

# Unlocking Intelligent Knowledge Search for Progress OpenEdge ISVs

How the Progress Agentic RAG Solution Transforms Internal Knowledge into Strategic Advantage

WHITEPAPER



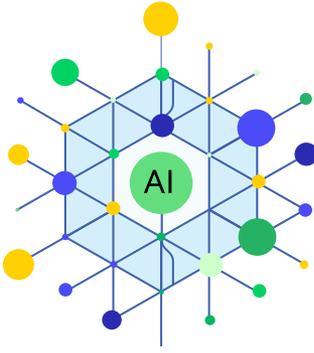
# Executive Summary

Independent Software Vendors (ISVs) building on the Progress® OpenEdge® platform have spent years, or even decades, refining deep business logic, vertical expertise and mission-critical applications. These systems are rich in data, rules and domain intelligence. Yet most OpenEdge-based solutions remain transactional systems of record rather than intelligent systems of reasoning. They store and process transactions exceptionally well, but they do not reason across them.

The rise of agentic retrieval-augmented generation (RAG) introduces a new opportunity as it enables applications to not only retrieve information, but to investigate, synthesize and reason across structured and unstructured data sources. For OpenEdge ISVs, this represents a path to transform internal knowledge search from a fragmented, manual process into an intelligent, context-aware system.

An architecture consisting of both the **Progress® Agentic RAG** solution and the **Progress OpenEdge MCP Server** becomes even more powerful. Business logic remains authoritative and secure, while AI agents gain structured, governed access to the domain intelligence embedded within the OpenEdge platform. The result is not simply an AI feature; it is a strategic evolution—from transactional platform to intelligent reasoning system.





# The Hidden Cost of Fragmented Knowledge

Within most OpenEdge ISV organizations, knowledge is widely distributed. Transactional data lives in the OpenEdge database. Business rules are encoded in Advanced Business Language (ABL) procedures. Customer context may be stored in CRM systems. Policies and contracts exist as PDFs. Support histories are recorded in case systems. Product documentation resides in knowledge bases or shared drives.

When an internal team member needs to answer a complex question—such as why a pricing override occurred, whether a customer is at compliance risk, why margin dropped in a specific region, etc.—the process is rarely straightforward. Multiple systems must be accessed, reports generated, documents searched, subject matter experts (SMEs) consulted and data manually reconciled and interpreted.

Even well-designed search systems do little more than return lists of results. They do not interpret findings, reconcile contradictions, apply business rules or explain outcomes. The intellectual burden remains on the employee. This creates inefficiencies, increases dependency on experienced staff, slows onboarding and introduces inconsistency in decision-making. The cost is not only measured in time, but is measured in cognitive load, risk exposure and lost strategic focus.

## From Retrieval to Reasoning: The Agentic RAG Difference

Traditional RAG improves upon keyword search by retrieving relevant documents and allowing a language model to generate a contextual answer. Agentic RAG advances this further by introducing autonomous reasoning capabilities. The system does not merely retrieve information; it determines what information is required, queries multiple sources, invokes tools, applies logic and iterates until it produces a coherent and grounded response.

This distinction is critical. In an OpenEdge environment, many questions require more than document lookup. They require structured data queries, rule evaluation, historical context and interpretation. However, agentic RAG orchestrates these steps in sequence. It can retrieve a contract, query transactional data, compare current values to historical trends and explain discrepancies in natural language. The process resembles how a skilled analyst works, except it happens within seconds.

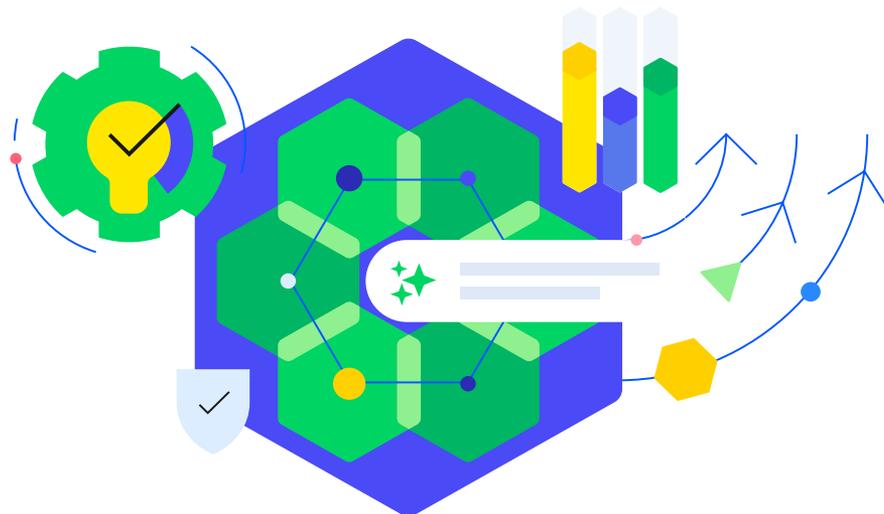
For internal knowledge search, this shift from static retrieval to dynamic investigation is transformative. Employees no longer receive fragments of information, they receive synthesized explanations grounded in system data and documentation.

