

# Mutually Reinforcing Strategic Priorities: Data and Knowledge Management Are Prerequisites for Achieving AI ROI

Authors:

Whitt Ellis, Director, Operational Excellence

Shobhit Saxena, Director, KM Products and AI Enablement

Venkatesh Coimbatore Kannan, Associate Director, Data Engineering

June 2026

# The Vision

Enterprises are recognizing the immense upside in treating knowledge as a strategic asset:

1. By building proprietary databases and expert knowledge banks, enterprises create a durable competitive advantage.
2. By enabling strong results from AI initiatives through strong data practices, enterprises maximize AI ROI and adoption.

The path to reaching this vision is not groundbreaking. However, it does require more structure and discipline than most organizations have today.

## The Root Causes

For the past decade, organizations have invested heavily in cloud migrations, system modernization, data lakes that are arguably oceans, and now AI. However, **AI ROI is lagging**. Many AI-enabled results disappoint, insights are wrong, and outputs are unreliable. Before further investment, organizations should fix the root causes.

### One root cause of poor AI results is the data layer.

- The complexities of integrating structured and unstructured data from diverse sources and multiple clouds
- The massive volumes of dark data that are unusable because of silos and human behavior
- The increasing risks from governance gaps in data lifecycle management

Without data management, AI leverages incomplete and incorrect data to generate its outputs.

### Another root cause of poor AI results is the knowledge layer.

- The failure to clean and enhance legacy content and structures
- The leakage of knowledge as professional hoard information or exit
- The lack of continuous and sustainable initiatives that reduce manual intervention and rework

Without knowledge management, AI lacks context and tacit understanding to constrain its outputs.

**Together, these gaps undermine the success and speed of AI implementation.**

## Objectives

This report seeks to briefly outline an end-to-end approach to data and knowledge management for enterprises seeking a solid foundation for AI deployment.

It is not a comprehensive technical report, rather designed to help business leaders understand and identify where to prioritize their investment.

We encourage leaders to ask themselves critical questions throughout:

- Can I train and then scale AI on our current knowledge base?
- Is our data protection sufficiently robust to manage access and protect confidentiality as AI begins accessing our knowledge base?
- How much knowledge is getting buried or lost due to poor structure, practices or turnover?

“Business leaders I speak with share similar issues; they face user resistance, are struggling to prioritize, deal with cumbersome legacy systems, or have underestimated the complexity of getting an enterprise organized.

**Transforming data and knowledge management has become intertwined with successfully implementing AI.**

That execution gap is where we step in. We build programs inside our clients' unique workflows, we ensure strong change management, and we deploy every solution with AI and business users in mind.”

– Whitt Ellis

# Structure and Discipline

The least exciting aspects of data and knowledge management—structure and discipline—are also the most critical. AI success sits at the intersection of accurate, accessible data and contextualized knowledge.

## The Foundations = Structure

### The Cloud Foundation

Before cloud platforms, data and knowledge were siloed into local servers, specific systems and rigid infrastructures. Many issues stem from this structural legacy.

Cloud platforms replace siloed infrastructure with elastic storage, scalable compute, centralized security, observability, and policy enforcement. They enable scale and reduce geographic and system fragmentation.

When assessing cloud readiness or current cloud state, we encourage organizations to take a “0 budget” approach. Envisioning from scratch helps you break free the hold of legacy structures.

*“How should the cloud enable our current and future strategy?”*

*“How does my current infrastructure need to evolve to meet those requirements?”*

The answers are unique to each business.

#### **Single Cloud: Speed + Scale + Simplicity**

Single cloud strategies provide operational simplicity, lower overheads, tightly integrated stacks across data platforms, ML/AI tooling and governance that speed up execution. It is an excellent entry point for early-stage adoption as it forces focus and clear intent.

#### **Hybrid Cloud: Residency + Asymmetry + Legacy**

Hybrid cloud strategies accommodate regulatory and data-residency constraints, and are relevant for asymmetric workloads, where certain workloads benefit from elastic cloud compute and while latency-sensitive or critical workloads remain local.

For many organizations, the hybrid approach is more inertia than strategy, stemming from legacy infrastructure investments and phased cloud migrations.

