



**KM**World

# 2026 State of KM & AI Report

# The State of KM & AI: Results of a KMWorld Survey

By Marydee Ojala, Editor-in-Chief, *KMWorld*

**R**apid adoption of generative AI (GenAI), combined with long-standing AI technologies such as machine learning, semantic search, knowledge graphs, and semantic layers, has had an outsized impact on the practice of KM. At the same time, there continues to be skepticism, as not all AI-driven projects have delivered on their promised value. From a KM perspective, however, the influence of AI on content creation, curation, organization, access, and dissemination is undeniable. AI aligns strongly with user expectations for speed and agility in knowledge sharing. Still, the human element remains critical: The expertise of KM practitioners is invaluable when organizations evaluate and implement AI-based products at an enterprise scale.

KMWorld found the relationship between KM and AI compelling enough to warrant a survey of KM practitioners. The insights gathered from 202 respondents covered a wide range of issues surrounding the intersection of KM and AI. Respondents were predominantly based in North America and represented a broad number of industry sectors. Organizations of varying sizes were represented, with 55% having less than 1,000 employees, and 45% having more than 1,000 employees. The largest organizations (more than 10,000 employees) made up 19% of respondents.

Roles were distributed relatively evenly across KM practitioners (32%), IT and technology professionals (29%), executive leadership (22%), directors and senior managers (21%), and operational managers (29%). This diversity provides a well-rounded view of both strategic and operational KM priorities.



## THE CURRENT KM AND AI LANDSCAPE

Looking at the current KM landscape, nearly half of organizations (49%) employ a hybrid model that combines centralized and decentralized elements, while only 15% use a highly centralized approach. It is notable—and somewhat concerning—that 23% of respondents report that no formal KM structure exists in their organization.

The survey results paint a clear picture of today's Knowledge Management (KM) environment: a rich, yet often overlapping and complex, array of tools and systems. Organizations are leveraging a broad spectrum of technologies, from long-standing staples to more advanced solutions, to manage their invaluable knowledge assets. This diversity reflects the multifaceted nature of knowledge itself, demanding various approaches for creation, storage, sharing, and discovery.

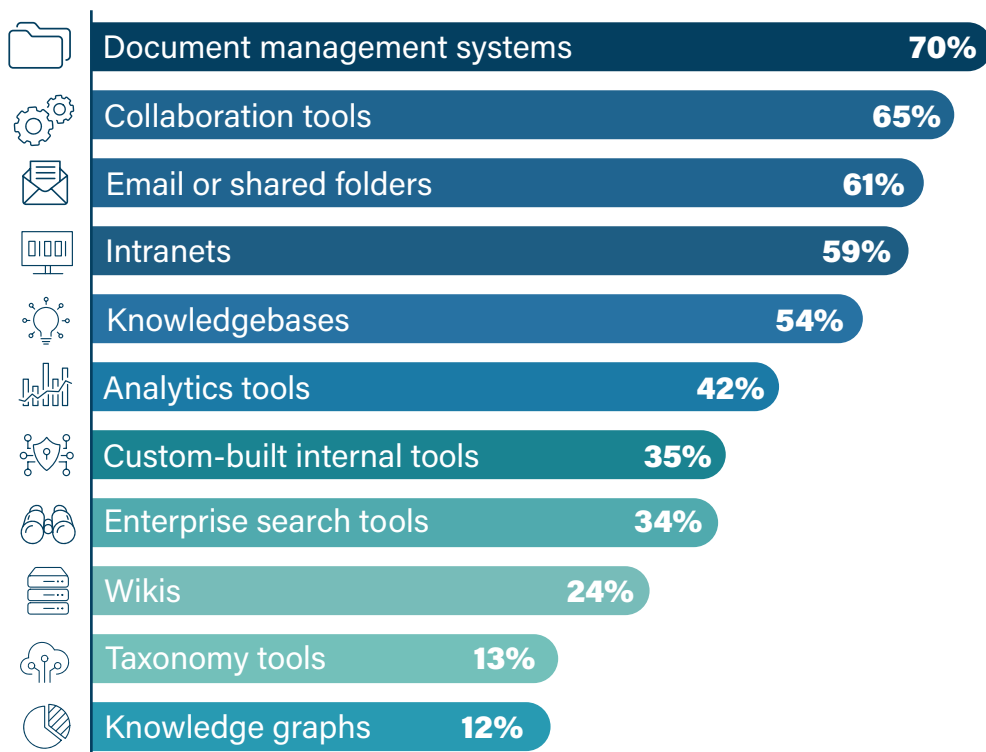
However, this extensive toolbox also presents its own challenges, leading to fragmentation and potential information silos if not strategically managed.

### Key Insights From the Survey:

- **Traditional Pillars Remain Strong:** Document Management Systems and Intranets continue to form the backbone of KM infrastructure in many organizations, proving their enduring utility.
- **Collaboration Is King:** Collaboration tools that facilitate real-time communication and joint work are widely adopted, underscoring the importance of social learning and shared workspaces.
- **Specialized Tools Emerge:** While still less prevalent, dedicated knowledgebases, enterprise search solutions, and even advanced tools like knowledge graphs are finding their place,

## The KM Toolbox: A Diverse Infrastructure

Organizations utilize a broad spectrum of tools and systems for knowledge management, reflecting a varied landscape



indicating a move toward more sophisticated knowledge organization and retrieval.

- **Email & Shared Drives Persist:** The continued high usage of email and shared folders highlights both their fundamental role and the ongoing need to integrate and formalize knowledge often trapped in these less structured environments.

This diverse landscape means that no single tool dominates; instead, organizations are building unique ecosystems tailored to their specific needs. The challenge moving forward will be to integrate these disparate systems into a cohesive and easily navigable knowledge environment.

Knowledge sharing itself appears active, with 61% of organizations sharing knowledge daily or multiple times per day, suggesting a generally strong knowledge-sharing culture.

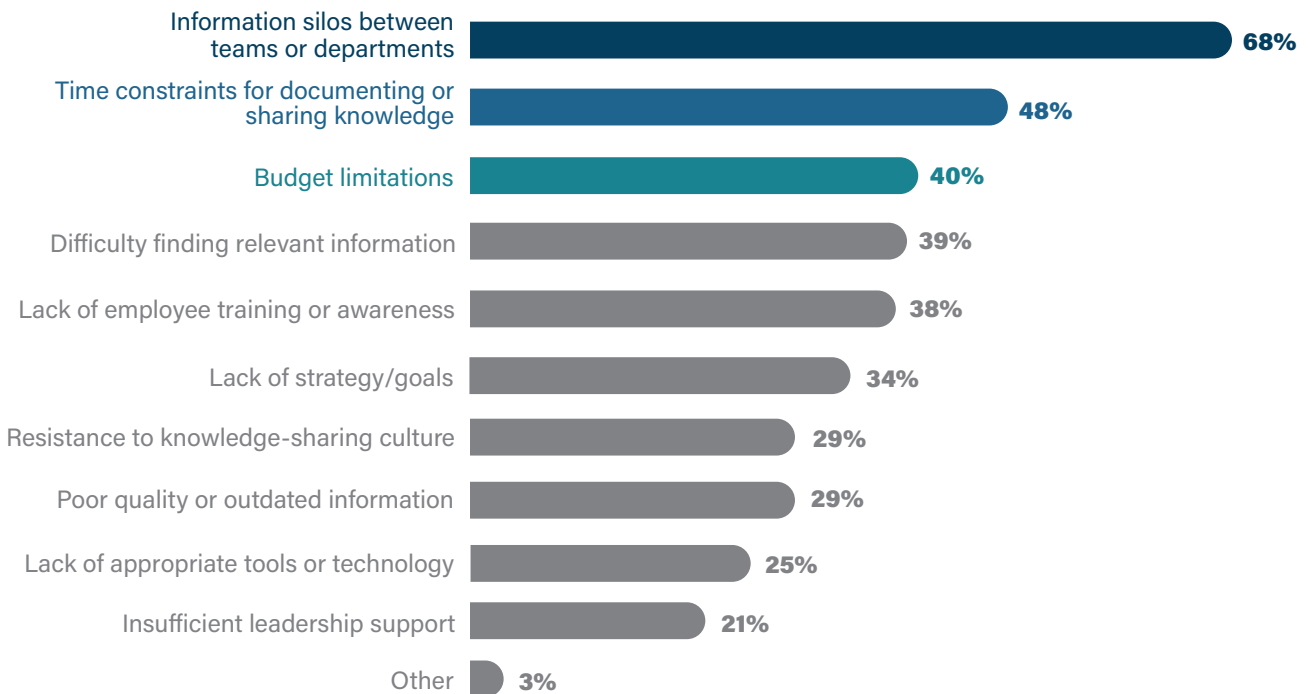
## KNOWLEDGE MANAGEMENT CHALLENGES

KM is not without significant challenges. The most commonly cited obstacles are information silos, time constraints for documenting knowledge, and lack of a clear KM strategy or goals. Difficulty finding relevant information and budget limitations also rank highly. Together, these findings show that organizations recognize the importance of KM but continue to struggle with both infrastructure and strategic alignment.

