Case Study

brillio

From Scattered Knowledge to a Single Answer

How we helped a global quick-service restaurant brand turn scattered knowledge into natural-language answers on AWS.

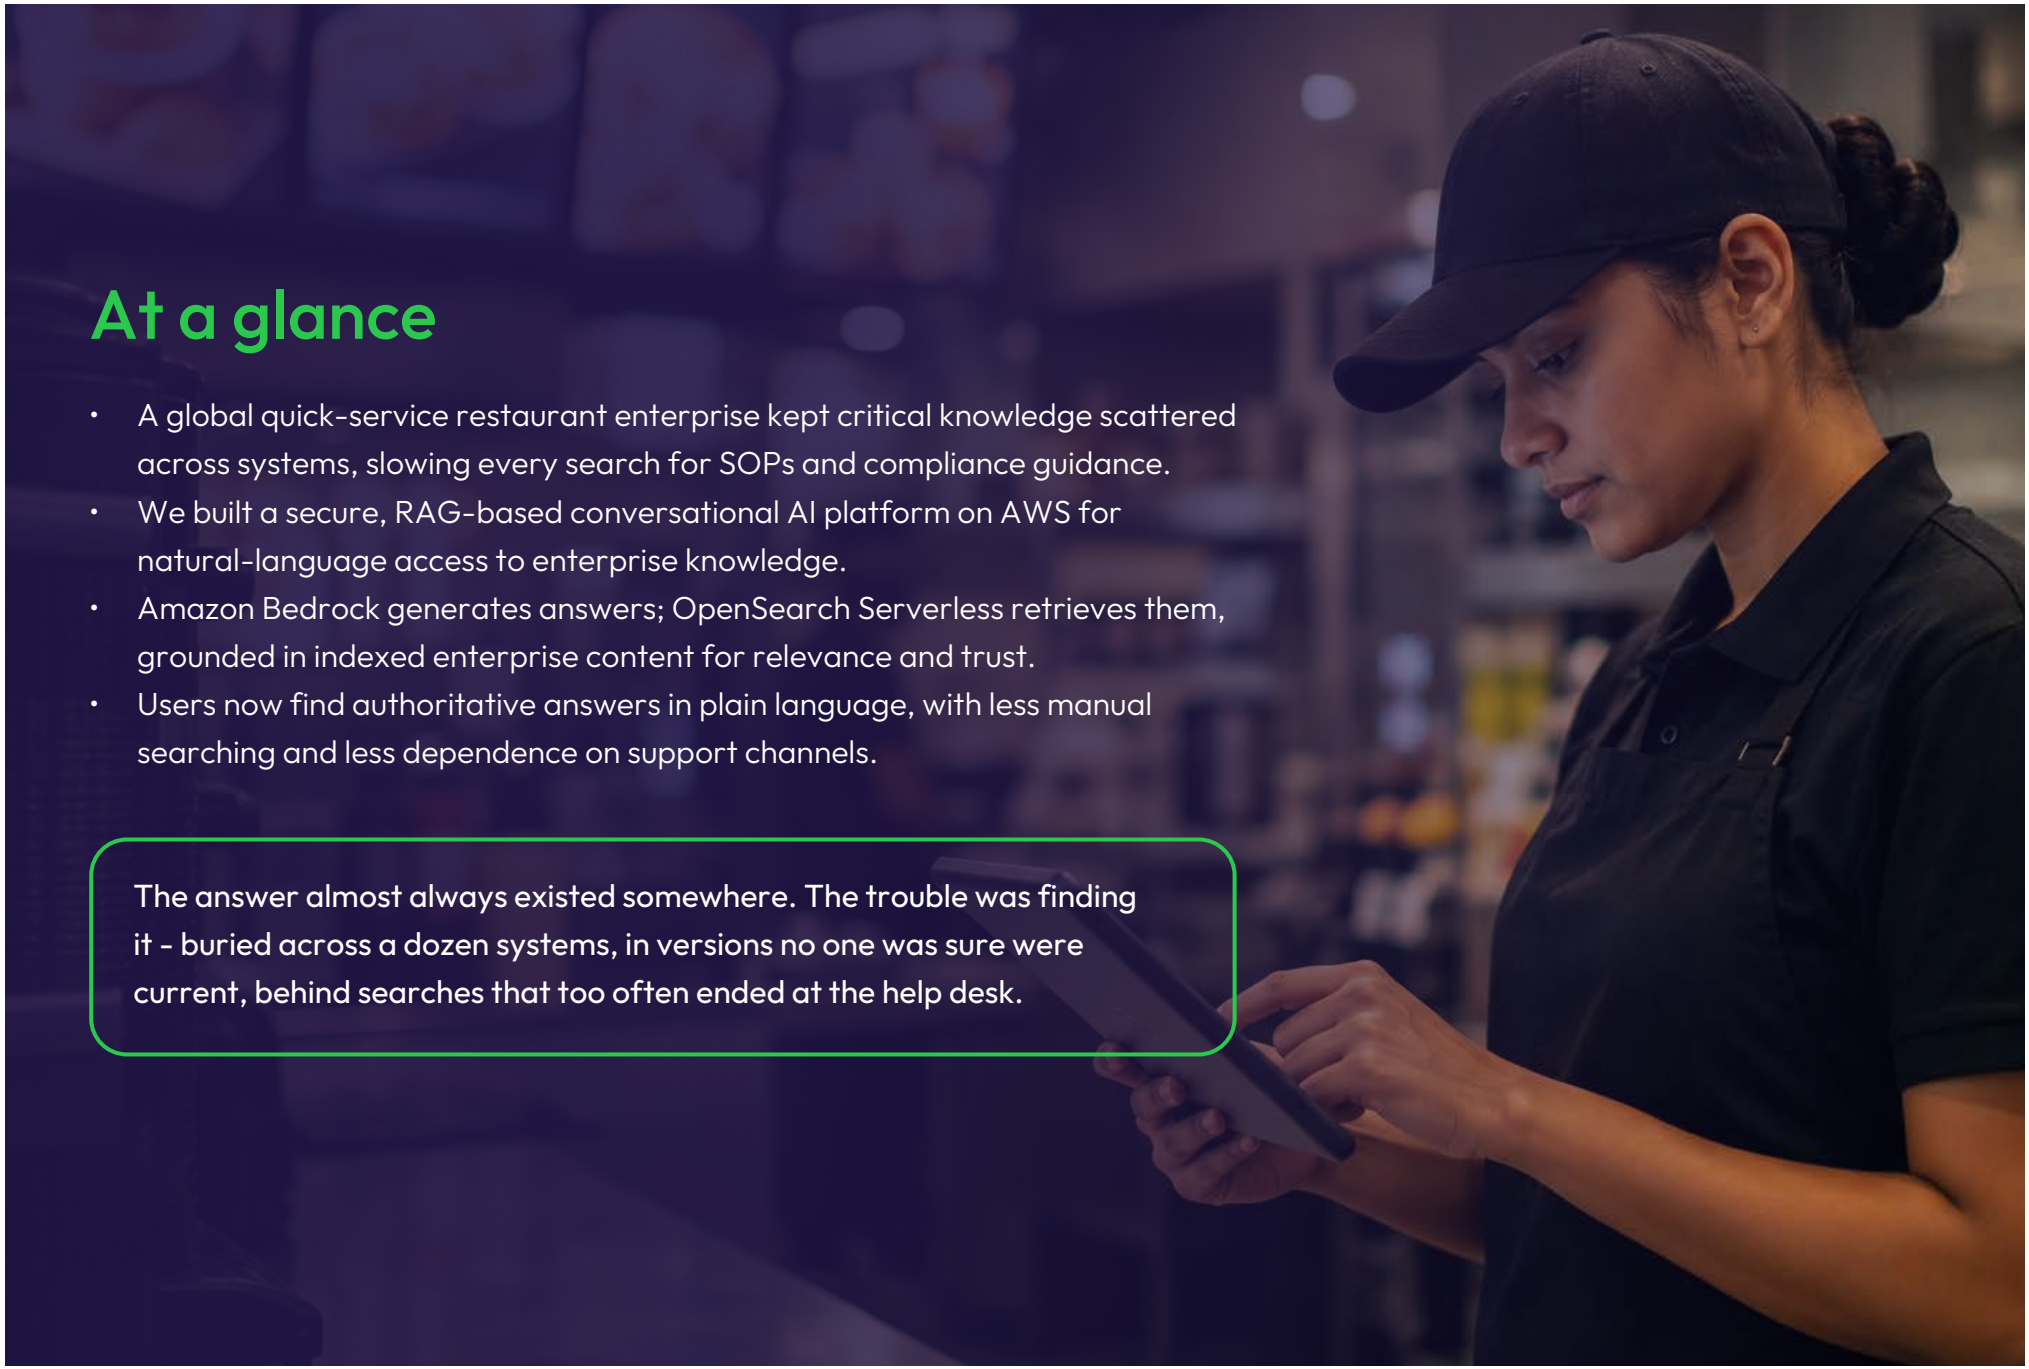

At a glance

- A global quick-service restaurant enterprise kept critical knowledge scattered across systems, slowing every search for SOPs and compliance guidance.
- We built a secure, RAG-based conversational AI platform on AWS for natural-language access to enterprise knowledge.
- Amazon Bedrock generates answers; OpenSearch Serverless retrieves them, grounded in indexed enterprise content for relevance and trust.
- Users now find authoritative answers in plain language, with less manual searching and less dependence on support channels.

The answer almost always existed somewhere. The trouble was finding it - buried across a dozen systems, in versions no one was sure were current, behind searches that too often ended at the help desk.

Knowledge without a map

Our client is a global enterprise in the quick-service restaurant industry, running a complex and highly distributed supply chain ecosystem. Its day-to-day depends on a constant flow of business-critical documentation - standard operating procedures, supplier onboarding materials, compliance guidelines, training resources - shared across a broad network of suppliers, franchise operators, and compliance teams.