#### **Multi-Cloud: Resilience + Flexibility + Maturity**

Multi-cloud strategies increase business continuity, avoid vendor lock-in, provide flexibility to run workloads on the optimal cloud, and accommodate regulatory and data-hosting requirements agilely. This requires mature DevOps and engineering teams to manage complex integrations and increases governance overheads.



Evalueserve helps clients across industries transform towards a modern data management framework. We manage complete data environment migrations from on-premise data warehouse to the cloud, and build analytical warehouses, data marts, and data pipelines to stitch together information for reporting, dashboarding and AI use at the enterprise level.

Recently, we've helped the safety team of an infrastructure development firm to perform enhanced data analysis across safety metrics, incident reports and compliance data, leading to an optimization of resource utilization and improvements in overall safety metrics. For a leading logistics firm, we managed the cloud migration, enabling optimization of operations and improvements in vendor management through better data structure, access and analytical workflows.

## The Information Foundation

AI requires data that is accessible, reliable, secure, and cost-effective. This demands scalable ingestion, strong data quality practices, governance, and DataOps/MLOps controls.

### **Structured Data**

Structured data from ERPs, CRMs and databases should follow documented data models and schemas and with consistent quality controls. These are ongoing activities that most organizations already have in place, but they tend to run uniquely across business units or geographies. Enterprises seeking to deploy AI at scale need to **establish modern data management (MDM) to enable semantic layers, BI and AI workflows.**

### **Unstructured Data**

A large focus in knowledge management programs centers around taming unstructured data. These are files like documents, PDFs, presentations, emails, recordings, and videos that remain hard to query using traditional relational data models.

**Success starts with diagnosing where knowledge resides, who owns it and how it is used.** Having a vision of the AI initiatives that will tap into that knowledge helps prioritize effort for consolidating content into scalable repositories and making it machine-readable through OCR, parsing, and embeddings.

### **Dark Data**

Queue the menacing music... Lurking in the shadows of every on-prem server and the corners of cloud storage, lies dark data!

Dark data is saved but is not actively used to drive business outcomes. It is a menace because it increases costs, is missed insight, and potentially hides compliance or security risks.

**Enterprises should run a one-time federated indexing, archiving and deletion initiative across legacy content.**

This surfaces what content exists and where, quantifies dark data and duplication, and helps understand the scale, formats, owners, and recency. Valuable dark data can be restructured to make it useful. Enterprises should create a sustainable process to avoid dark data happening again.

## **Tacit Knowledge**

Tacit knowledge is the knowledge individuals carry themselves based on their experiences, intuition and judgement. **While it's not effective to try to codify all tacit knowledge, high-value elements should be captured** through retrospectives, communities of practice, and the systematic storage of critical work artifacts (e.g., strategy meeting meetings). AI can surface the patterns in decision making and provide guidance based on better context.

Evalueserve supports a philanthropic program team in externalizing and retaining knowledge using tools already available within Microsoft 365. Given the high turnover of experts, critical context was at risk of being lost over time.

OneNote was established as a centralized repository for team-level meeting notes and discussion points, providing a single place for knowledge that requires broad visibility. OneNote supports both manual capture of insights and Copilot-assisted meeting notes, and can be referenced by Copilot to surface, summarize, and reuse documented information across the team. Loop was formalized for confidential conversations, which established clear information barriers as part of the approach to governance.

Modern data architecture is flexible and customizable.

What's needed is investment to ensure the base of information is correct and well structured, then the commitment to continue investing in discipline around information storage and quality processes.

## The Semantic Foundation

With an understanding of the data and knowledge available, where it is stored, who generates it, and who uses it, enterprises should then focus on establishing robust semantics. This is a quickly evolving area with impactful AI-enabled solutions for building and implementing semantics.

### **Metadata**

Metadata describes the attributes of a knowledge asset like a document or data point. These attributes include author, date created, file size, document type, department, version number, and confidentiality level. It helps AI know recency, authority, and sensitivity.

Additionally, **AI can automate metadata tagging at scale, with domain-led quality assurance.**

### **Taxonomy**

Taxonomy organizes content through categories and hierarchies. It can be visualized as a parent-child tree or reference nodes, such as: Industries, functions, and topics. Taxonomy helps AI label, filter and find information. It is necessary for effective retrieval of information.