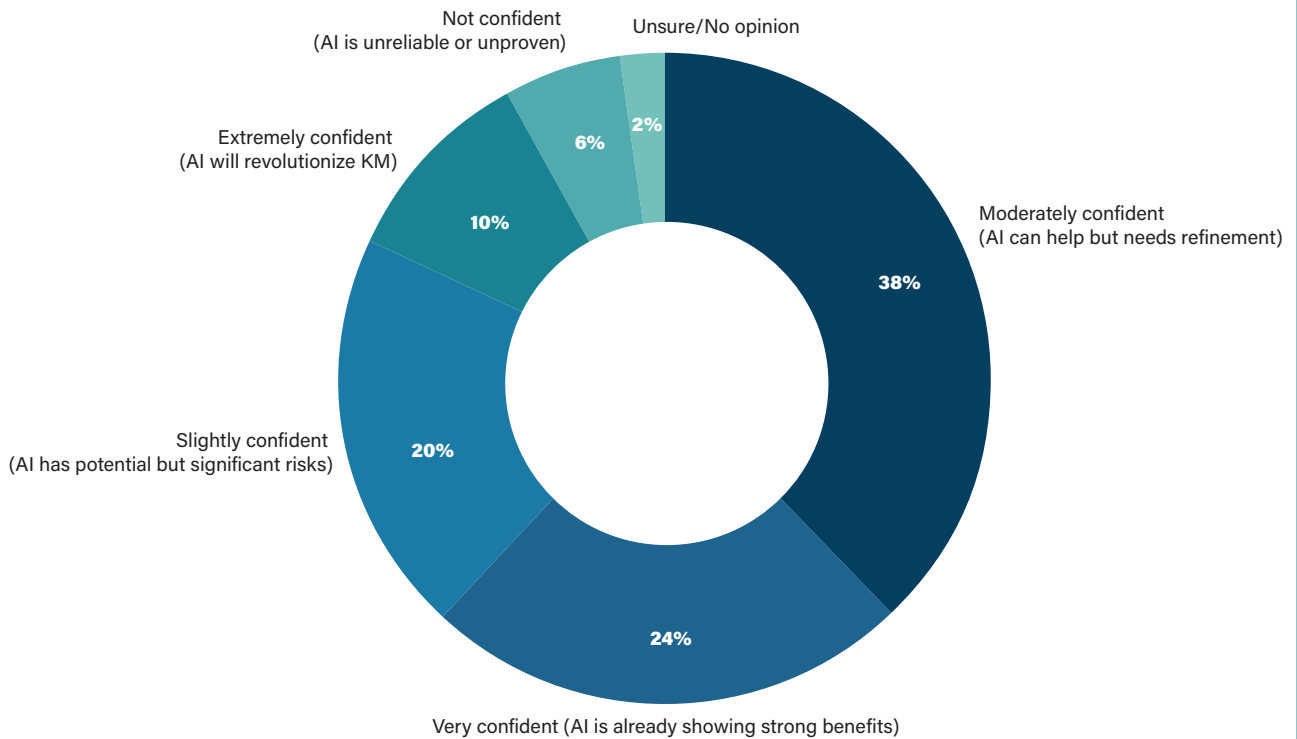
Effectiveness ratings reveal notable gaps. Only 25% of organizations rate their KM processes as “mostly effective” or better, while 32% describe them as “somewhat ineffective” or worse. Enterprise search performs similarly poorly: Just 7% rate it as “very effective,” and 53% rate it as “somewhat effective” or below.

## The Silo Problem Dominates KM

Organizations cite “information silos” as the number-one barrier to managing knowledge effectively, with time constraints and lack of budget next in line.



# Confidence in AI's Ability to Enhance Knowledge Management



## AI AND KM

Most organizations are still in the early stages of AI adoption for KM. About 37% use AI minimally through pilots or early testing, and 32% use it moderately for specific tasks. Only 5% report extensive, integral use of AI. Another 16% plan to implement AI soon, while just 8% have no plans, reflecting generally positive but cautious sentiment.

Where AI is in use, the most common applications are AI-powered chatbots (49%) and GenAI for content creation (48%). Intelligent search (39%) and natural language processing (38%) follow closely. More advanced use cases, such as predictive analytics, remain limited at 14%. Most respondents are relatively new to AI in KM, with 33% having used it for 6–12 months and 15% for 3–6 months, indicating that many organizations are still in learning and evaluation phases.

Confidence in AI's potential is high, with 62% expressing moderate to extreme confidence in its ability to enhance KM. Employee comfort is lower, with only 37% reporting moderate comfort and 27% expressing reluctance. This gap highlights the importance of change management and training.

Respondents see the highest-impact AI use cases as intelligent search and retrieval (70%), automated summarization (68%), knowledge discovery (60%), and document classification and organization (59). These priorities closely align with current adoption patterns, suggesting organizations are focusing on practical, high-value applications first.

While enthusiasm for AI is high, the survey identifies a "reality check" of significant hurdles that organizations must clear before achieving full-scale implementation. These barriers suggest that the transition to AI-driven KM is as much

# AI Adoption in KM: Navigating the Hurdles

## Top 5 Challenges Encountered by Organizations

51%

Concerns about data security or compliance



40%

Lack of clear use cases or strategy



35%

Compatibility issues with existing systems



33%

Insufficient technical expertise or IT support



32%

Lack of budget or funding



about **governance and strategy** as it is about technology.

### Key Implementation Roadblocks:

- **The Trust Gap (51%):** Data security and compliance remain the primary concerns. Organizations are wary of feeding proprietary knowledge into AI models without ironclad privacy guarantees.
- **The Strategy Vacuum (40%):** Many teams are eager to use AI but lack a clear roadmap or defined use cases, leading to “pilot purgatory” where projects never move past the testing phase.
- **Technical & Expertise Hurdles (36%):** Compatibility with legacy systems and a shortage

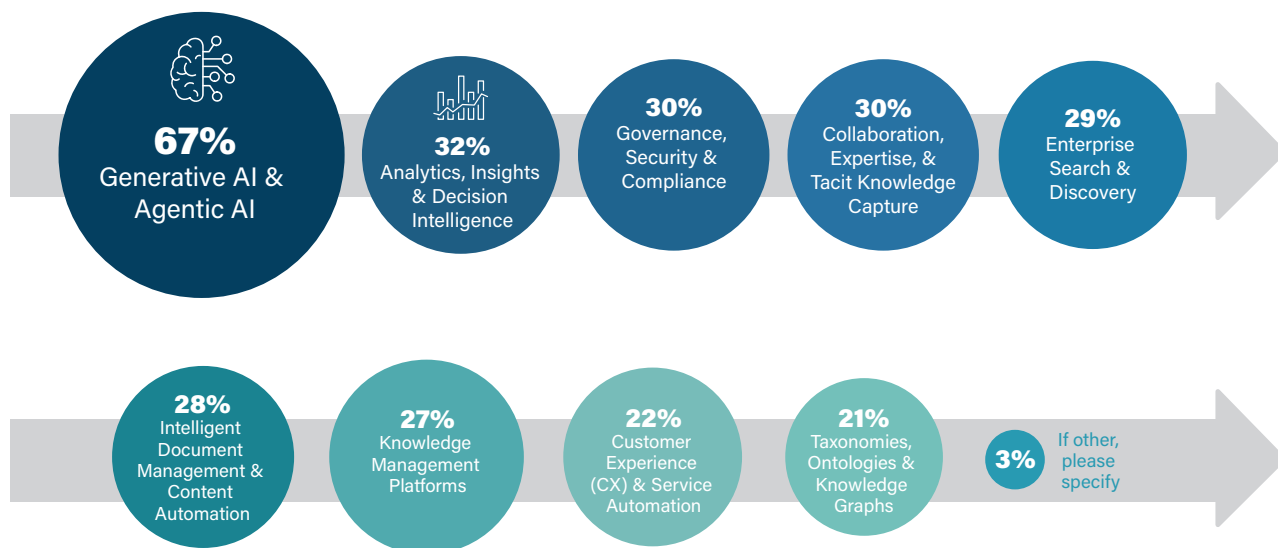
of in-house technical talent create a high bar for entry, especially for mid-sized firms.

- **Budgetary Friction (32%):** Despite the perceived value, securing dedicated funding remains a challenge, forcing KM teams to prove ROI earlier than in previous tech cycles.

Results from AI adoption are mixed. Only 30% report moderate to transformative improvements in knowledge retrieval efficiency, while 35% see minimal or no improvement. Success is most often measured by time saved (46%), improved accuracy (43%), and productivity gains (33%), yet nearly one-quarter (23%) do not measure AI success at all.

# Planned AI and Tech Adoption: The Next 12 Months

AI is now the top technology priority in 2026.



## ORGANIZATIONAL PRIORITIES FOR 2026

Top KM priorities for the coming year include enabling AI and automation initiatives (64%), improving access and discovery (49%), and boosting employee productivity (40%). These results clearly show that AI is central to KM modernization efforts.

Looking forward, 67% of organizations are researching or planning adoption of generative and agentic AI, making it the top technology priority. Other important areas include analytics and insights (32%), governance/security/compliance

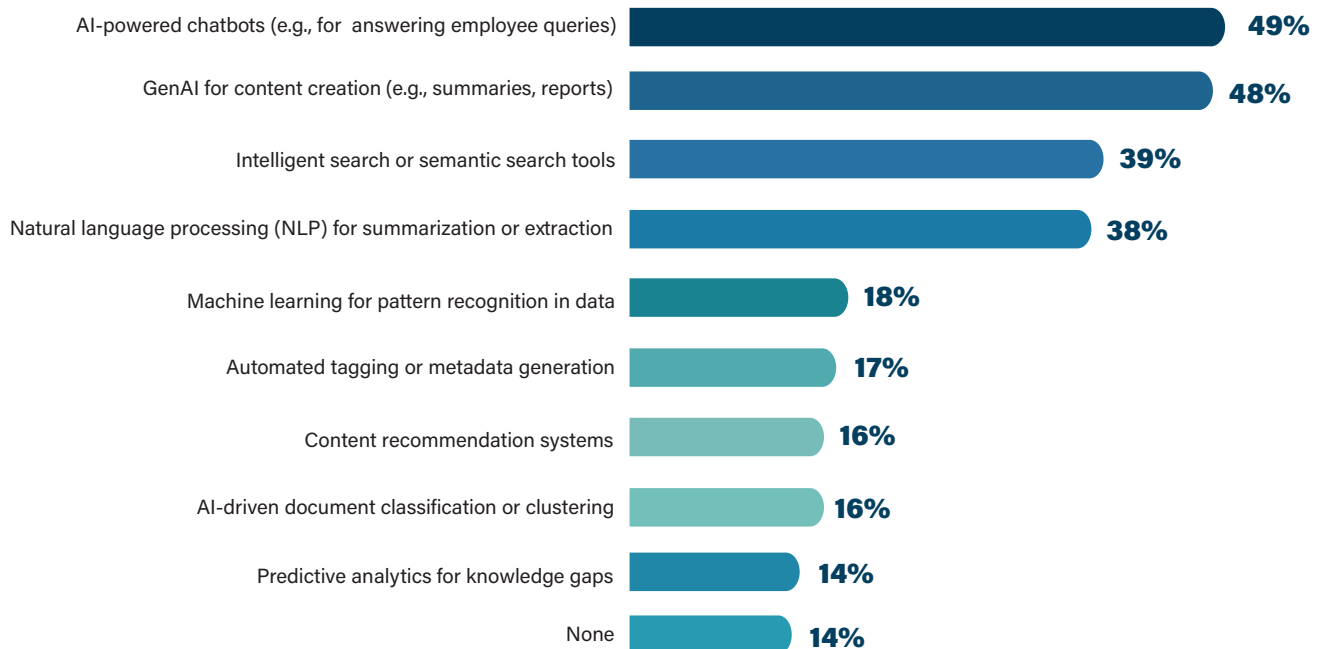
(30%), and enterprise search and discovery (29%).

