


 **BA Insight**



**eBook**


# **The Enterprise Knowledge Graph Playbook**

Building AI Systems  
That Actually Work




# Contents

- PG02** From Connectivity to Context
- PG04** Understanding the Knowledge Graph Imperative
- PG05** The Hallucination Issue is a Knowledge Issue
- PG06** The Cost of Inaction
- PG08** Knowledge Graphs Explained
- PG09** Breaking Down Knowledge Graphs
- PG13** The Business Case for Knowledge Graphs
- PG15** Industry-Specific Value



# From Connectivity to Context



## Discover the Missing Layer in Enterprise AI

Most enterprises have already invested in AI. You may have deployed Microsoft Copilot. You may be piloting Azure OpenAI or Amazon Q Business. Budgets have been approved, teams are experimenting, and expectations are high.

Yet many organizations are seeing a similar pattern. Early demos look promising... and then the real-world use begins. Answers are inconsistent and responses lack context. Security and compliance teams raise concerns about traceability. Adoption slows as user trust declines.

This is not a model problem. It's a knowledge foundation problem. This distinction is critical as the quality of the data foundation directly shapes AI trust, regulatory confidence, and the business value realized from existing AI investments.

95% of AI projects fail to move beyond experimentation into production. This is due to poor data quality and lack of preparation, which means AI is working with incomplete, fragmented, or unstructured information.

# 95%

of AI projects fail to move beyond experimentation into production.

*MIT, State of AI in Business 2025*

## What progress really requires

At the same time, there is growing confusion in the market around the term “graph.” Many vendors refer to their connectivity layer as a graph. That has led organizations to assume they already have a knowledge graph in place.

A connectivity layer is important. It provides secure API access to enterprise systems. It respects source-level permissions and allows applications to query across repositories.

What it does not do is create understanding. There is a significant difference between mere access and actual meaning.

A true enterprise knowledge graph goes further by connecting systems, enriching content, and mapping relationships across entities so AI can retrieve information with context and precision. It creates structure out of unstructured content, adds meaning, and preserves security. And it also guides retrieval so responses are explainable and traceable.

This difference becomes clear when you look at how AI retrieves information. The BA Insight team often demonstrates how a knowledge graph acts as the decision layer between raw connectivity and AI applications. This is seen in a side-by-side example of a complex query using traditional retrieval, which returns dozens of loosely related documents. The same query using knowledge graph-guided retrieval returns a small set of precise documents, along with the relationship path that explains selection.

For enterprise leaders, this is not an academic distinction, as it directly affects risk, productivity, and ROI.

The goal is straightforward: to help you build a data foundation that supports AI systems your organization can trust. As AI investments continue to grow, the organizations that see measurable returns will be those that treat information infrastructure as a strategic priority, not an afterthought.

It's time to stop seeing a knowledge graph as a 'nice to have' rather than an imperative for agentic AI. It's the structure that allows enterprise AI to work as intended.



## This playbook covers...

Why AI projects struggle and what those failures reveal about your data

The difference between connectivity and a true knowledge graph

The three-layer architecture required for enterprise-grade AI

Best practices to design enrichment and taxonomy strategies at scale

How knowledge graphs improve retrieval for generative AI systems

Steps to measure business impact and reduce risk

# Understanding the Knowledge Graph Imperative

## Why Your AI Is Failing, And Why It's Not the AI's Fault

### Symptoms of a Larger Issue

Most AI projects don't fail because the model is weak. They fail because the foundation underneath the model is not ready.


If you have deployed a generative AI assistant internally, you may have seen some version of this:

- | The system produces confident answers that are slightly off
- | It summarizes documents but misses key details
- | It gives different responses to similar questions
- | It struggles with complex, multi-part queries
- | Users begin double-checking every response

Over time, trust erodes. Adoption slows and the project gets labeled as "promising, but not ready."

Industry research shows that up to 95 percent of generative AI projects never move into full production. Poor data quality and lack of preparation are leading contributors.

