




Documents Are the Overlooked Layer of Enterprise AI Strategy



Executive Summary

As AI budgets keep flowing to models and agents, the documents that hold a company's knowledge remain overlooked. And while 62%¹ of enterprises run AI for enterprise search or knowledge management in production or pilot, that access alone does not make an answer trustworthy. An AI answer needs context: where it came from, how current it is, and whether it matches the original text. More documents without that context just mean more work to check.

At the same time, buyers want generative and agentic AI capabilities when they shop for new tools. But companies building AI in-house point to a different problem than compute or talent: the quality of their data, and whether an answer can be traced back to its source.

Closing that gap is only the start. The bigger opportunity is reasoning across whole document sets with every answer traceable to its source, and then carrying that intelligence into the content teams create and the systems where they already work.

AI Strategy Rarely Reaches the Documents It Depends On

AI budgets generally go to models, agents, and the infrastructure that runs them. The documents that those systems are supposed to read rarely get the same attention, even though they hold the contracts, playbooks, specs, and records that show how a company really works.

A model is only as good as what it can read. A contract sitting in a shared drive, a policy split across departments, or a spec locked in an old file format stays invisible to AI until someone connects it. Most AI strategies stop short of that step.

Treating documents like a filing cabinet, instead of a live input to decisions, is the default, and an expensive one. But closing the gap does not mean buying another AI model; instead, it means building the layer that lets AI reach the knowledge the company already owns, and trust what it finds there.

¹ ETR, AI Product Series, March 2026 (n=600).



More Documents Are Not the Same as Trusted Answers

Traditional enterprise search was built to find relevant documents. But finding a document isn't the same as giving an AI system the context it requires to reason over it. Retrieval pipelines usually index and return content as text chunks, so they may not keep or surface the context that makes a document trustworthy and useful. Context could be the inclusion of line items in a financial statement, the numbering of clauses in a contract, or the callouts in an engineering specification. When document retrieval lacks context, an AI model may locate the right document, but it will still fail to understand what makes its answer reliable.

Search alone does not solve the problem. Even after a system finds the right document, someone still has to check that the answer is current, complete, and from the right source. Teams make up for this by searching multiple systems, comparing versions, and checking a result by hand before they trust it enough to act.

Fragmented, unverified documents create three problems at once:



Cost:

Analysts and knowledge workers spend hours a week finding information, then double-checking it before they trust it.



Risk:

Decisions made on outdated or unverifiable documents cause problems later, like a mispriced contract or a compliance gap.



Unrealized Value:

Insight without clear context is too risky to act on, so it never reaches the people who need it.


Unstructured content is not dead weight sitting in storage. It is a cost, a risk, and a source of value that the company has already paid to create but cannot yet trust enough to use.

Adoption Is Outpacing Trust in the Data Behind It

Agentic AI is software that acts on a user's behalf instead of just answering questions. And AI for enterprise search or knowledge management now runs in production or is being piloted at 62% of enterprises, up from 57% a year earlier. Buyers are backing that shift with their purchase criteria: 43% rank generative AI capability as their top priority, and 38% rank agentic AI at the same level.¹

That said, trust has not kept up. Among companies building their own AI, three barriers get in the way (see Figure 1).

Figure 1. Data Quality, Privacy, and Availability Are Barriers to Scalable Enterprise AI Adoption



Source: ETR, AI Product Series, March 2026 (n=444).¹

Accuracy concerns have not gone away either, according to Futurum Research's 1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report. Nearly one in four buyers, 24.9%², still name accuracy and hallucination as their top worry about generative AI. It is far easier to summarize one document than to prove that an answer comes from the correct document across a large, often disconnected project repository.

Even more challenging is enabling AI to reason and take action across dozens or hundreds of related documents, such as a contract library, a set of specs, years of board decks, while keeping every answer traceable back to its source. An answer is also rarely the finish line. Most workflows need one more step: a proposal, a client-ready deck, or a finding sent to the system where a decision gets made.

² 1H 2026 Data Intelligence, Analytics, and Infrastructure Decision Maker Survey Report (Futurum Intelligence, March 2026).



How Adobe Helps Enterprises Turn Documents Into Trusted Intelligence

From Summarizing One File to Reasoning Across Many

The first wave of AI document tools focused on summarizing a single file. Although that is a useful capability, most enterprise knowledge, as well as data required to make a decision, spans more than one document. A pricing decision may rest on a master agreement, three amendments, and a rate card. A compliance review may draw on internal policies, audit findings, and regulatory guidance. The real value comes from reasoning across the whole set and showing where each part of the answer came from.

Adobe Acrobat Studio is built for that shift. It brings together document consolidation, AI-powered comparison and analysis, and content creation. A user can pull related files into one place, ask a question that spans all of them, and get an answer tied to the passage it came from. That way, they can check it before acting on it, which a summary without sources doesn't allow.

PDF Spaces brings the same approach to everyday teamwork. A team can link the documents for a specific deal, project, or review and ask questions inside that space. Answers stay focused on the right material, and nobody has to search the whole company each time a related question comes up.

Traceability like this depends on curating the document set on an ongoing basis. A one-time upload won't sustain it. The quality of an answer depends on the quality and currency of the documents behind it.