This reflects a balanced strategy that combines innovation with foundational improvements and risk management.

Organizations want AI tools that deliver higher accuracy and reliability (62%), integrate better with existing systems (55%), provide stronger privacy and security (49%), and handle unstructured data more effectively (41%). These expectations point to demand for more mature, enterprise-grade solutions.

# AI Tools Currently Being Used in Knowledge Management

## Percentage of organizations using specific AI-driven technologies for KM Processes



### AI Technologies in KM Adoption

The chart above reveals a clear hierarchy in how AI is being utilized today:

- **The Adoption Leaders:** AI-powered Chatbots (49%) and GenAI for content creation (48%) are the clear favorites, nearly reaching a majority in usage. These tools address immediate needs for interactive support and high-speed content drafting.
- **The Search & Analysis Tier:** Intelligent or Semantic search tools (39%) and Natural Language Processing (38%) are foundational technologies that help employees cut through information overload to find meaning and summaries.

- **Specialized Emerging Tools:** A cluster of tools ranging from 14% to 18% usage shows the emergence of more structured AI applications. This includes machine learning for pattern recognition (18%), automated tagging (17%), content recommendation (16%), document classification (15%), and predictive analytics (14%).

This data suggests that while the industry is heavily invested in “chat-based” and “generative” interfaces, there is a massive opportunity for growth in the more technical, structural aspects of KM, such as automated metadata and predictive gap analysis.

Expected benefits of AI in KM include improved customer satisfaction (19%), reduced operational

costs (18%), and faster response times (14%), although 16% of respondents still assessing impact.

## EMERGING TECHNOLOGIES AND FUTURE OUTLOOK

Open-ended responses show overwhelming focus on AI in all forms: GenAI, agentic AI, and large language models. Other frequently mentioned technologies include knowledge graphs, advanced analytics, and semantic search. A strong theme is the desire for AI that is trustworthy, secure, and genuinely useful, rather than “AI for AI’s sake.”

Overall, the survey reveals deep interest in AI across the KM community, but also a clear readiness gap. Organizations see AI as a strategic imperative, yet

remain cautious to avoid costly missteps. Concerns around quality, trust, security, hallucinations, and data integrity continue to slow broad adoption.

While enthusiasm is high, many organizations are still constrained by foundational KM issues such as governance, integration, and content quality. These weaknesses limit AI’s effectiveness. To fully realize AI’s potential, organizations need better tools, clearer strategies, and stronger change management.

The future of KM appears to be a blend of AI-powered discovery and automation, integrated systems, strong governance, and analytics-driven insight—a major evolution from today’s largely manual, siloed practices. ■

# KM for AI: From Human to Agentic AI Enablement

By Fabrice Lacroix, CEO of Fluid Topics

Artificial intelligence isn't just a technology; it's the new audience and purpose of Knowledge Management. For the past two years, conversations have centered on how AI might alter KM through smarter content creation, automated tagging, and classification. Useful, yes, but incremental. Reshaping KM with AI is inevitable. The real issue is reshaping KM for AI, a shift that will transform KM for a new mission: Agentic AI enablement.

## KM for AI—The New Paradigm

KM has been designed around human needs. Its systems conventionally organize, store, and display information for people via intranets, document libraries, portals, and search engines. Agentic AI is rewriting the rules. As technology matures at unprecedented speed, the primary consumers of enterprise knowledge are shifting from humans to robots and agents.

Agentic AI refers to AI systems composed of agents that can behave and interact autonomously to achieve their objectives.

At a basic level, AI uses corporate knowledge to answer employees' questions and deliver insights. More advanced agents orchestrate processes and complete complex tasks independently.

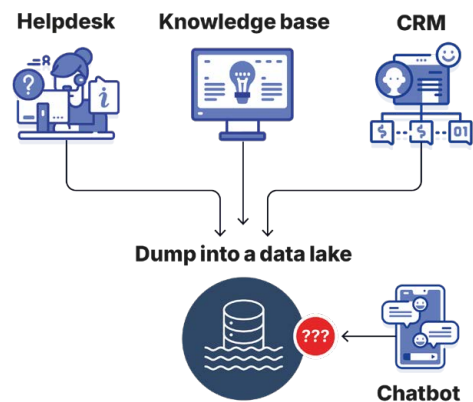
This isn't tomorrow. It's happening now: line of business (LOB) applications are already offering embedded assistants and autonomous workflows: CRM solutions deploy AI SDRs (Sales Development Representatives) that handle full prospection cycles, help desk tools rely on AI agents to resolve cases, and operations platforms have built-in AI copilots.

In this new division of labor where AI filters, analyzes, and contextualizes information for humans, KM will need to pivot from empowering people to powering agents. As a matter of fact, KM must reposition itself as Agentic AI Enablement.

## Implications for KM Information Architecture

The pre-requisite for agent performance is connecting to all available information, including siloed data and content scattered across the enterprise knowledge ecosystem.

Unified search has long been the KM solution to bridging silos: put everything into a data lake, then query across it. Early "universal" chatbots replicated this approach and failed. Outside its original context and system, undifferentiated content lacks meaning, making it largely unusable for AI. Additionally, access rules are hard to transfer, updates are asynchronous, and interactions are limited to read-only experiences.



In 2025, AI evolved, opening a new way. LLMs grew into reasoning-capable LRMs, and MCP enabled agents to interact directly with business applications. Rather than using irrelevant data lakes, an agent now reasons towards its goal, retrieves what it needs directly from each source, and when appropriate, acts directly on those systems.

Large Language Models (LLMs) are a category of AI trained on vast amounts of data, enabling them to understand and generate natural language as well as other types of content.

Large Reasoning Model (LRM) is an AI system that combines natural language understanding with logical reasoning to solve complex problems.

The Model Context Protocol (MCP) is a protocol that standardizes how LRMs can query existing applications (both read and write).

## The Need for Domain Knowledge

This technology is powerful, but it needs clear guidance in practice. Just as humans turn to subject matter experts for specific inquiries, agents need a canonical reference for each business domain. AI must be grounded in domain knowledge, and the underlying information architecture must reflect this.

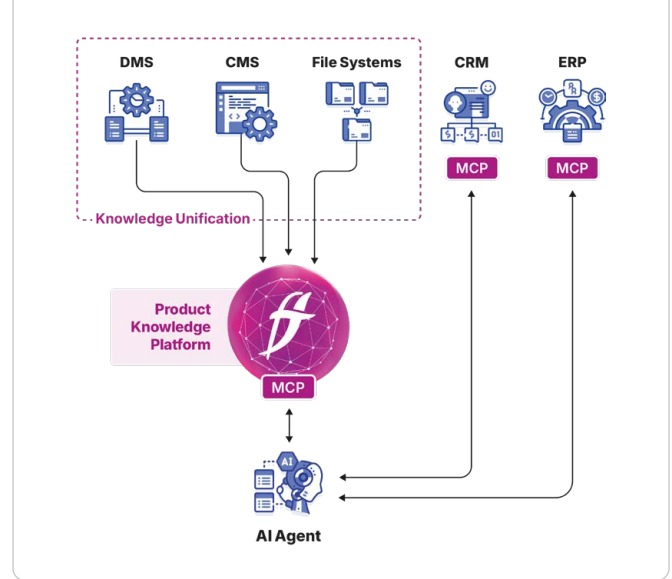
Systems that support LOB solutions, also known as systems of record, are the natural go-to applications: CAD and PLM for product design, CRM for sales and clients, ERP for HR and finance, and more. Their data is structured, enriched with domain semantics—ideal for AI—and they are now equipped with MCP servers.

However, beyond these business applications, content is loosely organized, mostly in files. Knowledge is fragmented across PDFs and slides, often lacking structure or metadata. Sources are numerous and spread across shared drives, wikis, DMSs, cloud-based CMSs, and many repositories that are unlikely to support MCP anytime soon.

To provide AI agents with an authoritative source covering these fuzzy zones, the infrastructure must evolve. KM must adopt modern solutions that span and unify fragmented information for a given domain. They should function as MCP-enabled gateways, enriching content with structure and context, and enforcing access controls. Such solutions are now available, each specializing in different flavors of knowledge (e.g. legal, marketing, product).

## The Product Knowledge Case

Product Knowledge is highly heterogeneous and dispersed across the enterprise yet critically needed for all customer-facing activities. Platforms like Fluid Topics consolidate and centralize this knowledge and enable it for AI.



## Takeaways

Agentic AI is redefining the purpose of KM. The user is no longer a human, but an ecosystem of software agents serving people. Organizations that structure their knowledge for AI will tap into unprecedented levels of efficiency and productivity. Those who don't will experience strong AI limitations. KM now has the opportunity and responsibility to lay a strong foundation for AI and for its own future. ■

# 2026: Moving Beyond “Data-First” to a “Knowledge-First” Architecture



The generative AI gold rush of previous years has matured into a pragmatic operational reality in 2026. The industry's focus, once fixed on model scale and compute power (a “data-first” race), has shifted. The fundamental truth is that scale alone lacks value; context is key. The defining shift of 2026 is from powerful models to meaningful ones—moving from “power to purpose” and adopting a “knowledge-first” architecture over “data-first.” Organizations are no longer content with successful pilots that fail to scale. They are dismantling the “digital quicksand” of fragmented, unstructured data and rebuilding on foundations of semantic interoperability. The following six key areas define this new enterprise horizon.

## Modern Knowledge Management Platforms and Tools

By 2026, the “symbolic” vs. “neural” AI debate is over. Successful Knowledge Management platforms use hybrid “neuro-symbolic” architectures, combining LLM statistical intuition with knowledge graph structured reasoning. Modern KM tools are “insight orchestrators,” no longer relying on a single model. They switch between open-source and proprietary systems, acting as enterprise backbones. Value creation hinges on blending GenAI's creativity with domain-specific governance, precision, and explainability, ensuring AI is a transparent, governed participant, not a black box.