**AI can support taxonomy creation** by proposing categories and hierarchies, identifying synonyms to support correct classification, and resolving ambiguity through context when terms express more than one meaning.

## Ontology

**Organizations are now moving beyond simple taxonomies to enterprise-level ontologies.**

Ontology defines the types of entities, their relationships and rules for how they interact. This resolves ambiguity and enables more reliable retrieval and reasoning. Upper- or mid-level ontologies such as the Common Core Ontology (CCO) provide a shared semantic foundation across systems, allowing enterprises to reuse common concepts.

## Knowledge graph

Knowledge graphs instantiate ontology with real data, enabling multi-step reasoning. Importantly, it is not a black box. It provides an evidence trail and rationale grounded in operational reality.

“Organization-specific ontology needs to be thought of from a point of scalability and ease of implementation. It’s not a static, one-time set, but rather needs to evolve with the organization.

We start by defining clear use cases, such as enterprise search and agentic workflows, and then set scope boundaries. From there, we use our domain and business expertise to create organization-specific ontologies around their business concepts, operating model, and sector terminology. It’s important to avoid over-modeling in the early stages and take an iterative approach.”

– Shobhit Saxena

---

**Companies should adopt a customized, pick-and-choose approach to semantics that meets the requirements of the use cases they expect to deploy and the reality of their industry.** For example:

- A financial services firm has different knowledge failure modes (regulatory knowledge, deal precedent, risk frameworks) than a manufacturing firm (engineering specs, maintenance knowledge, process documentation). Industry-specific.
- If an organization foresees only the requirement for enterprise search (data retrieval), then metadata and taxonomy could suffice. For companies seeking to deploy GenAI and Agentic workstreams, more comprehensive semantics are required.

Evalueserve partnered with a global professional services organization to modernize its knowledge management approach for business development. The engagement began with the design of a practical, business-aligned taxonomy to organize large volumes of unstructured proposal and RFP content distributed across regions and business units. Using a combination of business-aligned taxonomy design and AI-enabled tagging, classification, and redaction, legacy content was standardized and structured to improve consistency, relevance, and ease of reuse. This ensured that the KM foundation reflected how proposal teams actually work, rather than relying on abstract or theoretical classifications.

Once the foundation of knowledge was in place, AI and automation were applied to accelerate discovery and reuse at scale. Automated document parsing, contextual search, and recommendation capabilities enabled faster creation of zero-draft proposals while maintaining quality and governance. The solution was designed to be self-sustaining through clear workflows, embedded governance, and continuous AI-assisted ingestion, allowing new content to enrich the system over time. As a result, knowledge management evolved into a core part of day-to-day business processes rather than a one-time transformation effort.

## The Operational-Enablement Layer

Even with strong data and semantics, AI remains unreliable without operational enablement. This layer bridges “data readiness” to “AI impact” by transforming assets into AI-ready artifacts.

Semantics become executable once meanings are consistent. This is achieved through normalization, for example, the standardization of entity names, resolution of synonyms using taxonomies, and alignment with business concepts, and through enrichment of raw data, for example, with semantic tags or relationship inference from knowledge graphs.

Connections across sources for AI's use become feasible through cross-source indexing and retrieval structures. AI ingestion pipelines must be optimized for high-frequency access, chunk-level granularity, context preservation and latency sensitivity. The retrieval system is fundamental.

The RAG layer grounds AI responses on validated enterprise data that provides the LLM enhanced contextual understanding. While RAG is generally considered the go-to approach to resolve hallucination concerns, since responses are grounded, it does not mean AI responses will be correct.

**In fact, the failure of RAG is more subtle than a generic LLM since responses stem from validated data.** This is why having a well-structured approach to data and knowledge management has become so critical.


---

## Data Quality Ensures Operational Readiness for AI

A strong and structured foundation that is established through robust engineering is critical to supporting solid reasoning over imperfect and ever-changing enterprise knowledge.

Enterprises that seek a competitive advantage in the era of AI need to critically assess their foundations.

*Up next, Discipline*



It is because of the ever-changing nature of enterprise knowledge that companies need a disciplined approach to knowledge lifecycle management.