These are not surface-level issues. They point to deeper structural gaps in how enterprise data is organized and delivered.



### The 80 Percent Problem

In most large organizations, 80% of valuable content sits outside the systems directly connected to AI tools.

Teams work in specialized platforms. Legal teams use document management systems. Finance relies on ERP systems. Product teams live in Jira and Confluence. Sales operates in Salesforce. Each system has its own structure, access controls, security model, and metadata standards.

When an AI assistant pulls content from only a subset of these systems, the answers reflect that limited scope. Even when multiple repositories are connected, content is often ingested without consistent enrichment or normalization.

The result is predictable as the AI retrieves information based on text similarity rather than structured relationships. It misses context that exists across systems and treats loosely related content as equivalent. From a user perspective, this feels like inconsistency. From an architectural perspective, it's fragmentation.

# Hallucinations are a Knowledge Foundation Issue

Hallucinations are often described as a flaw in large language models. In practice, they frequently stem from weak grounding. When AI cannot find clear, structured evidence to support a response, it fills gaps with plausible language. That's what it's trained to do.

Let's consider a complex query processed through traditional retrieval, which returns dozens of loosely related documents. The system cannot clearly explain why those documents were selected. In contrast, knowledge graph-guided retrieval narrows the results and provides a visible relationship path between entities, clauses, and governing laws.

The second scenario reduces ambiguity. The AI works with explicit connections rather than guessing based on keyword overlap. So, if your AI is hallucinating, it is often because your data infrastructure lacks structure.

## The Compliance and Security Pressure

For regulated industries, the challenge extends beyond accuracy. AI must respect system and document-level permissions while supporting audit defense and regulatory explainability, all while avoiding exposure of sensitive information.

Many enterprises rely on connectivity layers to index data across systems. Connectivity is necessary, but it is only one step. And many organizations assume that secure connectivity equals a knowledge graph.

Connectivity ensures access, but does not ensure meaning.

Without enrichment and relationship mapping, AI retrieves chunks of content without full awareness of context. That creates risk. An answer may be technically sourced from a permitted document but interpreted outside of its intended scope. In legal, manufacturing, financial services, and life sciences environments, this is not a minor issue.

Existing security permissions, structured metadata, and semantic relationships must work together. Otherwise, AI becomes a liability rather than an asset.

**Existing security permissions, structured metadata, and semantic relationships must work together.**

**Otherwise, AI becomes a liability rather than an asset.**


# The Cost of Inaction

When AI underperforms, the cost shows up in several ways. These patterns are common. But don't think that these reflect a lack of effort; instead, they reveal a gap between AI ambition and data readiness.



## The results aren't there

First, productivity gains do not materialize. Employees become frustrated spending time fact-checking responses, rather than acting on them.



## Projects start to bottleneck

Second, parallel projects pile up. Teams experiment with different tools or models, hoping for better results. Costs increase while outcomes remain inconsistent.



## Adoption stalls, the project fails

Third, leadership confidence quickly declines. AI implementation moves from a strategic priority to a controlled experiment. Momentum stalls as trust in outcomes erodes.



## A Pattern We See Repeated



### STEP 1

Fragmented information across dozens of systems



### STEP 2

Unstructured content without consistent taxonomy



### STEP 3

Inconsistent metadata applied to existing content



### THE RESULT

AI struggles.

The model may be advanced and prompts may be well written. But the answers depend on what the system can reliably retrieve and interpret.



### STEP 4

AI ingesting content not prepared for machine interpretation




### STEP 5

Security concerns slowing deployment

## AI Reflects Your Knowledge Maturity

Think of AI as a mirror reflecting the state of your data infrastructure. If your content is siloed, the output feels incomplete. If your metadata is inconsistent, the responses won't be precise (or accurate). If relationships between entities are implicit rather than explicit, complex queries break down.

This is why knowledge graphs matter as they address structural issues rather than surface symptoms.