## Enterprise Search and Discovery

Enterprise search has evolved from a limited “strings, not things” keyword matching approach to a “knowledge-first” discovery engine driven by user intent. The former “data-first” approach created passive metadata lakes, leading to irrelevant results (e.g., a search for “resilience” yielding psychological definitions instead of risk standards for a banker). In 2026, search uses active semantic metadata to understand concepts and relationships (e.g., “ReactJS” is a skill, a form of JavaScript, and related to “User Experience”). This shift helps organizations overcome “digital quicksand” by linking complex, siloed data, allowing modern discovery tools to deliver results that are not just statistically probable, but semantically accurate and aligned with the user's domain.

## GenAI and Agentic AI

The rise of agentic AI marks the turning point in automation, but it has also exposed a critical “cart-before-the-horse” problem. Many enterprises rushed to build multi-agent ecosystems before establishing a stable knowledge foundation. In 2026, the standard for trustworthy GenAI is [GraphRAG](#) (Graph-based Retrieval-Augmented Generation). Unlike standard RAG, which relies on unverified chunks of text, GraphRAG grounds agents in a semantic backbone—a trusted, continuously updated web of facts. This allows agents to reason, negotiate, and collaborate with traceable logic. Whether it is a single agent assisting an engineer or a swarm of agents managing a supply chain, the difference between a useful agent and a hallucinating one is the quality of its knowledge graph. We are seeing a move toward “governed autonomy,” where [human-in-the-loop validation acts as a flywheel](#), continuously refining the knowledge assets that the agents consume.

## Taxonomies, Ontologies, and Knowledge Graphs

Taxonomies and ontologies have graduated from the backroom of library science to the boardroom of AI strategy. They are now

recognized as the essential infrastructure for trust and explainability. As regulatory frameworks mature, trust has become the currency of AI. Enterprises are using knowledge graphs to engineer this trust, making it possible to trace every AI-generated conclusion back to its source data. This structure turns “passive” data into “active” knowledge, adhering to [FAIR principles](#) (Findable, Accessible, Interoperable, Reusable). By using standard vocabularies and globally unique identifiers, organizations are creating a “semantic digital twin” of their enterprise—a living network of people, processes, and products. In 2026, a robust taxonomy or ontology is the essential guardrail for compliant, factual, and business-aligned AI.

## Intelligent Document Management

Intelligent Document Management in 2026 is defined by the ability to structure the unstructured. By applying semantic models to document repositories, organizations are converting static files into dynamic knowledge assets. This supports critical processes like compliance intelligence and technical knowledge management. For example, rather than just storing a maintenance manual, the system extracts the entities and relationships within it, linking a specific machine part to its failure history and regulatory requirements. This capability drastically reduces the “mean time to resolution” in engineering and manufacturing, transforming documents from dead weight into active participants in the decision-making process.

## Customer Experience and Customer Service Automation

The shift to knowledge-first architectures is transforming Customer Experience (CX). Legacy chatbots were frustrating due to a lack of context. Now, automation is driven by deep, semantic understanding via a knowledge graph, enabling AI to handle complex, domain-specific queries with expert precision. This is critical in regulated fields like finance and healthcare, delivering personalized, auditable, and compliant answers. The system's “traceable logic” explains why an answer is given, justifying it against policy. This restores trust lost to earlier, hallucination-prone bots. By combining automation with oversight, organizations ensure CX is efficient, empathetic, and accurate.

## Conclusion

The clear lesson from the years before 2026 is that building intelligent systems on fragmented, data-poor foundations leads to failure. In 2026, the winners prioritize knowledge, understanding that AI maturity stems from accessible meaning, not model size.

- **The Problem:** Intelligent systems fail on fragmented, data-poor foundations.
- **The Winning Strategy:** Leading enterprises prioritize a “knowledge-first” approach.
- **True AI Maturity:** Success hinges not on model size, but on the depth of accessible meaning and context.
- **Key Adoptions:**
  - Hybrid architectures
  - Governed knowledge graphs
  - Human-in-the-loop processes

**The Outcome:** This shift embeds intelligence into operations, moving from raw models to meaningful, actionable, and scalable intelligence. ■

Additional information found on page 19

# How M-Files and Microsoft Are Revolutionizing Context-first Document Management with AI and Automation



For years, executives from Microsoft and document management platform provider M-Files would meet across the aisles at industry events, often as competitors. Today, the narrative has shifted dramatically. The two companies are now strategic partners, collaborating to deliver integrated, context-first solutions that redefine productivity, security, and intelligence for the AI era. "All the years we would see each other at different events, we would be competing," recalls Ian Story, Principal Architect for OneDrive and SharePoint at Microsoft. "I'm excited just to be sitting here now as partners." For Antti Nivala, Founder and Chief Innovation Officer of M-Files, this partnership fulfills a vision that began with the company's inception: creating a frictionless, intelligent experience for end users by embedding content management into the tools they rely on every day. "Most of the customers we work with use Microsoft tools as their digital workplace," says Nivala. "So, it's been important for me to make sure that M-Files integrates with Office tools, with the Windows operating system, and today with Microsoft 365 in the cloud."

## From Integration to Intelligence: AI at the Core

This collaboration has reached a new level. M-Files now enables users to co-author and edit documents simultaneously in Microsoft 365 desktop applications while the M-Files platform organizes information, guides business processes, and automates security and compliance. "Co-authoring is very important for M-Files users," says Nivala. "M-Files serves customers best when it's not just the system of record, but also the system of work. Compliance becomes automatic. You can just do your work while everything else happens behind the scenes."

Story underscores the depth of integration:

"It provides 100 percent of the native co-authoring features of Microsoft 365—comments, mentions, presence awareness, working with Copilot—not partial functionality."

Underpinning these capabilities is SharePoint Embedded, an API-only solution that enables developers to leverage Microsoft 365's file and document storage functionality without exposing the user interface. Users interact through M-Files' interface, combining the best of both worlds.

## AI as the Game-Changer: Unlocking Context and Velocity

While integration is powerful, AI is transformative. Both companies see AI as the next frontier for knowledge management. With Microsoft Copilot and M-Files' context-first, metadata-driven architecture, organizations can unlock unprecedented value and gain a velocity advantage, reducing operational drag, and unlocking trapped business knowledge.

"Having Copilot natively access content is an exciting prospect," says Nivala. "By exposing business context and documents in M-Files to Copilot, alongside the Microsoft Graph, you really get to the right conclusion for whatever the user needs."

Imagine an AI agent that understands not only your documents but also the business context behind them, customer data from Salesforce, operational data from SAP, and collaboration signals from Teams. This is where M-Files shines: its ability to enrich

content with metadata and context, making AI-driven insights more accurate and actionable.

Story envisions specialized AI agents for industries like tax, accounting, engineering, and life sciences:

"With the business object context from M-Files and collaboration data from Microsoft Graph, it's a really powerful combination."

## AI-Powered Compliance and Security: Embedded, Not Bolted On

Security and compliance remain central to this collaboration. AI introduces intelligence into automation, ensuring sensitive information is protected without manual intervention.

"A very practical example is sensitivity labels," explains Nivala. "Now, thanks to the integration, we can create sensitivity labels automatically with metadata. If a document relates to a confidential project, metadata can trigger encryption and prevent oversharing."

Compliance is no longer reactive; it's embedded in the workflow. Microsoft Purview, FedRAMP, SOC 2, and HIPAA certifications apply seamlessly to content managed through M-Files. Advanced Threat Protection and Microsoft Defender extend to every document, ensuring enterprise-grade security without complexity.

## The Future: Smart, Context-Driven Knowledge Work

The Microsoft 365 and M-Files partnership isn't just about solving today's problems, it's about shaping the future of work. AI will play a pivotal role in making knowledge management effortless, predictive, and personalized.

"The more intelligence and automation we bring into the daily flow of work, the closer we get to making information management effortless," says Nivala. "That's the vision—and we're only getting started."

Story agrees:

"With AI, compliance, and seamless co-authoring built in, we're helping customers move towards a workplace where content is always available, always secure, and always in context."

## Why This Matters for KM Leaders

For KM leaders, this partnership represents a blueprint for the future:

- **AI-driven discovery:** From traditional search to conversational AI, users can find answers—not just documents—faster and more intuitively.
- **Contextual intelligence:** Metadata and business object integration ensure AI understands not just content but its relevance.
- **Effortless compliance:** Automated sensitivity labeling and security policies reduce risk while improving productivity.
- **Scalable innovation:** SharePoint Embedded and M-Files' architecture enable organizations to build custom solutions without sacrificing security or user experience.

As organizations navigate the State of KM + AI, Microsoft and M-Files are proving that the future of knowledge work is not just automated, it's intelligent, secure, and deeply integrated into the tools people use every day. ■

Additional information found on page 20

# For many enterprises, AI isn't delivering. The solution is knowledge graphs.



**metaphacts**

A Digital Science Solution

In recent years, the enterprise landscape has become saturated with AI tools promising unprecedented growth. Yet for many organizations, the return on their AI investment has been disappointing. Instead of driving transformation, current AI strategies often fall short, with applications proving more limited, expensive, and less trustworthy than expected.

We are now in what Gartner describes as the AI “trough of disillusionment”—a phase where enterprises question if AI can deliver on its promises. While it may be tempting to abandon new technologies, this is precisely when a clear strategic vision matters most. The primary issue is that popular tools like large language models (LLMs) are not sufficient for enterprise needs on their own. Neural networks are powerful, but they lack business context, often hallucinate, and produce unexplainable outputs.

*For AI to integrate successfully, it must be trustworthy, consistent, and accountable—this requires knowledge-driven AI.*

Leaders facing underwhelming results have a choice: persist with high-risk AI programs or retreat and fall behind. **Fortunately, there is a viable alternative: fortifying the existing strategy by incorporating knowledge graphs.** This symbolic layer provides the structure and context essential to unlocking AI's true potential within the enterprise.

## THE LIMITS OF LLM-ONLY APPROACHES

AI adoption has accelerated dramatically, with 78% of organizations reporting AI use in 2024. This was driven by claims of extreme productivity gains. However, the reality has been underwhelming. McKinsey found that 80% of companies using generative AI saw no significant bottom-line impact, and 42% abandoned their strategy.

