



CASE STUDY

Redefining Candidate Matching with Graph RAG and Multi-Agent AI for RightSourceAI



Index

3	—	Executive Summary
4	—	Client Background
5	—	Value Delivered
6	—	The Challenge
6	—	The Solution
7	—	Why Graph RAG Over Traditional RAG?
8	—	Multi-Agent AI: Performance Boost
8	—	Future Innovation Areas



Executive Summary

A concise overview of how RightSourceAI leverages next-gen RAG techniques; specifically Graph RAG and multi-agent architectures; to solve one of the most pressing challenges in recruitment: matching accuracy. With a sharp focus on speed, precision, and cost-efficiency, the platform delivers a powerful tool for engineering teams and recruiters in fast-scaling SaaS environments.

Industry

- Recruitment tech

Initial Go-to-Market Timeline

- 3 months

Tech Stack

- Graph RAG (graph-based retrieval-augmented generation),
- AWS Neptune (graph DB),
- AWS RDS,
- Bedrock (LLM orchestration),
- Multi-agent orchestration,
- LangChain, LangGraph, LangSmith,
- NextJs,
- AWS Fargate,
- SSO via Google Authentication

Team Composition

- 1 Product Manager
- 1 Frontend Engineer
- 1 Backend Engineer
- 1 Data Scientist & LLMOps
- 1 QA Analyst



Client Background

RightSourceAI was developed to address one of the most persistent and costly challenges in technical hiring: slow screening cycles driven by inaccurate candidate matching.

Most traditional applicant tracking and sourcing systems rely on keyword matching or vector similarity. While these approaches can surface large volumes of candidates, they introduce two critical problems:

1. **High false positives** – recruiters and hiring managers must manually screen many irrelevant profiles
2. **High false negatives** – qualified candidates are missed because resumes do not contain exact keywords

These issues compound into longer hiring cycles, higher recruiter costs, and excessive time spent by engineering leaders in interviews that should never have happened.

Originally built to support Athenaworks’ internal recruiting operations, RightSourceAI was designed to scale across both SaaS and enterprise organizations hiring for complex engineering, data, and product roles.

“RightSource uses Graph RAG and multi-agent AI to understand job descriptions and resumes the way a senior recruiter would—semantically, not by keywords. The result is fewer candidates to review, better matches surfaced faster, and significantly less time wasted by recruiters and engineering teams.”

Existing recruiting platforms failed to solve the core bottleneck: **screening efficiency without sacrificing precision.**

- Keyword-based systems surface volume, not relevance
- Vector-based RAG systems improve similarity but still lack contextual reasoning
- None adequately explain why a candidate matches, forcing recruiters to manually validate results

RightSource was built to close this gap by:

- Parsing job descriptions and resumes into **rich relationship graphs**
- Understanding implicit skills, experience depth, and role context
- Reducing both **false positives and false negatives** simultaneously

This approach enables faster, more confident hiring decisions while lowering total cost per hire.



Value Delivered

- **Reduced candidate screening time**

From over 2 minutes to just 30–40 seconds, enabling recruiters and hiring managers to work more efficiently and review more profiles in less time.

- **Improved match precision**

Significantly enhances candidate-job fit, reducing false positives and negatives, and improving interview-to-hire ratios, by using Graph RAG for parsing and matching JD.

- **High scalability and extensibility**

The platform can handle larger talent pools of tens of millions of candidates without degrading performance, thanks to being built on modular, cloud-native microservices infrastructure with multi-agent orchestration.

- **Frictionless early adoption**

Recruiters can access and review matched candidates for free, eliminating barriers to trial and encouraging organic product adoption.

- **Strategic alignment with both business and technical stakeholders**

The platform benefits both recruiting teams (faster screening, better matches) and engineering leads (higher-quality pipelines, less screening burden); accelerating hiring while maintaining technical quality.



The Challenge

Recruiting and engineering teams face two compounding issues:

- **Too many weak matches** from keyword/vector-based systems
- **Delayed discovery of strong candidates**, especially for specialized roles

Even modern AI-powered recruiting tools struggle to deliver context-rich matches at speed. For time-sensitive technical hiring, this leads to:

- Increased recruiter screening hours
- Lost productivity from engineering interview loops
- Slower delivery against product and revenue goals

The Solution

RightSource introduces a new candidate-matching architecture built on:

- **Graph RAG:** Uses graph databases (AWS Neptune) to model relationships between skills, roles, experience, and requirements
- **Multi-agent AI:** Parallelizes resume analysis and matching tasks to deliver near real-time results
- **Cost-aware LLM orchestration:** Selects models via AWS Bedrock to balance accuracy, latency, and inference cost

Together, this enables deep semantic matching at production scale.

Key Capabilities

- Upload or paste a job description
- Instantly return AI-matched candidates from a pool of 25,000+
- Secure, SSO-enabled access (Google login for MVP)
- Optional integration with Athenaworks' Gatekeeper service for curated, pre-interviewed candidates



Why Graph RAG Over Traditional RAG?

Feature	Traditional RAG (Vector DB)	Graph RAG (Graph DB - AWS Neptune)
Data Structure	Flat embeddings	Rich relationship graphs
Retrieval Logic	Similarity-based	Context-aware, connection-based
Accuracy	Often noisy results	Higher semantic relevance
Example Use Case	Keyword match on “DevOps”	Maps job responsibilities to candidate experiences & certifications

By introducing Graph RAG, RightSourceAI tackles one of enterprise AI’s biggest blockers: accuracy at scale.



Multi-Agent AI: Performance Boost

To meet user expectations of “real-time” results, RightSourceAI incorporates multi-agent orchestration:

- A supervisor agent breaks down search tasks
- Multiple specialized agents process candidate sets in parallel
- Responses are aggregated within user-defined time constraints

This mirrors how modern search engines and distributed systems optimize latency; bringing consumer-grade responsiveness to enterprise-grade AI.

Future Innovation Areas

- Enhanced graph-based personalization across candidate journeys
- Integration of user feedback to tune AI match scores
- Expansion of multi-agent orchestration to interview scheduling and post-match analytics

Contact Us

athenaworks.com