## The Lifecycle = Discipline

The knowledge lifecycle comprises the foundational aspects already discussed as well as governance, maintenance, and end-of-life management. These three aspects are where discipline pushes and pulls – the teams overseeing knowledge management need to push systematic solutions onto users and pull behavioral changes out of users.

“Our clients are quickly understanding how fundamental the knowledge base is to AI. What many teams across industries are now grappling with is how to make governance and maintenance functional.

No one wants to overburden teams with convoluted processes or endless SOPs, but everyone acknowledges that more structure and discipline are needed to achieve a return on the up-front investments being made into organizing data and knowledge.

Evalueserve’s approach to quality AI-enabled managed services fills this execution gap. We work with our clients to establish knowledge lifecycle solutions that enable successful AI deployment.”

– Whitt Ellis

### Unified Governance Committee

As organizations become more dependent on using their data and knowledge in AI applications, the governance questions become both more complex and more consequential.

- Who is accountable for the accuracy and recency of data?
- How is sensitive information protected from inappropriate AI retrieval?
- How do I audit a decision that an agent made autonomously?
- How do I control compute spend?

Deloitte's 2026 *State of AI*<sup>1</sup> report found that only one in five companies has a mature model for governing autonomous AI agents — a striking gap given the speed of deployment.

**There is a clear convergence of data, knowledge and AI. This should be reflected under the Unified Governance Committee.**

<sup>1</sup> [The State of AI in the Enterprise - 2026 AI report | Deloitte US](#)



## The Unified Governance Committee

The **Executive Council** should comprise the executive sponsors, key functional leaders such as information security and legal, leaders of the data and knowledge architecture, leaders of AI initiatives, and business user representatives. It sets enterprise-wide, cross-functional priorities, approves use cases, defines risk appetite, and sets forth a company's policies and standards.

**Working Groups** translate policy and strategy into action. One might deploy an AI-for-KM group to manage the lifecycle of all AI tools being used across knowledge management activities. Or a Semantics group to manage that lifecycle, ensuring ontologies are updated as organizations evolve. What's important is that these groups focus on concrete actions and remain agile.

**AI Use Case Owners** take end-to-end responsibility of that assets' lifecycle. These are the individuals that do the heavy lifting.

### Leaders should watch out for 3 warning signals:

1. Employee turnover and knowledge hoarding derail AI initiatives.
2. Agents trained on narrowly curated pilot datasets.
3. Agents with no clear owners or maintenance schedule.

Many companies and teams develop and deploy incredible solutions but do not provide ongoing maintenance or user adoption support. The initial results are promising, but as time passes – and this is a short period because data is so dynamic – the Agent struggles more and more to perform. Over time, the Agent must contend with semantic drift, changing data, evolving taxonomies and user behavior. Poor performance leads to usage drop off. These become Orphaned Agents, and they represent a waste of investment.

Enterprises need to invest in the lifecycle management of Agents and AI initiatives. This directly links to ensuring ongoing data and knowledge management practices are established and sustained.

**Formality supports continuity.** Under the leadership of a Unified Governance Committee, ideally an enterprise is able to monitor all AI and AI-enabling initiatives, establish structured processes for gathering SOPs and user feedback, and enable redundancy or succession planning for critical tasks.

## Governance Architecture

Under the guidance of a Unified Governance Committee, organizations can more effectively nurture their knowledge governance architecture.

- **Data classification and access controls:** Set up robust classification frameworks, combined with role-based access controls, to ensure that sensitive information such as client data, personnel records and proprietary methodologies are protected.
- **Regulatory and compliance alignment:** For multinational organizations, knowledge governance must navigate a complex and evolving landscape of data privacy regulations, sector-specific compliance requirements, and emerging AI governance frameworks. Establish continual monitoring for regulatory changes and a steering committee.
- **AI model management and auditability:** Maintain clear records of which AI models are operating on which knowledge assets, under what constraints, and with what accountability mechanisms supporting monitoring for model drift, hallucination risks, and training data quality.
- **Provenance and attribution:** As AI systems synthesize and generate content, maintaining clear records of the source knowledge underlying AI outputs is essential for trust, accountability, and audit.
- **Spend and optimization:** As the use of AI becomes more ubiquitous, centralizing functions should periodically re-assess their AI stack, negotiate enterprise pricing, and conduct internal audits to identify prompt optimization opportunities.
- **Documentation and User Manuals:** Knowledge hoarding and lack of documentation stem from human behavior. Establish a formal approach to how AI and AI-enabling activities are conducted, documented, improved and made sustainable over time (e.g., employee turnover).