This difficulty highlights a misunderstanding of LLMs. A chatbot may seem to understand context, but its outputs are

statistically likely responses, not rooted in certainty. An MIT study found that 95% of AI pilot projects failed, and the rise of “workslop”—polished but non-functional content—hinders productivity. When outputs require careful fact-checking, they decrease workflow efficiency.

## THE RISK OF RETREAT

Despite these challenges, abandoning AI implementation risks strategic irrelevance. Gartner predicts that by 2026, **enterprises adopting robust AI engineering practices will outperform peers by at least 25%**. Organizations that successfully operationalize transparency and trust will see a 50% improvement in AI adoption and business outcomes. When an enterprise scraps an AI pilot, it risks accelerating its own obsolescence. The high failure rates reflect architectural problems, not organizational incompetence. The solution is to build a more robust, knowledge-driven foundation.

## SUCCESSFUL IMPLEMENTATION RELIES ON NEURO-SYMBOLIC AI

For AI to integrate successfully, it must be trustworthy, consistent, and accountable. This requires knowledge-driven AI: an approach where AI systems are guided by structured, enterprise-specific knowledge. In practice, this means embedding explicit rules and relationships—often captured in knowledge graphs—into AI workflows.

A truly effective strategy is realized through neuro-symbolic AI, which merges two complementary technologies:

- Neural components: LLMs that excel at pattern recognition and language modeling.
- Symbolic components: Knowledge graphs that define the logical structure and rules of enterprise data.

This relationship is like the human brain. The neural network is the analytical “left brain,” while the symbolic layer provides the contextual “right brain,” setting constraints the model must respect. By anchoring AI in organizational knowledge, neuro-symbolic approaches ensure both innovation and reliability.



**metaphacts**

A **Digital Science** Solution

*Enterprises can't afford to settle for AI strategies that deliver only half the vision.*

### BUILDING TRUST AND REDUCING RISK

The "black box" nature of neural networks makes it difficult to explain how an LLM generates outputs, creating operational risk. Neuro-symbolic AI addresses this structurally. The knowledge graph defines the scope of data an LLM can reference, ensuring it interprets information only through approved rules. This leads to:

- **Reduced Hallucinations:** Outputs are grounded in verified data.
- **Traceability:** Compliance teams can trace outputs to their sources.
- **Streamlined Audits:** Risk assessments become defensible and transparent.
- **Controlled Bias:** Governance over data interpretation minimizes bias.

### EMPOWERING USERS AND PRESERVING KNOWLEDGE

Once trust is established, neuro-symbolic AI allows users to integrate information across siloed systems through a conversational interface. This approach also protects enterprises from the loss of institutional knowledge. By capturing not just facts but the relationships and context that give them meaning, knowledge-driven AI creates a living knowledge bank that assists with onboarding, upskilling, and strategic planning.

### UNLOCKING TRUE ENTERPRISE INNOVATION

An AI strategy that frustrates users will never yield transformation. By pairing LLMs with a symbolic layer, organizations can operationalize AI safely and at scale, moving beyond disillusionment into a period of sustained growth.

Enterprises can't afford to settle for AI strategies that deliver only half the vision. A workforce empowered with conversational access to trusted, contextualized enterprise knowledge is within reach.

metis is a knowledge-driven AI platform purpose-built for enterprise. The metis platform brings together the strengths of LLMs and knowledge graphs to deliver enterprise-ready AI agents that combine generative power with semantic precision and return explainable, trustworthy and contextual insights.

**Learn more about metis by metaphacts (a Digital Science solution)!**

Sources:

- **Stanford HAI (2024):** Data on AI adoption rates, reporting that 78% of organizations used AI in 2024.
- **McKinsey:** Findings that 80% of companies using generative AI saw no significant bottom-line impact, and 42% abandoned their AI strategies.
- **MIT Study:** Highlighting that 95% of AI pilot projects failed across major enterprises.
- **Harvard Business Review:** Documenting the rise of "workslop," or polished but non-functional content, and its impact on productivity.
- **Gartner:** Predictions that enterprises adopting robust AI engineering practices will outperform peers by 25% by 2026 and that operationalizing transparency and trust will improve AI adoption and outcomes by 50%. ■

Additional information found on page 20

# Bridging the AI Gap: From Knowledge Chaos to Compliant Intelligence



By Keith Berg, SVP, Upland RightAnswers

AI is changing every corner of the enterprise, from how customer support resolves issues to how IT teams scale knowledge and service. But as leaders move aggressively toward implementing AI, a sobering reality continues to surface: AI cannot perform reliably if the knowledge behind it isn't ready for AI.

This disconnect, known as the Knowledge Activation Gap, is now one of the most significant barriers to AI success. It's the gap between what AI should be able to do (deliver instant, accurate answers) and what it can do inside large enterprises with fragmented, outdated, or poorly governed knowledge foundations.

*Modern enterprises need  
knowledge that is structured,  
governed, and activated across every  
system, not just stored.*

According to the recent Knowledge Activation Report from RightAnswers and Metric Sherpa, most organizations are already experimenting with AI, yet only 29% of leaders consider themselves "AI-ready." In addition, just 18.6% report having highly structured and trustworthy knowledge that AI can effectively use. The result is predictable: inconsistent responses, compliance risks, low agent trust in AI-generated answers, and customers receiving guidance that varies from person to person.

The core issue isn't the AI technology. It's the knowledge infrastructure beneath it.

## From AI Hype to AI Readiness

In the first wave of AI, organizations were experimenting like it was the Wild West. Chatbots were launched, content was automated, and "AI-first" strategies wiggled into every boardroom conversation. Yet, as adoption has accelerated, the cracks have started to show.

Frontline teams, especially those in compliance-reliant industries, report that their AI tools don't always produce consistent or even correct answers. In many cases, the issue isn't the AI, it's the knowledge base it's pulling from. When data is outdated, poorly structured or perhaps not documented at all, even the best algorithm is bound to produce wrong answers.

To move from experimentation to measurable outcomes, organizations must modernize how knowledge is created,

governed, delivered, and continually improved. This is where AI-fueled Knowledge Management becomes essential.

## What a Modern AI-Ready Knowledge Architecture Looks Like

To bridge the Knowledge Activation Gap, organizations need a modern, AI-ready knowledge foundation that puts structured, accessible, and trustworthy content at the top of the list. It should seamlessly integrate knowledge creation, governance, and delivery across the enterprise, ensuring that AI tools can operate effectively and consistently.

A robust knowledge management platform makes this happen effortlessly from a user perspective, but there's a lot going on behind the scenes. By embedding AI throughout the knowledge lifecycle, these platforms enable organizations to create and maintain a single source of truth. These tools can streamline content creation, flag duplicate or outdated information, and ensure compliance through human-in-the-loop governance. This combination of automation and oversight builds a scalable, reliable knowledge foundation.

When enterprises adopt this kind of architecture, the benefits are clear. As one IT leader at Nestlé noted:

*"The RightAnswers platform and partnership from the Upland team has been critical to our success with knowledge and self-service, and the support of our dedicated Customer Success Manager is helping us drive continuous improvement."*

*—Jon Drake, IT Core Technology Manager, Nestlé UK*

This highlights how a well-designed knowledge system can empower teams to deliver faster resolutions, improve customer satisfaction, and scale operations efficiently.

## AI Knowledge Creation and Governance

In most organizations, knowledge creation is slow, inconsistent, and reliant on SME availability. Next-gen knowledge management platforms transform this process with AI assistance that accelerates creation while strengthening quality and compliance.

Good governance starts with structured content creation. AI tools can help authors by drafting content from verified sources, improving clarity, and generating articles for emerging issues without starting from scratch. Automated systems can also detect duplicate or conflicting information, helping maintain a single source of truth. Before any content is published, approval workflows highlight errors, missing fields, or compliance concerns, ensuring that human reviewers can focus on high-value tasks.

---

*AI success begins with knowledge maturity, and the organizations that close the Knowledge Activation Gap will lead the next era of intelligent operations.*

This blend of automation and human oversight exemplifies what effective governance looks like: scalable, compliant, and trustworthy knowledge that evolves with the organization's needs.

#### Connected Knowledge Across the Enterprise Ecosystem

Knowledge doesn't create value if people can't find it. In most enterprises, content is scattered across CRMs, ITSMs, wikis, SharePoint sites, and outdated systems. RightAnswers unifies these sources through a connected knowledge ecosystem that brings the right information directly into the workflows where people work.

With federated search, users can retrieve relevant knowledge regardless of where it's stored. And with RightAnswers X, our browser-based extension, employees from contact center agents to field technicians can access AI-powered knowledge inside any web-based application without complex integrations.

This eliminates the need to bounce between a half dozen applications or tabs, and ensures every employee receives consistent, accurate guidance.

#### Precise, Step-by-Step Answers with Gen Answers

When someone needs information, they don't want paragraphs; they want actionable steps. RightAnswers' Gen Answers feature uses generative AI to synthesize short, precise instructions based solely on verified internal content, from your trusted knowledge base, not whatever AI pulls from the web.

The result is instant, accurate guidance that boosts first-contact resolution, reduces handle time, and gives customers and employees the confidence that the information they're using is trustworthy.

For regulated environments, Gen Answers is particularly powerful because it maintains full traceability to approved knowledge sources and remains governed by human oversight.

#### Turning AI Into Measurable Operational Impact

When enterprises combine structured knowledge, automated governance, scalable delivery, and AI-powered assistance, the results are immediate and quantifiable. Organizations participating in the Knowledge Activation Report saw faster resolution times, dramatically improved FCR, and significant reductions in escalations. New frontline contact center agents onboard more

quickly, and experienced agents spend less time searching and more time resolving.

One of the most important insights from the research is that AI success is tied directly to knowledge maturity. Enterprises with clean, structured, governed knowledge are nearly twice as likely to achieve measurable ROI from their AI initiatives. This underscores a crucial message for organizations planning their 2026 AI roadmap: AI performance is only as good as the knowledge foundation beneath it.

#### The Path Forward in 2026: Intelligent Knowledge, Applied with Purpose

The next generation of AI-driven enterprises will not be defined by the number of AI tools they deploy, but by how effectively they activate knowledge across their entire organization.

RightAnswers provides the foundation to unify AI content creation, governance, KCS-aligned practices, connected knowledge delivery, and trustworthy generative guidance into one modern platform. By giving enterprises knowledge they can trust and AI they can control, we help teams move beyond hype and into measurable, sustainable impact.