# Knowledge Graphs Explained

## From Theory to Business Value

The term “knowledge graph” is used frequently in conversations about AI. At the same time, it is often misunderstood.

Many organizations believe they already have a knowledge graph because they use a platform that includes the word “graph.” To dodge the market confusion and move forward, it helps to reset the definition and also understand what it is not.

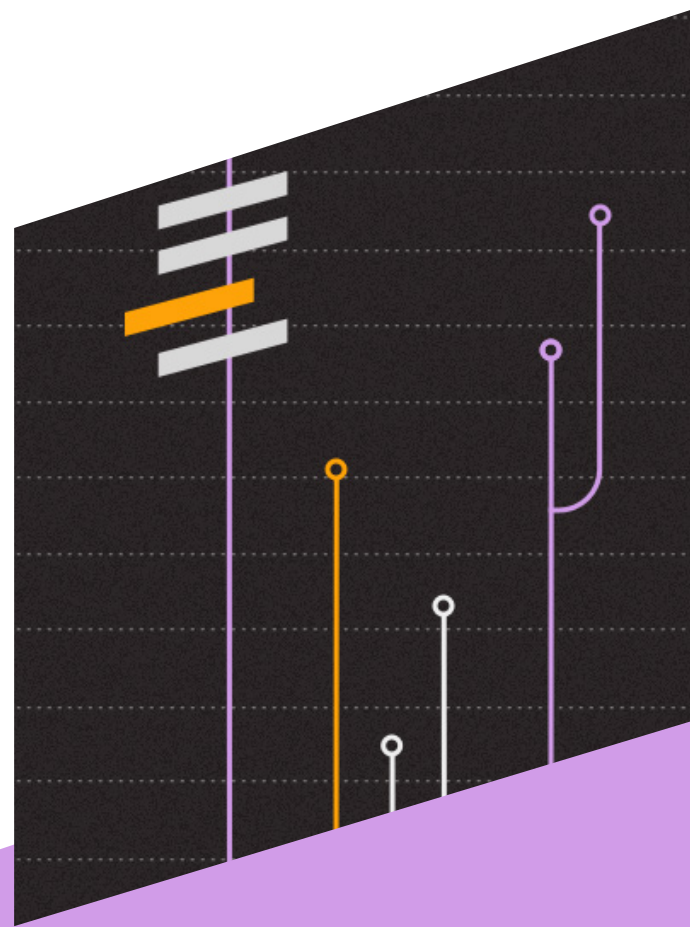
A knowledge graph is not simply a unified API. It is not a connector framework. It is not a search index.

A knowledge graph is a structured representation of the key entities in your organization and the relationships between them. It goes beyond access to capture meaning, with a structure that enables AI reasoning.

## The Four Building Blocks

At a practical level, every knowledge graph includes four core components.

1. Entities
2. Relationships
3. Attributes
4. Schema & Ontology



# Breaking Down Knowledge Graphs

## Four Building Blocks

### 1. Entities

These are the “things” that matter to your business. Examples: clients, contracts, products, subsidiaries, clinical trials, policies, employees, and projects.