**Properly designed, governance is what makes agility feasible,** because it means that as the organization moves faster, it does so with confidence that the knowledge underpinning its decisions is trustworthy.



The pillars of governance vary by industry.

### Examples

For regulated industries like life sciences, Unified Governance Committees need to address data minimization and consent traceability as well as regulatory compliance for clinical data.

In Financial Services, an additional pillar to address is model risk management.

In manufacturing, an additional pillar could be the threshold for maintaining operational health and safety while AI recommends efficiency opportunities.

## Reducing the Human “Middleware”

Knowledge repository maintenance is a daunting task because so many users are involved and manual maintenance does not scale.

**AI is a tool that enterprises should use to shift recurring and routine processes off humans.**

Fortunately, there is an ever-expanding list of AI point solutions and AI systems that autonomously execute multi-step knowledge workflows to support knowledge management teams. McKinsey's 2025 research found that 23% of organizations are already scaling at least one agentic AI system, with knowledge management and IT among the leading deployment areas.<sup>2</sup> For knowledge-intensive organizations, these agentic use cases offer truly exciting efficiencies.

Agent that automatically applies metadata tags when files are added to repositories or systems to drastically reduce manual effort and close this semantic gap.

Agent that scans across repositories to detect and recommend emerging terms, concepts and synonyms, categories rarely used, and content clusters to enable frequent taxonomy / ontology refreshes.

Agent that scans across repositories to identify duplicates / near-duplicates to keep repositories clean and ensure AI does not reference multiple versions of a similar file.

Agent that leverages metadata (e.g. dates) and inactivity (e.g., not opened in past 12 months) and automatically archives or recommends to archive in order to reduce manual effort and ensure AI references recent information.

Agent that extracts tacit knowledge from call transcripts or minutes and stores that into a repository to reduce dependency on humans for capturing tacit knowledge.

The critical dependency for reducing human middleware is, again, the foundation.

As a 2025 Deloitte survey found, nearly half of organizations cited searchability of data (48%) and reusability of data (47%) as the primary challenges to their AI automation strategy<sup>3</sup>. Solid knowledge infrastructure allows enterprises to replace portions of human middleware with Agents.

<sup>2</sup> [The State of AI: Global Survey 2025 | McKinsey](#)

<sup>3</sup> [Agentic AI strategy | Deloitte Insights](#)

## End-of-Life Management

All knowledge assets have a lifecycle. Consider that presentation that you have been agonizing on, pouring in hour after hour of your time to create the perfect storyline. It is relevant today but what about in six months? How about two years from now?

### As content progresses through its lifecycle, ROT creeps in.

- **Redundant:** Content is duplicated and exists in multiple places, often in slightly different versions.
- **Obsolete:** Content is outdated and has not been updated or archived.
- **Trivial:** Content has little or no value to begin with, for example, superseded templates and draft versions.

Users are good at creating content and bad at archiving or deleting old content. This needs to change. All data and knowledge need to pass through regular auditing, archiving and updating. Oversight from a Unified Governance Committee can help push activities to Working Groups, Use Case Owners and business teams, and can also roll out centrally beneficial tools like an Agent that recommends documents for archival.

## Data Quality Enables Operational Success for AI

Enterprise knowledge is dynamic by nature. It is shaped by changing businesses, evolving regulations, and human behavior. This is why discipline across the knowledge lifecycle essential.

AI initiatives succeed when governance, maintenance, and end-of-life management are treated as operational responsibilities, with clear ownership, unified governance, and repeatable lifecycle practices that keep knowledge accurate, current, and usable for AI systems.

*“We must spend more time ensuring that AI and humans surface and ingest the most relevant and recent information. When the content lifecycle is ignored, AI is propense to generate incorrect results.*

*This is why we focus on closing the execution gap. We help our clients operationalize data and knowledge with high quality, so AI investments translate into sustained business outcomes.”*

*– Venkatesh Coimbatore Kannan*

Evalueserve helps organizations achieve enterprise-grade knowledge management by turning the Microsoft SharePoint ecosystem into a governed and reliable knowledge platform. Rather than treating Copilot as a standalone capability, we work with clients to strengthen the foundations that KM depends on, including consistent knowledge capture, metadata quality, search relevance, lifecycle governance, and clear ownership. Using native Microsoft capabilities such as SharePoint, Microsoft Search, Syntex, PowerApps, and Copilot, we ensure AI operates on trusted, current, and permission-controlled content across the enterprise.

