategy (MI&S) has observed that few understand just how deep this problem runs. The surface symptoms are familiar: Employees can't find what they need, experts get pulled into questions they shouldn't have to answer, and new hires take too long to ramp. But the foundational cracks are more consequential than those symptoms suggest — like the tip of an iceberg. And it's getting harder to ignore these challenges, as AI takes on a larger operational role.

MI&S has identified several common dimensions to the enterprise challenge.

THE BUSINESS CHALLENGE OF THE ENTERPRISE 'INFORMATION ARCHIPELAGO'

Most organizations frame this as a search problem. It's not.

Over the past fifteen years, organizations have layered tool on top of tool — SharePoint, Confluence, Slack, ticketing systems, CRM platforms — and each addition made sense at the time. But together, they've created an information archipelago with no map.

The search problem is merely a symptom. The underlying condition is that enterprises have built knowledge systems optimized for storage and access control, not for retrieval, synthesis, or trust. Employees don't need a better search bar. They need a system that can tell them what the answer is, show where it came from, and be trusted to get it right. The operational cost of not having a trusted knowledge system is concrete, and it can easily run into millions of dollars per year in lost productivity for enterprises. Subject matter experts become informal help desks. New employees spend months learning who to ask before they can operate independently. And in regulated industries, acting based on a superseded SOP can move quickly from an efficiency problem to a compliance liability.

IT CAN'T KEEP PACE

IT organizations have invested heavily to address this problem. Enterprise search improved indexing and discoverability. However, it didn't solve the core issue: Employees need answers, not ranked lists of documents that *might* contain an answer. The cognitive work of synthesizing information across results still falls entirely on the person asking the question.

Custom RAG pipelines are the current state of the art for handling this problem. And in choosing RAG, the instinct is right: Ground an LLM in internal data, enforce access controls at the retrieval layer, and generate cited answers. Building this system well in an enterprise context requires ML engineers, data engineers, platform engineers, and careful governance design, with everyone working together for months. Most organizations do not have that bench. Those that do face ongoing maintenance overhead, and without a structured feedback loop, quality degradation is silent and sure. Nobody files a ticket when they stop trusting a tool. They just stop using it.

Research from Boston Consulting Group puts the challenge in sharp relief: Only 26% of companies have developed the capabilities to move AI investments beyond proof of concept. The rest have yet to show tangible value. Roughly 70% of implementation failures stem from people and processes, not the algorithms. And data readiness ranks as a problem above all else.

The Economic Status Quo

The economic cost of knowledge fragmentation is not a new discovery. As far back as 2012, McKinsey Global Institute research found that knowledge workers spent roughly 20% of the workweek searching for and gathering information. This amounts to one full workday per employee, per week. And the problem has only grown more complex since then, as SaaS sprawl multiplied the number of systems employees must navigate. Worse, the trend toward hybrid work has removed the informal channels that partially compensated for poor knowledge infrastructure.

Signal65's evaluation of the Lenovo Knowledge Super Agent found a 30% reduction in knowledge retrieval time across a 3,000-user deployment. Signal65 testing applied this to a workforce with an average salary of $100,000. This translates to roughly 120 hours saved per employee, per year. This amounts to approximately 360,000 hours and up to $17 million in potential productivity value. While these are order-of-magnitude figures, not guarantees, they illustrate why the status quo carries a real price tag. Even when nobody is tracking it.

The Knowledge Gap is Widening

As touched on above, broad SaaS adoption (and sprawl) adds new silos faster than organizations can retire old ones. And the hallway and desk-side conversations that partially helped get around this issue are not coming back. To make matters worse, AI-

generated content is now producing documentation at high volume without clear ownership or version control, creating a new category of knowledge debt.

The arrival of agentic AI only raises the stakes. AI systems that retrieve information and take action on their own are only as good as the knowledge layer beneath them. A weak foundation does not just slow things down — it produces wrong answers at machine speed.

Trust: Why This Is Harder Than It Looks

Productivity statistics don't capture the entire issue. Employees don't file tickets when they stop trusting a tool. They develop workarounds by asking colleagues instead, or by quietly reverting to manual processes. By the time IT sees adoption declining, the credibility damage is already done. One confidently wrong answer — let alone in a compliance-sensitive context — can trigger that loss faster than months of positive interactions can recover it.

This is the actual failure mode behind most of the enterprise AI pilots that have stalled over the past two years. Not model performance, not integration complexity, but the organizational judgment that the system cannot be trusted with real decisions.

By contrast, designing for trust from day one requires citation grounding, access controls that hold at the retrieval layer, and a feedback mechanism that closes the loop between what employees experience and what administrators can actually fix.

How Enterprise 'Fixes' Fall Short

Organizations have not ignored this challenge. Over the years, they've tried to solve for it in three distinct ways: investing in enterprise search, deploying general-purpose AI assistants, and building custom retrieval pipelines. Each approach addressed part of the problem. None closed the loop.

The Enterprise Search Investment

Better indexing, better ranking, more connectors. The investment has been real, but the returns have been inconsistent. As mentioned earlier, the fundamental limitation is architectural: Search returns documents, not answers. The synthesis still falls on the employee. Repeated cycles of adoption and quiet abandonment have left many of the organizations MI&S engages with skeptical of search-centric approaches.