Organizations that are modernizing with AI-powered KM like RightAnswers often see:

- Up to 35% faster resolutions
- 40% fewer escalations
- Significant reductions in Mean Time to Know
- Higher FCR and CSAT rates
- Faster onboarding and time-to-proficiency

Whether improving customer support, powering field operations, guiding internal IT workflows, or enabling digital self-service, RightAnswers helps organizations turn knowledge into an operational advantage and turn AI into a strategic asset.

The future of intelligent contact center operations starts with knowledge readiness.

And the organizations that close the Knowledge Activation Gap will be the ones that lead. ■

Additional information found on page 21

# From fragmented to foundational: enterprise AI search as the key to information discovery

By Dhananjay Bapat, Principal Product Marketing Manager, Simplr

Every day, knowledge workers lose valuable time hunting for information that should be at their fingertips. Emails pile up. Slack threads multiply. Documents scatter across SharePoint, Google Drive, and other systems. Even with the promise of AI assistants, more than a third of employees still can't find what they need.

This isn't just frustrating—it's expensive. When employees can't quickly locate the right policy document, customer history, onboarding asset, or product specification, decisions are delayed, meetings get rescheduled, and work is duplicated. The potential of AI-powered productivity hits a wall when the underlying search infrastructure can't deliver accurate, relevant results.

That's why Gartner's recent [Market Guide for Enterprise AI Search](#) is so valuable. The research firm has established a key position, naming enterprise AI search as the foundational layer that powers effective AI assistants and agents.

Simplr was named as a Representative Vendor in the report's in-application search category—which we consider to be validation that our approach to intranet-centered information discovery solves real problems for organizations worldwide.

Here's what Gartner's framework means for IT leaders, where the market is headed, and how to think about building a workable enterprise AI search foundation.

## The information crisis in today's workplace

The statistics are straightforward: 34% of employees struggle to find information at work, according to Gartner. Organizations responded by deploying AI tools—49% of workers now use platforms like Microsoft 365 Copilot and Google Gemini primarily to find data. Yet the [2024 Gartner Digital Worker Survey](#) reveals the uncomfortable truth: 36% of these AI tool users still can't access the information they need.

This creates a paradox at the heart of [AI adoption](#): Companies invest in cutting-edge technology expecting it to solve information discovery problems, only to watch employees continue struggling with the same findability gaps.

## Why current search solutions don't work

Enterprise employees juggle information spread across multiple applications. Customer data lives in Salesforce, product documentation sits in Confluence, HR policies [hide in SharePoint](#), and tribal knowledge gets buried in Slack threads. Each system comes with its own search function, its own permissions model, and its own quirks.

*Context switching between these systems doesn't just waste time—it kills productivity. It delays decisions, duplicates work, and frustrates employees.*

Here's what should worry IT leaders: AI assistants based on current retrieval-augmented generation (RAG) often underperform when scaled across diverse enterprise information. Gartner identifies the culprits as data source quality issues and weak retrieval relevancy mechanisms.

Your [AI assistant](#) is only as good as the quality of information it can find and how relevant the retrieved content is to the user.

Traditional search was built for a different era. Most systems treat search as a feature rather than infrastructure. They offer basic keyword matching without understanding user intent, user-document relationships, or organizational context.

These systems return hundreds of results with unclear relevance ranking. They surface documents users can't actually access. They mix outdated content with current information with no way to distinguish between them. They were designed for surfacing information based on narrow criteria, often without prioritizing relevance.

## Enterprise AI search is the foundation for effective AI assistants

Gartner makes a clear strategic recommendation: Reposition [enterprise search](#) as the foundational platform that powers AI assistants and agents. Search quality directly determines AI quality.

The reasoning is straightforward. AI assistants retrieve information before they synthesize answers. If retrieval fails—pulling the wrong documents, missing context, or ignoring permissions—everything downstream will be flawed.

## How RAG connects search to AI answers

RAG combines search with large language models to generate grounded responses. The search component retrieves relevant documents, then passes that context to the AI to generate an answer. This keeps AI responses tethered to actual enterprise content rather than hallucinating information.

But RAG only works if retrieval works. The search system needs semantic understanding—the ability to grasp meaning, not just match words. It needs to combine traditional lexical search with vector search that leverages the power of LLMs. And critically, it must enforce permissions—users should only see results they're authorized to view, and AI should only synthesize answers from content each user can legitimately access.

Gartner's Market Guide details the specific capabilities and governance frameworks required to make this work at enterprise scale, including how to systematically apply metadata enrichment to ensure information remains accurate, pertinent, and trusted.

## Understanding the enterprise AI search marketplace

Gartner organizes the market into three distinct categories, each addressing different organizational needs. Understanding these categories helps IT leaders map solutions to their specific use cases—and most organizations will deploy multiple approaches. In fact, Gartner predicts that by 2028, 60% of organizations will have more than six enterprise AI search platforms deployed across the business.

## Search platforms: comprehensive and configurable

Search platforms like AWS Kendra, Elastic, Glean, Google Vertex AI Search, and Microsoft Azure AI Search provide broad connectivity across diverse data sources with maximum flexibility and customization. These

solutions work well for complex needs that span many different systems and content types, though they require significant implementation effort and ongoing tuning.

### In-application search: embedded and context-aware

In-application search brings capabilities directly into the applications where work happens. This is where Simplr sits, alongside vendors like Atlassian Rovo, Microsoft 365 Search, Salesforce Agentforce, and SAP Joule.

*The key advantage is context. These solutions understand the specific content, their types, user workflows, and relevance signals within their domain.*

For [intranet-focused solutions like Simplr](#), that means a deep understanding of corporate communications, [knowledge bases](#), and digital workplace content patterns. The system knows not just what content exists but also how employees actually use it, which makes results more relevant from day one.

Deployment tends to be faster because core integrations and relevance models are prebuilt for the primary application environment. Rather than starting from scratch with a blank search platform, organizations get solutions tuned for their specific use case.

Gartner also identifies a third category of solutions that provide AI-powered access across multiple existing systems, including tools like Anthropic Claude and OpenAI ChatGPT Enterprise.

The Market Guide provides detailed evaluation criteria for each category, including specific capabilities to assess and implementation considerations based on organizational needs. For organizations evaluating intranet search solutions, understanding where specific vendors fit within this category—and what distinguishes their approach—becomes the next critical question.

### Simplr's position in the enterprise AI search market

Being named in Gartner's [Market Guide for Enterprise AI Search](#) as a Representative Vendor within the in-application search category puts Simplr alongside companies like Atlassian, Microsoft, Salesforce, and SAP. For organizations [evaluating intranet solutions](#), that recognition signals enterprise-grade capabilities and market relevance.

More importantly, it validates the approach: Intranet search works best when it's purpose-built for intranet content, not retrofitted from general-purpose search technology.

### Why intranet search needs specialized context

Corporate intranets have unique characteristics that generic search engines miss. There's the mix of structured and unstructured content, e.g., policies, announcements, knowledge articles, event pages, employee profiles. There's the relationship between content and organizational structure, e.g., departments, locations, roles. And there's the temporal dimension—some content is evergreen, some is time-sensitive, and some becomes obsolete quickly.

Simplr's in-application search understands these patterns because it's built for them. The system knows that a policy document should rank higher than a three-year-old blog post when someone searches for guidelines. It understands that announcements from someone's own department are more relevant than companywide updates. It recognizes that content engagement patterns—what people actually read and find useful—are strong signals for relevance.

### Governance and permissions built in, not bolted on

One of Gartner's key recommendations in the Market Guide is establishing robust governance with clear policies for managing enterprise information. For intranet search, this means respecting the permissions and access controls already established in the platform.

Simplr's approach treats permissions as foundational, not optional. When an employee searches, they only see results they're authorized to view based on their role, department, and group memberships. When AI assistants generate answers, they only draw from content that specific users can access.

The Market Guide details the importance of metadata enrichment and content quality for effective AI assistants. Simplr gives organizations the tools to maintain accurate, pertinent, and trusted content through governance [workflows](#), content lifecycle management, and analytics that show what's being used and what's going stale.

### The future of enterprise information discovery

Gartner's predictions for 2028 paint a picture of distributed, specialized search rather than one-size-fits-all solutions. By then, 60% of organizations will deploy more than six different enterprise AI search platforms across their business—multiple tools for different needs. Separately, 60% of enterprise applications will have search capabilities embedded directly into them, up from just 20% today.

This proliferation makes sense. Different types of content and different use cases benefit from different search approaches. Customer service teams need search optimized for support knowledge bases. Sales teams need search that understands CRM data and customer history. And employees need intranet search that makes sense of corporate communications and workplace resources.

*The challenge for IT leaders is ensuring these multiple search platforms work together rather than creating new silos.*

Gartner's Market Guide provides frameworks for evaluating interoperability, governance across platforms, and total cost of ownership as search capabilities multiply.

### What this means for your intranet strategy

For organizations that have treated intranet search as an afterthought—good enough if employees can eventually find what they need—the stakes are rising.

*As AI assistants become standard tools for information discovery, the quality of underlying search determines whether those assistants deliver value or frustration.*

Gartner recommends positioning enterprise search as foundational infrastructure that powers AI assistants and agents. For intranets, this means investing in search that understands your specific content, respects your governance model, and improves as employees use it. Organizations that get this right will see measurable improvements in search results, decision speed, and employee satisfaction. Those that don't will watch their investments in AI assistants underperform while employees continue hunting for information the old way.

### Ready to evaluate your enterprise search strategy?

[Gartner's Market Guide for Enterprise AI Search](#) provides the detailed framework IT leaders need to assess current capabilities, identify gaps, and build a roadmap for improvement. The report includes vendor evaluation criteria, implementation considerations, and governance recommendations based on research across hundreds of organizations.

Simplr can help you apply these insights to your intranet strategy. We'll show you how in-application search designed for intranets delivers faster deployment, better relevance, and stronger governance than generic search platforms. And we'll walk you through what a pilot implementation looks like, from connecting your content sources to measuring improvements in findability.

Find out how Simplr's AI-powered enterprise search can change information discovery in your organization. [Request a demo today.](#) ■

Additional information found on page 21

# Why Enterprise AI Is Forcing a Rethink of Knowledge Management

**O**ur indefatigable KMWorld conference program coordinator, Brian Pichman, employed an AI tool to capture the essence of conference presentations. He recorded the presentations, then rolled those into an LLM from which he could query in a number of different ways.