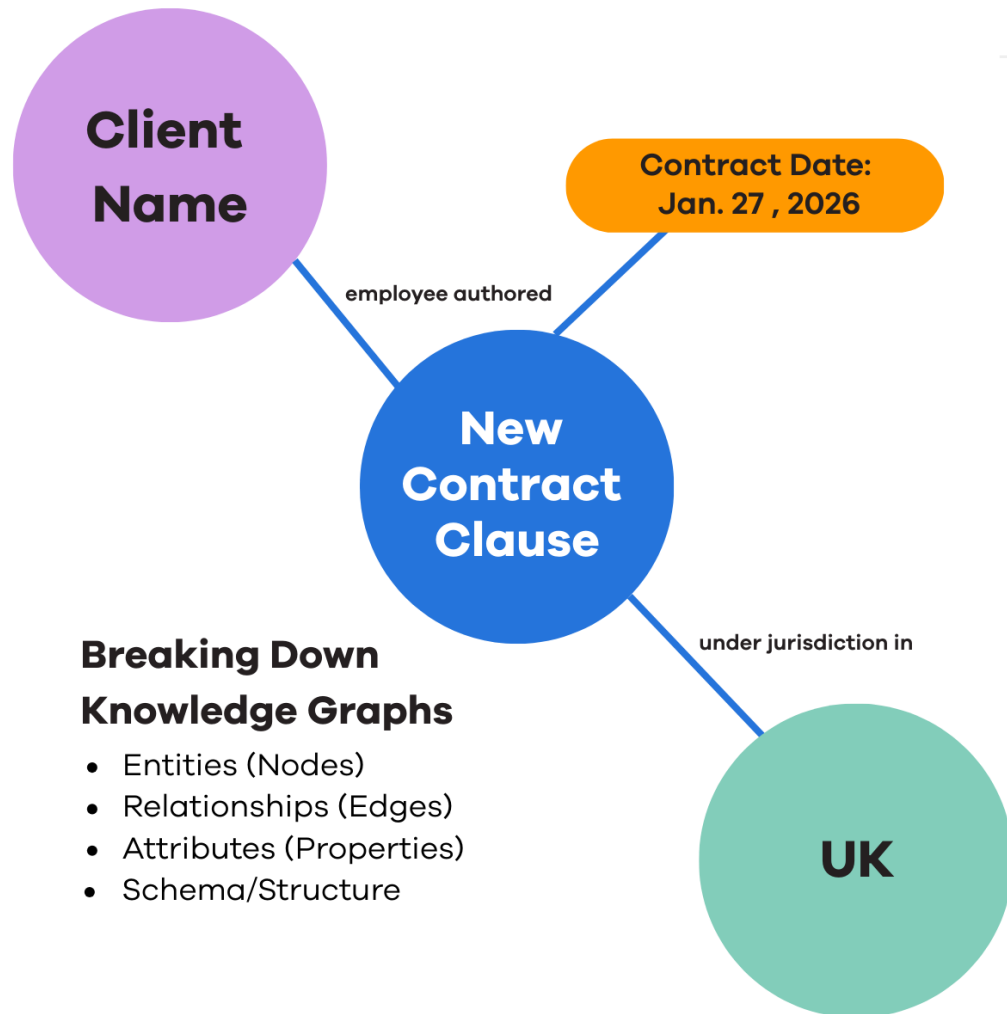
Instead of treating documents as isolated files, a knowledge graph identifies the entities inside them. A merger agreement PDF contains parties, clauses, governing jurisdictions, and impacted subsidiaries.

### 2. Relationships

Relationships describe how entities connect:

- A contract contains a clause
- A clause is governed by a jurisdiction
- A subsidiary reports to a parent company
- An employee authored a policy

These connections allow the system to answer complex, multi-part questions with precision.



## Breaking Down Knowledge Graphs

- Entities (Nodes)
- Relationships (Edges)
- Attributes (Properties)
- Schema/Structure

### 3. Attributes

Attributes enrich each entity with the descriptive context AI needs to interpret and prioritize information accurately.

### 4. Schema & Ontology

A schema defines what types of entities exist and how they are allowed to relate. It establishes consistency. Without a schema, every system defines metadata differently, and cross-system reasoning breaks down.

Together, these components transform unstructured content into structured knowledge.

# How Knowledge Graphs Differ from Traditional Databases

## Enabling Connected Intelligence


It's common to ask whether a knowledge graph is simply another database. The difference is in flexibility and context.

Traditional databases store structured records in predefined tables. They work well for transactional systems. Knowledge graphs are designed to represent complex, evolving relationships across diverse data sources.

In an enterprise setting, information rarely lives in one place. A contract may sit in a document management system. Financial data may be stashed in an ERP system. Project details may live in Jira. A knowledge graph connects these sources conceptually, not just technically. This matters for AI.