Traceable Document Intelligence Matters More as AI Takes Action

Agentic AI raises the stakes around document management, knowledge retrieval, and insight generation. Security and data privacy are the top concerns about agentic AI at every stage of deployment, and concerns about transparency more than double, from 7% to 15%, as organizations move to multi-agent orchestration.³ The bottleneck is no longer producing AI output. It is the ability to govern, integrate, and verify that output: only 15.8% of organizations report sustained ROI at scale from AI built in-house, and 12.6% from AI they buy.⁴

Much of that risk sits with the documents an agent draws on, not with any single platform. Adobe's role is to help organizations work from governed, curated knowledge, so any system acting on it can trace back to an official source of information. A platform earns trust through the reliability of the documents underneath it.

³ The Futurum Group, 1H 2026 AI Platforms Decision Maker Survey Report, March 2026 (n=820).

⁴ The Futurum Group, 2H 2026 CIO & Technology Buyers Global Enterprise Decision Maker Survey Report, September 2026 (n=1,636).

Knowing where an answer came from is also what lets a team hand it to an autonomous agent with confidence. Knowledge Base, part of Acrobat Studio, helps companies access organizational knowledge that is already being curated in an accurate and verifiable way. Traceability also lets intelligence travel. An answer that carries its source can move into other systems without losing what makes it credible.

Turning Intelligence Into New Work and Action


Once document knowledge can be reasoned over and traced, the next question is where it goes. Too often the answer is another AI interface: a chat window or assistant panel where insight sits until someone copies it into the system that needs it. Each handoff adds delay and a chance for error, and it breaks the line back to the source document.

Document intelligence is worth more when trusted insight flows straight into the systems and workflows where decisions are made and actions are taken. A verified contract term is most useful in the system that governs the contract. A validated price belongs in the system that quotes it. A confirmed spec matters most in the workflow that builds or ships against it. When insight arrives in those systems with its source still attached, teams can act on it without rebuilding the trail of evidence.

Acrobat Studio supports two ways to put that intelligence to work. The first is creation. Users can turn cross-document analysis into new content, such as converting technical specs into a client-ready proposal or a set of board materials into a briefing deck, while staying grounded in the same sources.

The second is connected action. Analyzer, part of Acrobat Studio, connects to the enterprise resource planning (ERP), customer relationship management (CRM), and contract lifecycle management (CLM) systems a company already uses. An extracted clause, price, or spec can move directly into the system where a team works, so trusted document intelligence becomes part of the operational record.

Figure 2. Turning Enterprise Document Knowledge Into Trusted Intelligence and Action



Trusted Document Intelligence Completes the Loop

Consolidation, traceable reasoning, creation, and system integration are designed to work together as one system. Together they give companies a way to work from curated and governed documents, reason across them with a clear line back to source, and carry that intelligence into the decisions, content, and workflows teams depend on.

As model capabilities mature, enterprise AI purchase decisions will increasingly depend not only on which model to use, but on the quality and trustworthiness of the document knowledge behind it. Companies that treat their document library as governed, connected, and traceable can put AI to work on decisions already in front of them: checking supplier invoices against negotiated terms, flagging renewal risks buried in customer agreements, or pulling verified answers from past submissions to respond to a new RFP. As such, enterprises can rapidly start unlocking value from their existing document knowledge.

Trust is the foundation for all of it and makes everything that follows possible. The bigger opportunity combines trusted document intelligence, cross-source reasoning, creation of new work, and connected action within the systems teams already rely on. The aim reaches beyond making documents trustworthy enough for AI. It means turning enterprise knowledge into intelligence that shapes what people decide, create, and do next, across the workflows and systems where work already happens. Documents are the overlooked layer of enterprise AI strategy, and Adobe is helping turn that layer from a static repository into a working source of trusted intelligence across the enterprise.

Important Information About This Report

AUTHORS

Keith Kirkpatrick

Vice President & Research Director, Enterprise
Software & Digital Workflows | The Futurum Group

PUBLISHER

Futurum Research

INQUIRIES

Contact us if you would like to discuss this report and
The Futurum Group will respond promptly.

CITATIONS

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

LICENSING

This document, including any supporting materials, is
owned by The Futurum Group. This publication may not be
reproduced, distributed, or shared in any form without the
prior written permission of The Futurum Group.

DISCLOSURES

The Futurum Group provides research, analysis, advising,
and consulting to many high-tech companies, including those
mentioned in this paper. No employees at the firm hold any
equity positions with any companies cited in this document.



ABOUT ADOBE

Adobe is a global leader in digital media and digital marketing solutions, empowering individuals and organizations to create, manage, and engage with compelling content. The company's comprehensive ecosystem bridges document management, artificial intelligence, and creative tools to reimagine modern workflows. Through flagship offerings like Adobe Acrobat Studio, Adobe enables users to move fluidly from gathering information to analyzing insights and taking connected action across enterprise systems. By focusing on data security, content traceability, and intelligent document platforms, Adobe helps businesses unlock the full value of their institutional knowledge. Today, Adobe continues to pioneer the future of digital work, creativity, and customer experiences worldwide.



ABOUT THE FUTURUM GROUP

The Futurum Group is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.



CONTACT INFORMATION: The Futurum Group LLC | [futurumgroup.com](https://www.futurumgroup.com) | (833) 722-5337