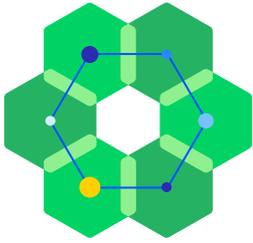
# Elevating OpenEdge Applications from System of Record to System of Intelligence

OpenEdge applications already contain immense domain expertise. Pricing algorithms, eligibility rules, regulatory logic, approval workflows and financial calculations are embedded deeply within ABL code and database structures. This intellectual property differentiates each ISV in its market.

Agentic RAG unlocks that embedded intelligence by allowing AI agents to reason within the domain context of the application. Instead of producing generic responses, the system operates with awareness of the structures and semantics that define the business. When asked why a customer received a discount, the agent can correlate contract terms, historical overrides, approval logs and relevant policies before producing an explanation.

This transforms the OpenEdge environment from a passive repository into an active reasoning partner. The application does not merely hold information—it participates in interpreting it.





# Unifying Structured and Unstructured Knowledge

One of the longstanding challenges in enterprise systems is the divide between structured and unstructured knowledge. Databases excel at storing transactions and numeric values. Documents capture policies, legal language and historical narratives. And traditional tools treat these domains separately.

Agentic RAG bridges that divide—it can retrieve a policy document, extract relevant clauses, query structured data for supporting evidence and merge both into a coherent explanation. For industries served by OpenEdge ISVs, such as manufacturing, healthcare, distribution, financial services and human capital management, this capability is especially valuable. Many decisions depend not only on numbers, but on interpretation of policies and agreements.

By unifying these knowledge domains, agentic RAG reduces fragmentation and increases contextual accuracy.

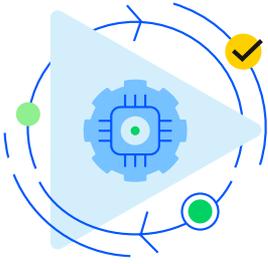


# Productivity and Organizational Impact

The impact of intelligent internal knowledge search is measurable. Knowledge workers frequently spend significant portions of their day searching for information, validating assumptions and confirming interpretations. When this effort is reduced, productivity increases immediately.

Support teams resolve cases faster because they can access historical resolutions and contextual customer data in a single interaction. Sales teams gain rapid visibility into contract terms and pricing exceptions. Finance teams can investigate anomalies without manually cross referencing reports and documents. Operations teams can identify compliance risks through natural-language queries rather than complex reporting tools.

Beyond time savings, the organization benefits from consistency. Answers generated through agentic RAG are grounded in authoritative data sources, reducing variation between individuals. Employees onboard quicker because institutional knowledge becomes accessible via conversation, rather than tribal memory.

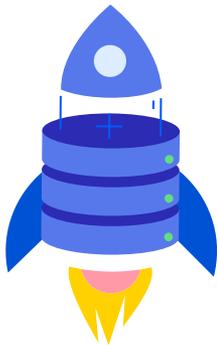


# A Strategic Path for Modernization

Many OpenEdge ISVs face modernization pressures. Customers expect AI capabilities. Competitors advertise intelligent assistants and automated insights. And rewriting a mature OpenEdge application to meet these expectations is costly and risky.

Agentic RAG offers a different path. Rather than replacing the transactional core, ISVs can augment it. The existing application remains the system of record. It also becomes the reasoning layer above it. This layered approach preserves investment, minimizes disruption and accelerates innovation.

Then internal deployment provides a low-risk proving ground. By enabling internal teams with intelligent knowledge search, ISVs can refine governance, security and performance before extending capabilities externally. Once validated, the same architecture can power assistants, AI features and analytics.



# Enhancing the Architecture with the OpenEdge MCP Server

While agentic RAG provides the reasoning engine, the OpenEdge MCP Server strengthens the architecture by supporting AI agents to interact with OpenEdge applications through controlled, governed interfaces. Rather than bypassing application logic and querying raw tables, agents invoke defined services and procedures exposed through the Model Context Protocol (MCP).

This preserves the integrity of the application. Business rules remain centralized and authoritative. Pricing calculations, eligibility validations and compliance checks continue to execute in the same logic trusted by the transactional system. The risk of duplicating or drifting from established rules diminishes.