His queries focused on the emerging trends in knowledge management, enterprise search, taxonomies, enterprise AI, and text analytics. His LLM provided a comprehensive overview of all the co-located conferences since it, unlike humans, has the capability to be in multiple places at the same time.

## WHAT FOLLOWS IS AN AI-WRITTEN RESPONSE TO PICHMAN'S QUERIES.

Enterprise AI initiatives are surfacing a hard truth that many organizations have long avoided. The biggest barriers to trustworthy AI are not model limitations or tooling gaps, they are weaknesses in how knowledge is created, structured, governed, and measured across the enterprise.

What once appeared to be a collection of separate challenges, including search relevance, taxonomy design, content governance, and analytics accuracy, is now being understood as a single, systemic issue. Findability problems, low-quality AI outputs, and lack of trust all trace back to fragmented knowledge systems that were never designed to operate as an integrated whole.

This shift is not theoretical—it is being driven by real-world deployments of generative AI, retrieval-augmented generation, and enterprise analytics platforms. As AI is layered onto existing environments, it amplifies both strengths and flaws. Poorly named content, missing metadata, inconsistent information architecture, and weak governance quickly surface as hallucinated answers, irrelevant retrievals, and outputs that cannot be defended or reproduced.

As a result, organizations are being forced to rethink knowledge management itself. The focus is moving away from individual tools and deliverables and toward operating models that integrate content operations, information architecture, metadata, analytics, and governance. The goal is no longer simply to make information easier to find, it is to ensure that

AI systems produce authoritative, contextual, and trustworthy outcomes at scale.

The following emerging themes reflect this shift. Together, they illustrate how enterprise AI is accelerating a transition from artifact-centric knowledge management to capability-driven systems that support reasoning, accountability, and measurable impact.

## REFRAMING THE CORE PROBLEM: FROM SEARCH TO SYSTEMIC FINDABILITY

For years, organizations diagnosed findability failures as search problems. The typical response was another search engine, tuning effort, or interface refresh. That framing is now breaking down.

Practitioners increasingly recognize that poor search results, weak AI answers, and low return on investment stem from deeper issues such as misnamed content, missing or inconsistent metadata, fragmented information architectures, and uneven governance, which can all contribute. Each domain may be optimized in isolation, yet the overall system fails to deliver reliable outcomes.

This convergence is emerging now because AI makes upstream knowledge debt impossible to ignore. Large language models expose weaknesses immediately through hallucinated summaries, irrelevant retrievals, and answers that lack authority or context. In response, organizations are turning to maturity models and cross-functional diagnostics to connect user frustration with root causes.

This shift matters most to knowledge management leaders accountable for business outcomes, AI platform teams that depend on high-quality inputs, and governance and compliance stakeholders responsible for risk and lifecycle control. The core insight is both simple and disruptive: Findability is not a feature, it is a property of the entire knowledge system.

## MINING PATTERNS AT MOMENTS OF NEED

Another emerging shift is methodological. Rather than investing broadly in knowledge management or AI capabilities, organizations are seeking high-leverage interventions. These are points where targeted improvements unlock benefits across multiple workflows.

This has led to renewed focus on moments of need such as situations involved in daily work where people struggle, stall, or create workarounds. Instead of relying on abstract personas or generic surveys, teams are observing real usage patterns to identify repeatable pain points and opportunities for improvement.

Text analytics plays a central role in this shift, but its positioning has changed. Once treated as a legacy

*Another important theme challenges the prevailing copilot metaphor. While AI excels at assisting with tasks, overreliance can lead to shallow engagement, skill atrophy, and loss of professional judgment.*

analytics function, it is now viewed as an AI-enabled form of instrumentation. It helps organizations understand what content is used, when it is helpful, and where it fails to support task completion.

This perspective resonates with product owners prioritizing internal knowledge platforms, operations leaders seeking measurable process improvements, and analytics teams responsible for demonstrating value. The emphasis moves away from content volume and toward content leverage.

## FROM PASSIVE REPOSITORIES TO ACTIVE REASONING SYSTEMS

There is early enthusiasm for generative AI centered on summarization and thematic analysis. In practice, these capabilities have proven necessary but insufficient.

Organizations are learning that enterprise-grade answers require more than probabilistic text generation. Authority, relevance, provenance, and organizational context are critical, especially in regulated or high-stakes environments.

As a result, knowledge platforms are evolving from passive repositories into active reasoning systems. These systems combine retrieval policies, semantic layers, authoritative source control, and contextual grounding to produce answers that can be trusted, explained, and defended.

This shift is particularly important for knowledge platform architects designing semantic and retrieval layers, legal and compliance teams insisting on reproducibility, and business units that rely on accurate answers in areas such as support, research and development, and finance. The value of AI shifts from faster summaries to grounded reasoning.

## HUMAN AND AI COLLABORATION BEYOND ASSISTANCE

Another important theme challenges the prevailing copilot metaphor. While AI excels at assisting with tasks, overreliance can lead to shallow engagement, skill atrophy, and loss of professional judgment.

In response, practitioners are articulating richer collaboration models that emphasize human development rather than replacement. One emerging framing describes ex-formative dialog, in which AI interactions support externalization, encounter, recognition, and integration of human insight.

In this model, AI is not simply a shortcut, it becomes part of a learning loop that strengthens expertise over time. This approach resonates with learning and development teams, research and analytics functions, and leaders concerned with preserving judgment quality and long-term capability.

## DUAL-LAYER RISK: CLASSICAL RIGOR PLUS AI-SPECIFIC FAILURE MODES

Many organizations assume that strong statistical methods and analytical rigor are sufficient to ensure trustworthy AI outcomes. Experience increasingly suggests otherwise.

Traditional risks such as poor study design, bias, and cherry-picking remain critical. At the same time, AI introduces new failure modes, including dependency, automation bias, erosion of expertise, and shifts in knowledge hierarchies. These risks persist even when analytical foundations are sound.

As a result, risk management is becoming dual-layered. Analytical expertise remains necessary, but it is no longer sufficient on its own. AI-specific vigilance must operate alongside it. This view is gaining traction among research leaders, audit and model risk teams, ethics boards, and workforce strategy experts. Trust now depends as much on human behavior as on technical correctness.

## GOVERNANCE MOVES UPSTREAM

Governance is also being redefined. Rather than focusing primarily on downstream controls, organizations are pushing governance upstream into data preparation and cross-system reconciliation.

This shift is especially visible in external and third-party data domains such as Environmental, Social, and Governance (ESG) reporting. Double counting, incompatible geospatial taxonomies, and inconsistent vendor codes undermine both analytics and AI outputs.

AI accelerates the problem because bad inputs scale faster. In response, enterprises are investing in AI-enabled deduplication, normalization, and semantic harmonization before data reaches a model. This approach matters most to data governance and ESG teams, platform owners managing AI pipelines, and legal and compliance functions concerned with auditability. Governance increasingly enables scale rather than constrains it.

## PORTFOLIO THINKING FOR ENTERPRISE AI INNOVATION

Another notable shift is the application of innovation portfolio thinking to AI initiatives. Organizations are recognizing that overconcentration on business-as-usual use cases limits long-term value.

AI surfaces adjacencies where methods transfer across domains and where previously unarticulated problems become visible. To capitalize on this, leaders are building balanced portfolios that span incremental improvements, novel solutions, novel problems, and situations where both the problem and the solution are new.

*Maturity models increasingly integrate content operations, information architecture, governance, and analytics into a single capability stack. This evolution is especially relevant for knowledge and content operations teams, AI application owners, and executive sponsors demanding clear return on investment*

This mindset resonates with strategy and innovation teams, line-of-business leaders balancing risk and scale, and vendor and partner ecosystems operating at the frontier. The objective is disciplined exploration rather than unfocused experimentation.

## FROM TAXONOMY AS ARTIFACT TO EMBEDDED CAPABILITY

Taxonomy and information architecture are undergoing both a linguistic and practical shift. Once treated as static deliverables, they are now framed as embedded capabilities tied directly to outcomes.

A recurring refrain captures the change. Organizations do not need a taxonomy for its own sake, it needs to solve a problem. In practice, this means evaluating taxonomy and information architecture by their impact on search effectiveness, retrieval quality, and content performance.

Maturity models increasingly integrate content operations, information architecture, governance, and analytics into a single capability stack. This evolution is especially relevant for knowledge and content operations teams, AI application owners, and executive sponsors demanding clear return on investment.

## GUARDRAILS AGAINST AI DEPENDENCY

As AI becomes easier to use, its risks become increasingly behavioral rather than technical. Practitioners are already

seeing signs of dependency drift, where AI becomes the default decision-maker even in contexts that require human judgment.

In response, organizations are calling for explicit guardrails. These include role-based usage policies, warning signs for overreliance, and norms designed to preserve human agency in high-stakes decisions.

This issue matters most to people leaders, professional standards bodies, risk and compliance functions, and teams operating in safety-critical or judgment-intensive domains. Responsible AI increasingly requires intentional restraint alongside capability.

## TEXT ANALYTICS AS INSTRUMENTATION, NOT REPORTING

Text analytics is also being repositioned as a core instrument for understanding problem-solution fit. Rather than counting documents or tracking superficial metrics, organizations are using it to assess whether content actually helps users complete tasks.

These insights are fed back into retrieval tuning, ranking, and content design, creating a continuous improvement loop. This approach resonates with documentation and support teams, product and customer experience leaders, and AI teams optimizing retrieval effectiveness. The result is measurement shifts from volume to value.

## WHAT THESE SHIFTS COLLECTIVELY SIGNAL

Taken together, these trends point to a coherent transformation in how enterprises approach knowledge and AI.

Technologically, organizations are moving from model-first approaches to data-, semantics-, and governance-first architectures, with AI layered on top of authoritative and instrumented foundations. Organizationally, integrated operating models are dissolving silos between knowledge management, data, analytics, and governance, while explicitly assigning ownership for AI-specific risk.

The language is changing as well. Isolated deliverables are being replaced by concepts such as patterns, leverage points, authority, context, and capability stacks. There is growing recognition of the gap between generic language model output and enterprise-grade answers, along with a pragmatic focus on the pipelines required to close that gap.