We design ingestion workflows across SharePoint and enterprise data sources, implement quality and compliance controls at the point of entry, and deploy agent-led processes for classification, sanitization, lifecycle management, and ongoing content health monitoring. Alongside technology, we help establish the right operating model, so knowledge remains accurate, compliant, and actively reused as AI adoption scales.

Through long-term engagements across consulting, financial services, and large corporations, we help clients move from fragmented repositories to living knowledge environments. By applying a repeatable KM framework on the Microsoft stack, we enable faster deployment, lower risk, and clearer returns on AI investments.



## The Vision Emerges = Intelligence

In the ideal world, enterprises would address structure and discipline around knowledge before deploying AI solutions. The reality is that these run in parallel.

### *Semantic search*

Semantic search is the first visible manifestation of intelligence for most users. Instead of forcing employees to know *where* information lives or *how* it is labeled, semantic search allows them to express their need in natural language and retrieve relevant knowledge from many sources based on meaning, context, and relationships. This is thanks to the effort made to establish embeddings, ontologies, and knowledge graphs.

In practice, semantic search becomes the connective tissue between curated enterprise knowledge and everyday decision-making. It dramatically improves discoverability, trust, and reuse of information. It is often the precursor to more advanced AI capabilities because it establishes a reliable retrieval layer upon which agents and automated workflows operate.

### *Agents*

Agents interact directly with enterprise systems by calling databases, invoking APIs, performing calculations, triggering workflows, and chaining multi-step actions to achieve outcomes.

**In this model, AI becomes an active participant in business processes, just like a team member. For employees, this shifts AI from a tool they query to a collaborator.**

### *Front-end platforms*

For enterprise-grade deployments of AI solutions, companies need to deploy front-end platforms that unify access to semantic search, agents, and AI-enabled workflows within a consistent, secure, and intuitive user experience that supports adoption.

These platforms act as the orchestration layer between users and the increasingly complex AI ecosystem operating beneath the surface.

An effective front-end platform abstracts technical complexity while reinforcing governance. It ensures that users see only what they are authorized to see, that AI responses are grounded in approved knowledge, and that interactions are logged and auditable. It provides a channel through which feedback, corrections, and usage patterns can be captured. This should be captured by the Unified Governance Committee and cascaded to the relevant owners to enable continuous improvement of the underlying data, semantics, and AI solutions such as Agents.

At scale, a front-end platform becomes the place where enterprises manage rollout, measure impact, control spending and evolve capabilities in line with business priorities.

**These platforms become the hub for operationalizing AI.**

Intelligence emerges when employees can reliably access, trust, and act on AI-enabled insights as part of how work actually gets done. Proactively providing these tools to employees significantly accelerates AI adoption across an organization; it is far more effective than asking everyone to innovate and deploy in pockets.

## A Necessary Cultural Shift

The primary cause of failure will not be because of technology but because of people.

For an organizations' knowledge ecosystem to become collaborative, structured and disciplined, the humans creating, managing and using content need to change their behaviors.

“Whenever we start a knowledge management engagement, we first sit down with stakeholders from every function, every team, and every level. We get to know them as people, people that create and work with these knowledge assets and processes every day. We treat their assets and processes with respect. And we generate their buy-in to the change by clearly highlighting the benefits that individual or group will gain by embracing the changes.”

– Whitt Ellis

## 6 Musts for Change Management

Since knowledge management initiatives require employees to change how they work, strong change management and intentional, sustained investment are fundamental to nudge change and move toward sustainable processes.