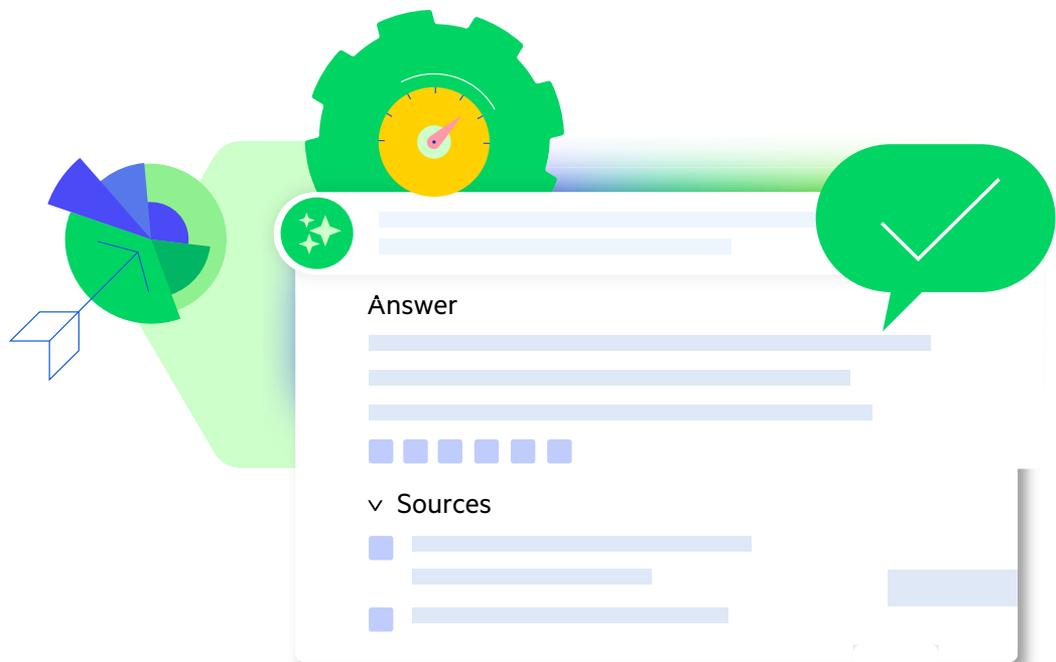
For ISVs, this is crucial. Intellectual property resides in ABL code and domain-specific processes. The MCP Server allows that IP to be leveraged by AI without being exposed or replicated. Governance controls define which procedures are accessible, what parameters are permitted and how actions are audited. As a result, security models remain intact. In practical terms, the combination of agentic RAG technology and MCP Server supports AI-driven reasoning that's not only powerful, but aligned with the application's core logic and security framework.

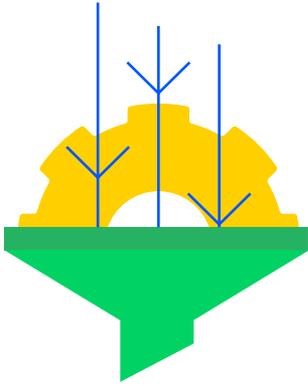
# Long-Term Competitive Differentiation

The integration of the Progress Agentic RAG solution with the OpenEdge MCP Server positions ISVs for long-term differentiation. As markets increasingly demand intelligent, proactive systems, ISVs can deliver capabilities that extend beyond reporting and dashboards. They can offer systems that explain, recommend and act accordingly.

Internally, this translates to greater efficiency and insight; externally, it opens pathways to AI-enabled product tiers, self-service support copilots and intelligent analytics embedded directly within applications.

The strategic advantage lies not in adding a conversational interface, but in transforming accumulated domain intelligence into a living reasoning system.





# Conclusion

Progress OpenEdge applications represent decades of investment, refinement and expertise. Yet without an intelligent reasoning layer, much of that expertise remains locked deep within code and data structures.

The Progress Agentic RAG solution enables OpenEdge ISVs to unlock that intelligence by transforming internal knowledge search into a dynamic, context-aware investigative process. When paired with the MCP Server, this capability is delivered securely, governed by existing business logic and aligned with the system of record.

The result is a clear evolution, transitioning from storing transactions to interpreting them, retrieving data to reasoning across it and moving from system of record to system of intelligence.

For OpenEdge ISVs seeking to modernize without rewriting and innovate without compromising control, agentic RAG represents not merely an enhancement, but a strategic inflection point.



**Learn more about about the Progress Agentic RAG solution today.**

## 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 digital experiences with agility and ease. Customers get a trusted provider in Progress, with the products, expertise and vision they need to succeed. Over 4 million developers and technologists at hundreds of thousands of enterprises depend on Progress. Learn more at [www.progress.com](http://www.progress.com)

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.  
All rights reserved. Rev 2026/03 | RITM0353749

## Worldwide Headquarters

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

 [facebook.com/progresssw](https://facebook.com/progresssw)  
 [twitter.com/progresssw](https://twitter.com/progresssw)  
 [youtube.com/progresssw](https://youtube.com/progresssw)  
 [linkedin.com/company/progress-software](https://linkedin.com/company/progress-software)  
 [progress\\_sw\\_](https://instagram.com/progress_sw_)