Trust and control are being addressed through dual-layer risk frameworks, authoritative grounding, and explicit practices designed to counter AI dependency. At the same time, human and AI collaboration is becoming more intentional, with interaction models designed to expand human capability rather than replace it. ■



**Fluid Topics**  
 Fluid Topics  
 58 Av. Debourg,  
 69007 Lyon, France  
 Phone: +33 4 72 76 03 80  
<https://www.fluidtopics.com>

## The Product Knowledge Platform for People and AI

Product knowledge is at the heart of your success—yet it's one of the hardest resources to share. Don't let content chaos slow you down and frustrate your users.

Fluid Topics' Product Knowledge Platform is the layer that instantly connects your product and service information to users, applications, and AI:

- **Unified product knowledge**  
 Fluid Topics aggregates and unifies all your publications, documentation, product and support content—from any source and format—into a central, governed platform.

- **Smart user experiences**  
 Enable personalized knowledge experiences tailored to user roles, regions, product lines, and entitlements—across all your channels and touchpoints.
- **The Fast Lane to AI**  
 Feed your AI applications and autonomous agents with trusted product knowledge, prepared and transformed by Fluid Topics for seamless AI use.

[www.fluidtopics.com](http://www.fluidtopics.com)




**Graphwise**  
<https://graphwise.ai>  
<https://graphwise.ai/contact/>

## Unlock Enterprise AI ROI with Knowledge Graphs

Graphwise enables organizations to unlock ROI for enterprise AI by delivering the most comprehensive and trusted industry solution in the field of knowledge graphs and semantic AI.

As enterprises invest millions in AI, Graphwise provides the critical infrastructure to ensure these technologies are ready for scale. This commitment to excellence is consistently validated by Gartner and reflected in several recent accolades, including being named "Data Integration Innovation of the Year" at the 2025 Data Breakthrough Awards and being



recognized in the KMWorld "100 Companies That Matter in Knowledge Management."

Born from the merger of tech visionaries Ontotext and Semantic Web Company, Graphwise operates with over 200 employees across North America, Europe, and APAC.

# M-Files

The Smarter Way to Work.

M-Files

[www.m-files.com](http://www.m-files.com)

(972) 516-4210

## M-Files: Context-First, AI-Native Document Management for the Modern Enterprise

M-Files is reshaping how work gets done by putting documents to work, capturing the context, knowledge, and interrelationships that drive business processes. As a context-first, AI-native platform, M-Files delivers purpose-built solutions that span universal and industry-specific workflows, empowering secure collaboration, automating governance, and enabling intelligent insights.

In the face of explosive growth and complexity, traditional systems leave context trapped in silos and people's heads, creating operational drag and compliance risk. M-Files solves the lost context crisis with a metadata-driven architecture that connects documents to the people, projects, and processes they support, making information instantly findable, governable, and AI-ready.

With native content storage in Microsoft 365, M-Files combines Microsoft's collaboration and security with its own intelligence and governance, eliminating duplication and middleware. This delivers a unified, scalable solution that boosts efficiency, reduces risk, and unlocks trapped

value, giving organizations a velocity advantage.

Built for industries like manufacturing, finance, and professional services, M-Files helps over 6,000 customers in 100+ countries turn documents into strategic assets, ready for AI, compliance, and the pace of modern business.

Key capabilities include:

- **Workflow Automation:** Streamline processes and reduce errors.
- **Compliance Management:** Embed security and regulatory controls into daily work.
- **Content Collaboration:** Enable secure sharing across teams and partners.
- **Knowledge Management:** Capture and reuse institutional knowledge with metadata and AI.

Recognized as a Customers' Choice in the June 2025 Gartner Peer Insights 'Voice of the Customer' report, 94% of reviewers would recommend M-Files.

Learn more at <https://www.m-files.com>.



## metaphacts

A Digital Science Solution

metaphacts

<https://metaphacts.com/solutions/products/metaphactory>

## Digital Science: Advancing Enterprise Knowledge

A part of Digital Science, metaphacts is a semantics and AI company delivering innovative solutions that help global enterprises transform data into consumable, contextual, and actionable knowledge—supporting customers across industries in building semantic layers, creating digital twins, and developing trustworthy AI apps for knowledge discovery. Digital Science is an AI-focused technology company empowering enterprise knowledge workers and decision-makers in research-intensive industries with enhanced data, strengthened collaboration, and powerful insights. Our core solutions, including metaphacts, ReadCube, Dimensions, and Overleaf, leverage AI and Knowledge Graphs to transform fragmented data into structured, actionable intelligence. These platforms unify data for discovery and insights generation and enhance team collaboration. We ensure enterprise teams can fully dedicate themselves to what matters most: advancing knowledge and driving progress for their organizations.

Learn more





## RightAnswers

RightAnswers

<https://uplandsoftware.com/rightanswers/>

## Building AI That Works: RightAnswers Turns Knowledge Into Enterprise-Ready Intelligence

RightAnswers is the first knowledge management platform to be KCS v6 verified and the trusted choice by leading enterprises to transform collective knowledge and improve customer outcomes across their support organizations.

Purpose-built for large customer support and IT service teams, RightAnswers empowers agents to resolve issues 4x faster with a 49% faster search speed. Achieve 80% AI-generated search response accuracy, and scale operations without increasing headcount. Enterprises achieve these results using a proven combination of KCS-aligned workflows and next-generation tools like Gen Answers and RightAnswersX.

By combining AI authoring, duplication detection, and automated knowledge retirement with KCS-aligned workflows, organizations can keep their knowledge base clean, current and compliant, achieving

80% faster review cycles, 85% faster content creation, and 90% faster quality control.

RightAnswers unifies all knowledge sources through federated search, supports bring-your-own AI (BYOAI) for trusted model flexibility, and integrates seamlessly with CRM, ITSM, and self-service portals.

Backed by 25 years of leadership in contact-center innovation and a team of KCS-certified experts, we help enterprises create knowledge ecosystems that continuously learn, adapt, and deliver operational clarity.

Recognized among *KMWorld's AI 100*, RightAnswers combines knowledge with trust and AI with purpose, so every customer interaction builds confidence, loyalty, and efficiency. Learn more at [uplandsoftware.com/rightanswers](https://uplandsoftware.com/rightanswers)



Simplr, Inc.

3 Twin Dolphin Dr  
#160, Redwood City, CA 94065

<https://www.simplr.com>  
PH: (877) 750-8330

## Turning fragmentation into flow: How AI is creating the next-generation digital workplace

Enterprises have spent years adding tools meant to improve communication, collaboration, and productivity. Instead, these disconnected systems created a fragmented digital workplace where employees lose time navigating apps, chasing information, and piecing together workflows. Fragmentation is no longer a minor inefficiency, it is a barrier to clarity, alignment, and performance.

AI is beginning to change this. Rather than functioning as another layer of technology, AI acts as the connective intelligence that unifies communication, workflows, and knowledge into one coherent system of work. It interprets signals across the organization, guides action in context, and personalizes experiences so employees can focus on meaningful work. With governed and transparent AI, the digital workplace becomes adaptive, intuitive, and far more effective.

Simplr brings this vision to life by integrating people, tools, and intelligence into a single unified platform. Instead of juggling siloed applications, organizations gain one workplace hub that reduces redundancy, strengthens adoption, and ensures every interaction moves work forward. This simplifies IT management, improves employee experience, and increases the return on technology investments.

To see the platform in action, visit <https://www.simplr.com>.

As economic pressure and rapid AI-driven change intensify, legacy systems and tool sprawl cannot keep up. Simplr's AI-native approach restores coherence across the digital workplace so organizations can execute with clarity, resilience, and confidence.



### Lexsoft Systems

Calle Lagasca 33,  
1º izq, Madrid 28001  
Spain

info@lexsoft.com  
https://lexsoft.com

## Lexsoft T3 GenAI for Knowledge Research

Legal KM solution, Lexsoft T3 GenAI, is powered by Microsoft OpenAI technology, to deliver Copilot-like functionality and user experience to the KM function. With this out-of-the-box and ready-to-deploy solution, Lexsoft is leading the way in leveraging generative AI to deliver new KM capabilities.

T3 GenAI has evolved legal KM into a “knowledge research” discipline, addressing a longstanding challenge caused by a lack of industry-wide standardization of

taxonomies and associated terminologies – i.e., how to make internal KM systems talk to third-party content platforms. Firms can now connect and integrate their internal knowledge repository with external legal publishers and content providers, enabling users to seamlessly search and research across data sources. There is no violation of client agreements, data protection policies, or compromise of firms’ intellectual property.



**Openstream.ai**  
Delivering Intelligence

### Openstream.ai

Phone: 1 (908) 685-5844  
Email: info@openstream.ai  
Web: www.openstream.ai  
LinkedIn: https://www.linkedin.com/company/openstream

## Explainable, Trustworthy AI—Built for Enterprise Reality

Deploy AI with confidence. [Openstream.ai](https://openstream.ai) delivers explainable, transparent, and hallucination-free multimodal agentic and conversational AI for enterprises where accuracy, compliance, and trust are non-negotiable. Powered by Eva™, our Enterprise Virtual Assistant, [Openstream.ai](https://openstream.ai) uses a neuro-symbolic foundation to ingest, reason over, and activate enterprise knowledge—ensuring grounded, verifiable outcomes.

Empower experts with agentic AI that understands complex processes, conducts research, and generates decision-ready insights. Engage customers, prospects, and employees through empathetic multimodal AI avatars—without sacrificing trust. Innovate faster and compete smarter with AI you can explain, audit, and rely on.



Eva™ from  
Openstream.ai



**RUTGERS—NEW BRUNSWICK**  
**School of Communication**  
**and Information**

mi.rutgers.edu  
mi@rutgers.edu

## Shape the future of knowledge management.

Advance how organizations find, trust, organize, and manage knowledge with Rutgers University’s Master of Information (MI). Rooted in science and grounded in practice, the MI prepares leaders for every stage of the information lifecycle—from creation and curation to retrieval, governance, and use.

Choose your focus:

- **Data Science:** Wrangling AI for insight-driven decisions

- **Technology, Information and Management:** Systems, cloud strategy, and information governance
- **Human-Computer Interaction:** UX research, prototyping, accessibility
- **Library and Information Science:** Metadata, taxonomy, and the organization, preservation, and discovery of knowledge

More information available at [mi.rutgers.edu](https://mi.rutgers.edu).