- 1. Personalization:** Map the different roles and understand how they interact with knowledge assets and processes. Tailor solutions at the enterprise and personalized levels.
- 2. Overcommunication:** Change creates concern and concern creates resistance. Overcommunication at all levels across the organization helps identify resistance early on and identify who needs more attention.
- 3. Leading by Example:** When leaders visibly use and contribute to the organization's knowledge systems, adoption follows.
- 4. Sustained Reinforcement:** Behavior change is not a one-time event. It requires ongoing coaching, feedback loops, recognition of good practice, and regular recommunication of standards.
- 5. Incentives:** Knowledge hoarding provides an individual with a sense of status and security but ultimately harms knowledge continuity and AI initiatives. Business culture, structural interventions and incentives help dismantle the hoarding culture.
- 6. Benefits First:** Ensure that employees understand how the changes will benefit them. Paint a picture of how their upfront effort will translate to more efficiency, greater ease, and the opportunity to have more time for value-added tasks.



# Assess Your Organization: 8 Questions Every Leader Should Be Asking

Data and knowledge need to be viewed as strategic assets, and no longer as standalone solutions. Their management needs to be undertaken with cross-platform, cross-function, cross-location collaboration in mind.

Use these questions to diagnose the maturity of your organization's knowledge management and your readiness for AI:

- 1. Can any employee in your organization find any document, data set, or piece of institutional knowledge they legitimately need without asking a colleague?** If the honest answer is "rarely," assess your foundational infrastructure and architecture.
- 2. What percentage of your organizational data would you classify as clean, current, and trustworthy enough to train an AI model on?** If this number is uncomfortably small, consider training AI models on curated data (to not slow down progress) while in parallel, launching an initiative to improve the information foundation.
- 3. Do you have consistent naming conventions, tagging standards, and metadata protocols across your major knowledge repositories and are they actively followed?** If this is inconsistent or nonexistent, first confirm the vision for AI deployment and necessary constraints, then invest in designing a functional semantic layer.
- 4. Do you have a data governance framework that classifies knowledge by sensitivity, defines access controls, and establishes clear retention and retirement policies?** Without this, compliance and AI governance risks accumulate rapidly. If the answer is no, closing this gap should be an urgent priority.
- 5. Who in your organization "owns" knowledge management? Do they have the mandate, budget, and cross-functional authority to enforce standards?** If the answer is no, assess the scale of knowledge management required and use that to inform the scale of governance and formality needed.
- 6. When you last launched a KM initiative (e.g., a new platform, a new taxonomy), what was your adoption rate 12 months later?** If the adoption was low, diagnose what could have been done better and use that to inform future initiatives. Focus on gaps in change management and clear ownership.
- 7. When a senior expert leaves your organization, how much of their knowledge leaves with them?** If gaps in tacit knowledge retention creates risks to your business continuity or competitive edge, identify the priority knowledge to retain, then establish systematic and one-off mechanisms for capturing that tacit knowledge.
- 8. Is data and knowledge management clearly included in your organization's AI strategy?** If not, pull together the AI strategy and relevant data and knowledge leaders to create a unified vision. This should be a robust strategy tear-down.



## Return on Investment: Getting This Right Delivers AI Success

Investing in data and knowledge management leads to clear outcomes - faster innovation, stronger client relationships, lower operational risk, and a workforce that can actually realize the potential of the AI tools the organization is investing in.

The practical path forward requires moving through three stages, with structure and discipline:

- **The Foundations:** Cloud, data, semantics and AI-enablement layers
- **The Lifecycle:** Governance, automation and lifecycle management
- **The Intelligence:** AI and BI modernization

For a large enterprise, this is a multi-year transformation program that requires continual investment. The organizations that start now - those that treat knowledge management as the strategic priority - will see greater RoI from AI investments and ultimately, will compound their advantage with every passing quarter.

**The organizations that win in the AI era are not the ones with the most sophisticated models.**

**They are the ones with the best-managed knowledge.**

## About Evalueserve

Evalueserve is a global quality AI-led managed services company that helps enterprises build and operate the data, knowledge, and workflow foundations required for AI at scale.

Combining deep domain expertise with data engineering, AI enablement, managed services, and change management, we move organizations from fragmented repositories and isolated pilots to governed, reusable, AI-ready environments.

Across ecosystems including Microsoft and Google Cloud, we operationalize complex, knowledge-intensive workflows within clients' existing environments to deliver measurable outcomes that scale.

## Contact us:

Whitt Ellis, Director  
[whitt.ellis@evalueserve.com](mailto:whitt.ellis@evalueserve.com)