When generative AI systems rely on basic retrieval methods, they often search for text similarity across indexed chunks. The process works, but it lacks deeper awareness of relationships. Traditional retrieval may return dozens of loosely related documents.

A knowledge graph-guided approach retrieves content based on explicit relationships between entities. The result set is narrower and more relevant, while the reasoning path is visible and the output is easier to trust.



## Why Connectivity Alone Is Not Enough


Connectivity is foundational. You cannot build a knowledge graph without access to source systems. Many enterprises have already invested in connectivity layers, and those investments are important. But connectivity alone does not create understanding.

Many organizations assume they have solved the problem because they can connect systems securely. What is missing is enrichment and intelligent delivery.

Without enrichment:

- Documents remain unclassified or inconsistently tagged
- Entities are not extracted or normalized
- Relationships remain implicit in free text
- AI retrieves content based on proximity rather than meaning

The result is AI output that feels generic or inconsistent.




# The AI Accuracy Multiplier

Large language models generate text. They do not store enterprise truth; they depend on retrieved context. When retrieval is guided by a knowledge graph, users see constant improvements.



## Precision

The system retrieves information tied to specific entities and relationships, rather than broad keyword matches.



## Traceability

The path between entities can be traced. If an answer references a clause, the system can show how that clause connects to a specific agreement and subsidiary.





## Consistency

Structured relationships reduce variability in results. Similar questions produce similar evidence sets.



*For enterprise leaders, this translates into measurable business outcomes:*

- | Higher user trust
- | Reduced manual verification time
- | Improved compliance posture
- | Stronger adoption of AI tools



## A Business Perspective

From a business standpoint, a knowledge graph is part of your knowledge infrastructure. It supports multiple use cases:

- | Enterprise search
- | Generative AI copilots
- | Agent-based automation
- | Expert discovery
- | Compliance audits

In regulated industries, where security and traceability are critical, structured information is especially important. AI tools that are ingesting unorganized content lead to poor output and reduced productivity. A knowledge graph directly addresses this issue by introducing structure and consistency.



## A Thought Exercise


Consider a core business question. For example:

“Which merger agreements executed between 2019 and 2021 include RSU acceleration clauses governed by Delaware law and impact European subsidiaries?”

If your current AI system struggles with that query, the issue is unlikely to be the language model alone. The challenge is whether your infrastructure understands:

- | What a merger agreement is
- | What constitutes an RSU acceleration clause
- | Which jurisdiction governs a clause
- | Which subsidiaries are European
- | How those elements connect

A knowledge graph encodes that understanding.



# The Business Case for Knowledge Graphs

## Analyzing the Investment

By this point, it's easier to recognize that connectivity alone does not create understanding and simple retrieval does not create trust. The next question is the one that matters most to executive leadership: Is this worth the investment?


For many CIOs and CTOs, AI initiatives already carry significant budget commitments. There is pressure to show progress and scrutiny around risk. Additional infrastructure must be justified in business terms, not technical ones. It's important to note that knowledge graphs protect and extend existing AI investments rather than introducing incremental infrastructure cost.

Rather than mere academic exercise, a knowledge graph is a practical response to three measurable issues: productivity loss, risk exposure, and underperforming AI investments.

## Productivity and Time Savings

Across enterprises, employees spend a meaningful portion of their day searching for information across fragmented systems. According to surveys, 56 percent of employees report never finding the information they are looking for without a strong internal search experience.

When AI tools are layered on top of poorly structured knowledge, the problem shifts rather than disappears. Instead of searching manually, employees validate AI responses. They check sources and cross-reference documents, but the friction remains.



**The Big Question:**  
*Is this worth the investment?*

A knowledge graph improves precision at the retrieval layer. Fewer irrelevant results, clear relationship paths, and better contextual grounding. Over time, this reduces verification time and increases confidence.

