



Websites Supercharged: Content Storage Transformed into an Answer Engine

The Hidden Costs of Elusive Content

People no longer want to “search” your site—they expect to ask a question and get a clear, confident answer that feels like a ChatGPT-style response, not a list of blue links.

For marketing leaders, that expectation now defines what “good” looks like for web and digital experiences. When visitors cannot quickly get the answers they need, decisions slow down, campaigns stall, customers get frustrated and revenue opportunities slip away.

Your teams feel the pain too, spending time tuning search rules, maintaining systems and recreating content that already exists but cannot be found when needed.



“40% of B2B buyers abandon their search if they cannot find what they are looking for.”

- The State of B2B Commerce: How To Win in the Digital Age, Forrester

The irony is that most organizations already have everything they need: webpages, PDFs, videos, reports, presentations and more. The real problem is that traditional content management systems (CMS) and site search are built on rigid metadata and keyword matching, returning irrelevant results and hiding your best content from the buyers and internal teams who need it.

We will show you how pairing your CMS and generative search capability with the Progress® Agentic RAG solution can turn that content into a trusted answer engine, while keeping potential buyers and customers engaged, empowering sales and support and delivering AI search that works within weeks.

Generative Search: Trusted, Grounded Answers, Not Long Lists of Links

Simply put, generative search turns your website into the business-driving tool it should be. Most organizations already have a rich content ecosystem that includes web pages, PDFs, videos, reports and presentations. The problem is that most content is underutilized because it's unstructured and too hard to surface when needed.

With the rise of AI tools, web developers and CMOs alike are finding they can reframe the CMS to be a core knowledge engine. And when paired with generative search using agentic retrieval-augmented generation (RAG) technology, websites can answer real questions from both employees and customers in mere seconds.

Traditional CMS and site search rely on keyword matching and rigid metadata, forcing users to waste time guessing exact terms and digging through irrelevant results. However, generative search with agentic RAG technology is tailored to the user's intent, instantly producing relevant, in-context answers.

Generative Search Answers questions in natural language, grounded in your content—not generic web data.

The Power of Agentic RAG

Generative search is the intelligent interface layer that sits on top of your CMS and other content sources, replacing link-heavy results pages with direct, grounded answers. Users ask natural-language questions and get the exact passage, snippet or multimedia timestamp they need. This dramatically reduces scrolling, guessing and the tiresome back-and-forth so users can search for and quickly find what they want. People want fast, contextual and reliable answers, and that's exactly what the Progress Agentic RAG solution delivers.

It does this by ingesting content from your CMS and beyond, including documents, text, video, audio, images and more. It then structures it all so it can be searched and reasoned over with AI. With modular pipelines, 30+ interchangeable LLMs and embeddable components, Progress Agentic RAG search can be deployed on your website in weeks, without custom infrastructure or deep AI expertise required. Plus, its LLM-agnostic approach will keep your website future-focused from day one.

Meet the Progress Agentic RAG Solution: The Intelligence Layer That Powers Generative CMS Search

Generative search is now simple to add to your CMS with the Progress Agentic RAG solution. This revolutionary software ingests content from your CMS and beyond and structures it so AI can search, retrieve and reason over it.

Additionally, the Progress Agentic RAG AI Search feature understands user intent across multiple languages and can return answers in the language of the user's prompt, regardless of the foundational language of the underlying content. This enables organizations to offer consistent, high-quality search experiences to worldwide audiences without duplicating content or maintaining separate search configurations per language.

It is an LLM-agnostic tool, so you can swap models as the landscape evolves or as your project needs change. Instead of multiyear AI integration projects, organizations can connect their CMS and a few high-value repositories, configure retrieval strategies and embed search components into existing sites or portals quickly. This allows for a pragmatic rollout path: start with one audience and a subset of content, measure adoption and answer quality, then expand sources and use cases as value is proven.

Agentic RAG

The AI layer that retrieves, reasons over and assembles grounded answers derived from all of your content sources.



Web Workflow Examples for Agentic RAG

Here are some workflow examples for how agentic RAG could be applied in various web search scenarios:

Workflow	1: Answer-Level Search in Your Web Content	2: Turning Unstructured Files into Usable Intelligence	3: Video, Audio and Visual Content at Your Fingertips
	Website content becomes a live knowledge base instead of a static library of pages. Marketers, product owners and visitors can ask questions in plain language and get a synthesized answer, backed by citations to specific sections of articles, landing pages and support content inside the CMS.	Agentic RAG brings PDFs, slide decks, reports and notes out of shared drives and combines it with data living in systems of record into an intelligent retrieval pipeline that understands context, not just keywords. Teams can ask questions and instantly receive concise answers with links to the original documents and the exact passages used.	The Progress Agentic RAG solution indexes video and audio assets, enabling search—not only across transcripts, but also at the timestamp level for specific segments and quotes. Readers can jump directly to the exact moment in a webinar, demo or training video where a topic is covered. Diagrams, infographics and images can also be surfaced as part of a multi-modal answer.
Example Scenario	A prospect lands on your pricing page with questions about usage-based tiers.	Product Marketing seeks to refine messaging that directly addresses prospect concerns.	Your content team needs to develop a customer compilation video to support your next product launch.
Sample Question	“How does your usage-based pricing compare across regions?”	“What were the main customer objections in last year’s RFPs?”	“Provide me with a script for social media videos where customers are speaking positively about their experience with [product X]. Please include the video source and timestamp.”
Result	Generative search synthesizes content from pricing pages, FAQs and documentation into a single, grounded answer with citations.	Agentic RAG would retrieve information from CRM opportunity notes, proposal management systems, document libraries and more to deliver a succinct, detailed response.	Agentic RAG retrieves relevant video testimonials, identifies positive segments and generates a concise script with source links and timestamps included.

Beyond Marketing: Cross-Functional Generative Workflows

Because agentic RAG is modular and LLM-agnostic, the same platform that powers generative search on your website can be extended to sales enablement, customer support, product documentation and more. And if you think this is just a “nice-to-have,” think again. According to the [Harvard Business Review](#), “...the financial impact of poor knowledge management is staggering,” and siloed knowledge can slow cross-functional collaboration by up to 30%. The research also cites that employees spend an average of 21% of their work time searching for information and another 14% recreating what



“Inefficiency directly costs a business an average of 25% of its annual revenue.”

- Value of Enterprise Intelligence 2025 Report, Bloomfire

they cannot find, which results in lost time that directly erodes productivity and revenue.

The Future Is Now: Generative CMS Search Powered by Agentic RAG

Generative search powered by agentic RAG technology enables organizations to move toward a generative CMS that doesn't just store and publish content, but actively participates in delivering personalized, answer-first experiences. And that's just the beginning.

When you combine the Progress Agentic RAG solution with your CMS, your digital experiences can evolve to be adaptive and dynamic to meet customers' demands. It produces context-aware and compliant experiences that meet user expectations while maintaining trust and governance in the age of intelligent, conversational systems.

By combining enterprise-grade RAG intelligence with dynamic experience generation, marketers will be able to transform siloed content and data into conversational, hyper-personalized

and continuously optimized digital experiences without the complex rules or tech barriers that delay time to value.

If you've been wanting to reimagine your visitor journeys, leveraging generative AI is a powerful way to make them more personalized while still retaining their reliability. A generative CMS with the Progress Agentic RAG solution offers simple setup and the powerful, grounded answers needed to make this a tangible reality

Generative CMS: Your CMS + Agentic RAG + Generative AI experiences that adapt to each visitor's intent and context.

Implementation Checklist: Where Do You Start?

Now that you understand the power of generative search in your CMS, how can you bring it to your organization? This handy checklist is a great place to start:

1. Assemble Your Collaborators: Which people or departments need to be involved in your project?

- Who are the natural stakeholders?
- Do you have the development expertise you need in house?
- Who are the system administrators?

2. Scope Your Project: What are you trying to accomplish on your website?

- What are your UX and QA considerations?
- What does success look like, and how will you measure it?

3. Build Your Knowledge Box: What assets are you including and where are they?

- Have you identified your specific sources and assets?
- Are they current, relevant and accurate?

4. Implement on Your Website First:

- Have you set up your knowledge box?
- Have you established system prompt guardrails and guidelines for responses?
- Have you tested answer quality with sample questions?

5. Consider for Ongoing Optimization and Expansion:

- Who will be responsible for ongoing optimization?
- How can you improve the user experience and responses?
- Has the system identified any content gaps?



Learn More about Generative Search in the Progress Agentic RAG Solution

About Progress Software

[Progress Software](#) (Nasdaq: PRGS) empowers organizations to achieve transformational success in the face of disruptive change. Our software enables our customers to develop, deploy and manage responsible AI-powered applications and personalized digital experiences with agility and ease. Businesses of all sizes get a trusted provider in Progress, with the products, expertise and vision they need to turn AI disruption into a competitive advantage. Millions of developers and technologists at hundreds of thousands of organizations depend on Progress every day. Learn more at www.progress.com

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/03 | 121348377700380

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA 01803, USA
Tel: +1-800-477-6473

- facebook.com/progresssw
- x.com/progresssw
- youtube.com/progresssw
- linkedin.com/company/progress-software
- [progress_sw_](https://instagram.com/progress_sw_)