General-Purpose AI Assistants

AI assistants are strong for drafting, summarizing, and reasoning across content you provide. But they're limited for questions that require knowledge of internal policy, proprietary process, or organizational context without significant additional configuration.

Major enterprise software vendors and hyperscalers are actively building retrieval and workflow orchestration into their AI strategies. However, the current gaps for regulated environments are specific: deep customization to a particular data environment, governance enforced at the retrieval layer, and on-premises deployment where data residency requirements apply.

While these are not permanent limitations, they represent where purpose-built platforms currently hold an advantage.

Custom-Built RAG Pipelines

RAG pipelines are architecturally sound, but operationally difficult. Production-grade retrieval systems require specialized engineering across multiple disciplines, and the governance work (RBAC at the retrieval layer, data quality management, feedback and monitoring) is consistently underestimated.

The common thread across all three approaches is this: None consistently close the gap between the information employees need and the systems that hold it in a way that is trusted, governed, and scalable.

Lenovo xIQ AI Agent Platform Is Built For This Opportunity

When most organizations evaluate enterprise knowledge platforms, they focus on employee productivity. This is the right starting point, but the wrong ending point.

The more consequential shift is architectural: As enterprises move toward agentic workflows, the knowledge layer stops being a productivity tool and becomes a control plane. AI agents that retrieve information, reason across it, and take action downstream are only as reliable as the data layer they operate against. The quality, governance, and accessibility of that layer — and not the model, orchestration framework, or interface — determine what the agents can actually do. With that in mind, getting the knowledge substrate right is now a precondition for AI that performs at enterprise scale, not a follow-on improvement.

The Lenovo Knowledge Super Agent is built on the Lenovo xIQ Agent Platform, a no-code environment designed for enterprise AI agent development, configuration, and management. The xIQ Platform handles the underlying infrastructure — model orchestration, retrieval pipeline management, governance controls, and monitoring — so organizations can deploy and customize knowledge agents without sustained involvement from ML engineers.

The Knowledge Super Agent is one of several domain-specific agents available through the Lenovo AI Library, a curated catalog of prebuilt agents built on validated enterprise AI patterns. Agents can be deployed independently or used together as organizational AI needs expand.

At the retrieval layer, the Knowledge Super Agent connects to enterprise data sources, including SharePoint, Confluence, and S3 object storage. It indexes content into a vector database that powers semantic search and retrieval-augmented generation. When an employee asks a question, the system retrieves relevant content from the indexed knowledge base, synthesizes an answer, and surfaces supporting citations so the response can be verified.

Role-based access controls and source-level permissions are enforced at this retrieval layer (not at the interface), which means an employee will not receive answers derived from content they're not authorized to see. This distinction matters more than it might appear. Most homegrown implementations apply access controls at the surface, which creates edge cases that erode trust when they surface in production.

What the Platform Delivers

Deployment flexibility is foundational. The platform supports on-premises, cloud, and hybrid configurations. For regulated industries, that is frequently not a preference — it is a requirement that eliminates cloud-only alternatives. The xIQ Agent Platform provides no-code configuration and management, reducing the engineering dependency that makes custom builds expensive and slow. Initial deployments can be operational in weeks rather than the months a custom build requires.

What the Numbers Tell Us

Signal65 evaluated the Lenovo platform against a live enterprise deployment of approximately 3,000 users, covering backend setup, feature capabilities, and real user outcomes. The results warrant examination in some detail.

- In this study, 81% of employees using the Lenovo system reported reduced time spent searching for information, and measured retrieval time dropped by 30% across the deployment.
- Expert escalations fell by 35%, meaning the system was delivering answers that employees trusted enough to act on without checking with a human.
- 94% said it helped them begin tasks they previously struggled to start —– a practical measure of how much organizational ambiguity the platform was absorbing.
- Sales teams prepared proposals twice as fast.
- The platform scaled to 3,000 users without architectural changes and held a 4.4 out of 5 satisfaction rating across that base.
- 85% or more of responses included supporting source citations.

That last figure matters as much as any of the productivity numbers. When employees can see the source behind an answer, they can evaluate it. Citation grounding is what makes trust rational rather than assumed. And it is what drives sustained adoption, instead of a usage spike that fades after the first month.

Is This Differentiated?

In short, yes. A few things stand out when looking at how the Lenovo xIQ Agent Platform is positioned relative to the alternatives. The Signal65 evaluation was conducted against a live production deployment of 3,000 users, not a controlled demonstration environment. That matters because production deployments expose failure modes that controlled tests do not.

On-premises deployment is native to the architecture, not retrofitted. For organizations in regulated industries, that is the difference between a viable option and one that looks good in a briefing but can't clear a compliance or security review.

The feedback loop is worth calling out specifically. Most enterprise knowledge deployments lack a structured mechanism to catch quality degradation before it becomes an adoption problem. But the Lenovo system has that structured mechanism. Users rate responses, negative feedback gets triaged, and issues are routed to the teams that can fix them. It sounds straightforward. But in practice, it's what keeps a deployment useful in year two, when the novelty is gone and the system has to keep earning its place on its own merits.