There is also a multiplier effect. When AI outputs are trusted, adoption increases. When adoption increases, productivity gains compound.

## Reducing Risk and Strengthening Compliance

In industries such as legal and consulting services, financial services, manufacturing and life sciences, AI risk is not theoretical. Sensitive information must remain permission-aware and outputs must be traceable. Audit teams must be able to understand how conclusions were formed.

AI projects often fail due to poor data quality, weak connectivity, and security concerns. Connectivity is necessary, but without structured enrichment and relationship mapping, AI can misinterpret context even if security trimming is in place.

A knowledge graph strengthens governance in several ways:

- Document-level security inheritance remains intact across systems
- Entities and relationships are explicitly defined, reducing ambiguity
- Reasoning paths can be surfaced to explain why a document was selected
- Taxonomies create consistency across repositories

This reduces the likelihood of misleading or contextually incorrect responses. It also makes AI easier to defend in front of compliance teams and regulators.


## Maximizing Existing AI Investments

As noted, many enterprises have already invested in Microsoft Copilot, Azure OpenAI, Amazon OpenSearch, Amazon Q Business, or similar platforms. The expectation is that these tools will drive measurable ROI.

This is a common misconception. Organizations believe they have a knowledge graph because they have connectivity. They expect strong results, and when output quality disappoints, confidence in the broader AI strategy declines.

A knowledge graph protects those investments. It improves the quality of the input layer, increasing retrieval precision and supports more advanced use cases (such as agent-driven workflows).

The alternative is continued experimentation without structural improvement.



# Industry-Specific Value

The business case varies by industry, but the pattern is consistent.



## Legal

Legal teams manage contracts, clauses, case law, and regulatory materials across multiple repositories. Complex queries often require understanding how clauses relate to governing jurisdictions and specific entities. A knowledge graph enables precise clause-level retrieval and traceability.



## Life Sciences

Clinical trial data, regulatory submissions, product documentation, and research reports are highly interconnected. Security requirements are strict. Relationship mapping across studies, compounds, and geographies improves both AI output and compliance confidence.