CONSIDERATIONS FOR IT LEADERS

Over the course of ongoing conversations with enterprise IT and business leaders, a few consistent themes emerge around where knowledge deployments run into trouble. The challenges are rarely technical. They're organizational — governance decisions deferred, data quality assumptions that don't hold, and pilots scoped for optics rather than real-world stress.

Based on that experience, and on what the Signal65 evaluation reveals about what drives adoption and what kills it, MI&S suggests the following considerations for IT leaders evaluating this space.

Map Your Knowledge Topology First. Most enterprises have a reasonable picture of their formal systems and a murky picture of what actually gets used, who owns it, and where tribal knowledge gaps are. Inventory authoritative content sources and ownership before configuring anything. AI cannot navigate what has not been mapped.

Treat Governance as a Prerequisite. RBAC at the retrieval layer must be right before the first user query. Retrofitting access controls after deployment erodes trust immediately. Employees who see content they should not have access to — or suspect others can see theirs — lose confidence in the system. Decide the authorization model before configuration begins.

Pilot on a Hard Problem. General Q&A hides weaknesses. Compliance workflows, sales enablement, and onboarding are better tests: Answers need to be specific, current, and verifiable. A pilot that passes those tests is demonstrating something real.

Establish Baselines Before Launch. Citation rate, escalation volume, and user satisfaction should be tracked from week one. Without baselines, quality degradation is invisible until adoption has already collapsed. Define what good looks like before you have data to measure against.

Evaluate Deployment Model Fit Early. For organizations in regulated industries, cloud-only processing of sensitive internal data is often not viable. Treat deployment model fit as a first-order requirement, not a late-stage procurement question.

Address Data Quality Before You Deploy. This is where most implementations run into trouble. Deploying a knowledge agent on top of poorly maintained documentation does not solve the knowledge problem — it scales it. Duplicate repositories, stale

documents, conflicting SOPs, and unclear content ownership are organizational problems that AI will surface and exacerbate, not resolve.

There is a related problem that organizations rarely anticipate. RAG-based systems degrade over time in ways that are not immediately visible. Document indexes grow stale as source content changes. Embeddings generated against an older version of a knowledge base do not automatically update when the underlying documents do. Permission structures that were clean at launch drift as organizational roles change and source systems evolve. Conflicting “authoritative” sources accumulate — two documents making different claims about the same policy, both indexed, both surfaced with equal confidence. The system doesn’t know which one is current. It returns both, or worse, synthesizes across them.

This isn’t a deployment problem. It’s a lifecycle problem that requires ongoing governance investment, not just a pre-launch cleanup. The organizations that sustain knowledge system performance treat index hygiene, source synchronization, and content ownership as operational disciplines with assigned accountability, not IT maintenance tasks that happen when someone remembers.

Audit the highest-priority knowledge domains before deployment. Establish a content governance model that sustains quality over time, not just during the initial rollout.

Call to Action

Enterprise knowledge fragmentation is not a new problem. What is new is the context. AI systems that can reason, retrieve, and act are moving from experimentation into production. And organizations that build those systems on top of a governed, well-structured knowledge layer will get materially different outcomes than those that don’t.

Most enterprises can’t yet generate tangible value from AI at scale. A significant share of that failure can be traced to the knowledge foundation — data quality, governance design, and access control. These aren’t glamorous problems. But they determine whether AI investments pay off or accumulate as a portfolio of expensive pilots that never shipped.

The Lenovo Knowledge Super Agent addresses the failure modes that matter most in practice. Governance is enforced at the retrieval layer, not bolted on afterward. Deployment flexibility is real, not a cloud product with an on-premises workaround. And the feedback loop is structured to catch quality erosion before users lose confidence.

The enterprises that are pulling ahead on AI aren't those with the most sophisticated models. They're the ones that have built a clean, governed, well-instrumented knowledge layer that agents can reliably act on. That layer is becoming the operational substrate for everything downstream — from workflow automation to decision support, and customer-facing intelligence. Organizations that treat it as foundational infrastructure, govern it properly, and measure it from day one will be better positioned to realize strong returns on the AI investments they have already made.

The risk of not doing so is worth stating plainly. A general-purpose AI assistant that hallucinates produces a wrong answer. An agentic AI system operating against a poorly governed knowledge layer can operationalize that hallucination — routing it into a workflow, a customer response, or a compliance decision before anyone catches it. The stakes are different at that level, and the knowledge foundation is where exposure is either managed or ignored.

For more information on the Lenovo xIQ Agent Platform and Lenovo Knowledge Super Agent, please visit this link.

Important Information About This Paper

Author

Matt Kimball, Vice President and Principal Analyst, Datacenter Compute and Storage

Publisher

Patrick Moorhead, CEO, Founder and Chief Analyst at Moor Insights & Strategy

Inquiries

Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

Citations

This paper can be cited by accredited press and analysts but must be cited in-context, displaying the author's name, author's title, and "Moor Insights & Strategy." Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

Licensing

Disclosures

Lenovo commissioned this paper. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

Disclaimer

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.