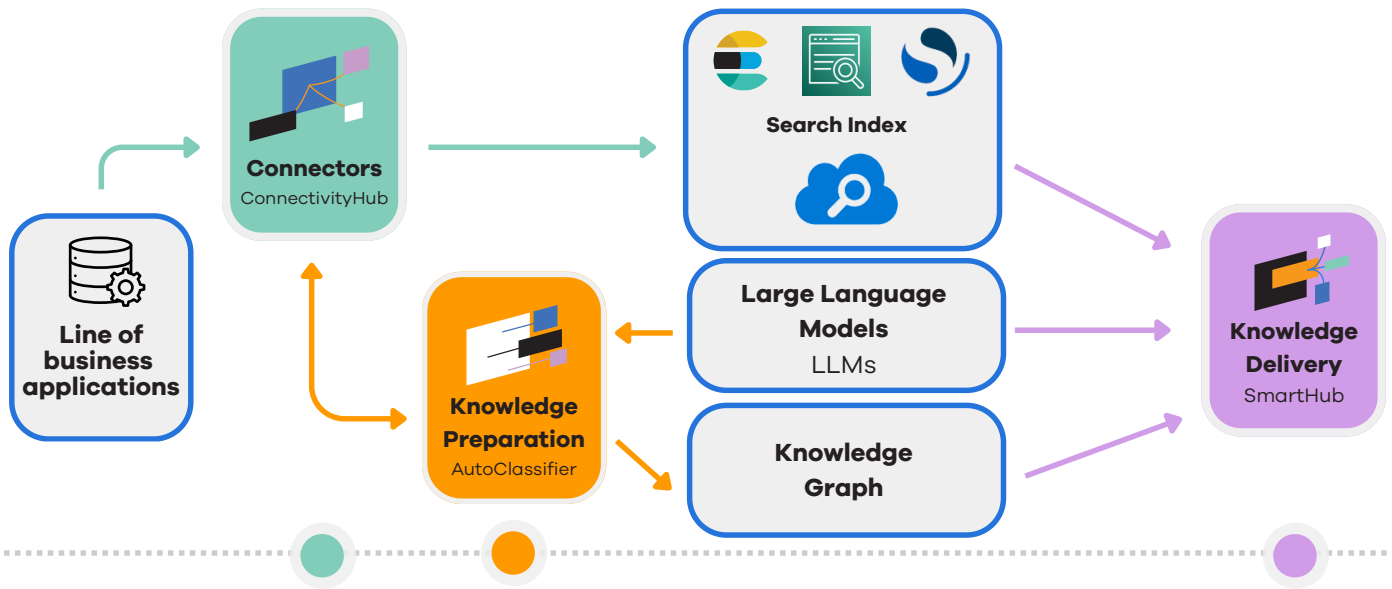
## Professional Services

Client engagements generate large volumes of project documentation. Poorly structured information often leads to hallucinations and inconsistent chatbot responses. Enrichment and structured connectivity restored trust and improved productivity.



## Financial Services

Risk models, regulatory guidance, and client data must be tightly controlled. AI systems must operate within well-defined boundaries. A knowledge graph provides the structural layer needed for consistent, secure retrieval.



*This architecture shows how shared data preparation and delivery layers support consistent, scalable AI experiences across multiple applications.*

## Competitive Advantage

Early adopters of structured data infrastructure are seeing a clear advantage. They move from pilot projects to production systems faster. Internal friction is reduced between IT, compliance, and business teams. They create reusable assets that support multiple AI initiatives.

In contrast, organizations that rely solely on connectivity and raw retrieval often cycle through tools without addressing the root cause.

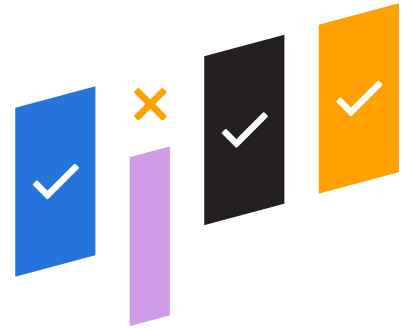
## Frame Your Investment

From a financial perspective, a knowledge graph should be viewed as core infrastructure. It supports enterprise search, strengthens AI copilots, and enables agent-based automation, all while improving findability and contextualization across systems.

The organizations that treat their information as strategic infrastructure will see faster time to value from AI. Those that treat it as an afterthought will continue to troubleshoot symptoms rather than solving the structural issue.

# Readiness Checklist

Before investing in a knowledge graph initiative, it's useful to assess your current state. Use this checklist as a starting point for cross-functional conversations between IT, compliance, and business leaders to align on readiness, risk, and priorities for scaling enterprise AI.



Ask the following:




## Question 1:

Are at least 80 percent of your critical systems securely connected?



## Question 2:

Do you have consistent taxonomies across repositories?



## Question 3:

Are entities (clients, products, and contracts) defined and extracted?



## Question 4:

Can your AI system explain why a specific document was retrieved?



## Question 5:

Do similar queries produce consistent evidence sets?



## Question 6:

Does your compliance team have confidence in AI traceability?

Take note of your 'no' answers to determine how to prioritize next steps.

Answered 'yes' to all six questions? You're in the minority who have matured beyond experimentation and are seeing value from your AI Agents and Assistants. Congratulations!



## Protecting AI investments across the enterprise

Upland BA Insight's AI Enablement Platform allows enterprises to access, prepare, and expose the knowledge needed to power AI projects. BA Insight is AI agnostic; using vector search, data chunking, ML, NLP, and security trimming to index, enrich, and display knowledge from all connected sources to supercharge enterprise AI.



**Connect with us.**

## Request a demo.

See how BA Insight connects to all your enterprise data sources to power whatever AI engine you choose to build.

[Learn more](#)