The problem was never a shortage of knowledge. It was that the knowledge had no center. Critical information sat scattered across multiple systems and repositories, so finding the right answer meant knowing where to look and searching by hand. That fragmentation rippled outward: inconsistent access to standardized knowledge, heavier reliance on support channels, a user experience that shifted depending on who was asking and where, and real compliance risk in a business where an outdated answer carries consequences.

Critical information lived everywhere and nowhere - spread across systems no single search could reach.

Ask in plain language

We built a production-grade conversational AI platform on AWS, designed around a Retrieval-Augmented Generation (RAG) architecture so users could simply ask for what they needed in natural language and get an authoritative answer back.

The design pairs generation with retrieval so nothing is invented. Amazon Bedrock powers the language model that generates responses, while Amazon OpenSearch Serverless handles semantic search and vector retrieval across indexed enterprise content - so the platform answers from the company's own authoritative sources rather than guessing.

Amazon S3 stores the underlying documents, AWS Lambda and API Gateway provide serverless orchestration, Amazon DynamoDB holds structured metadata, and Amazon ElastiCache for Redis maintains session state and conversational context, so a follow-up question knows what came before.

We delivered the engagement end to end, from strategy and architecture design through model selection, implementation, and operationalization. The result is a scalable, cloud-native, serverless deployment - semantic retrieval and contextual grounding wired into integrated backend services - built to serve knowledge access across the entire enterprise.

We replaced the hunt with a conversation - one secure platform that understands the question and grounds every answer in source.

From hunt to answer

The platform gave the client something it never had: a single, intelligent layer for knowledge access across every repository. Instead of hunting through systems, internal users and external partners alike can now ask in plain language and get a relevant, authoritative answer drawn from indexed enterprise content.

The downstream effects follow from that. The chatbot not only provides answers but also cites the source from which the information was derived. Time spent locating critical information drops. Dependence on manual support channels eases. And because every response is grounded in the company's own source content, answers stay relevant and consistent - strengthening trust in enterprise knowledge and making it genuinely usable across the ecosystem.

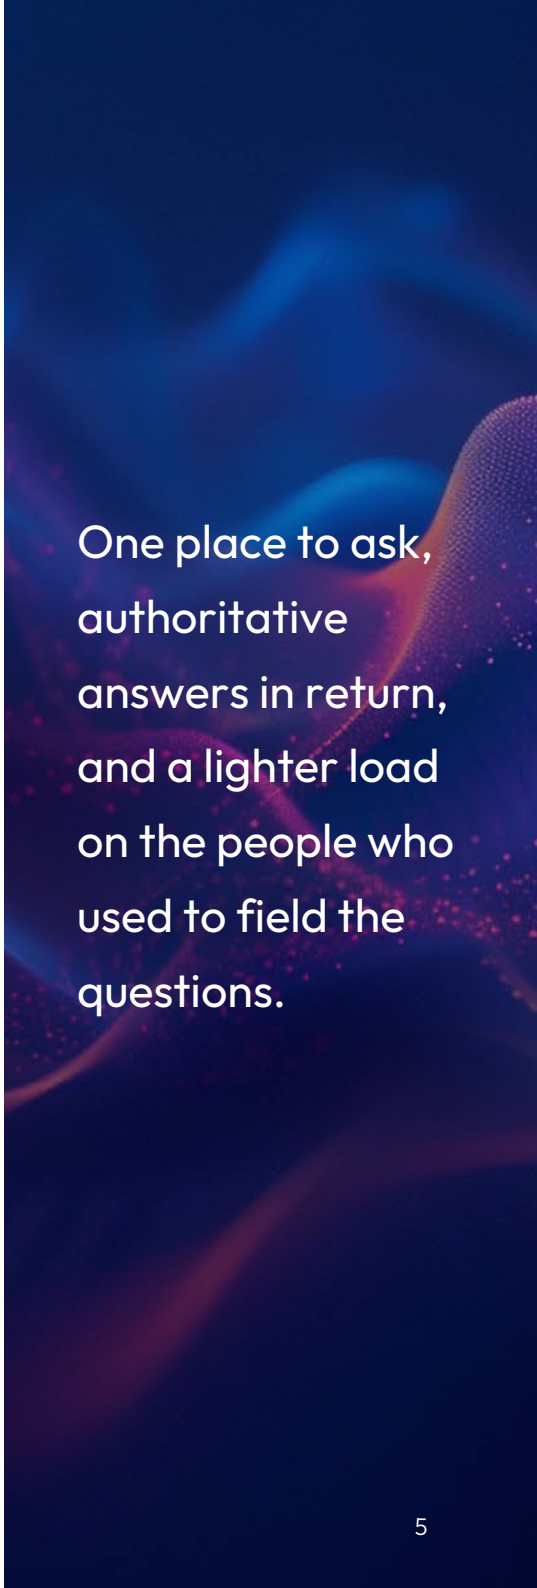

What the client achieved:

Centralized, intelligent knowledge access
across multiple repositories

Natural-language search
in place of manual hunting

Reduced time
locating critical information

Lower dependence
in place of manual hunting

Answers grounded
in authoritative enterprise sources

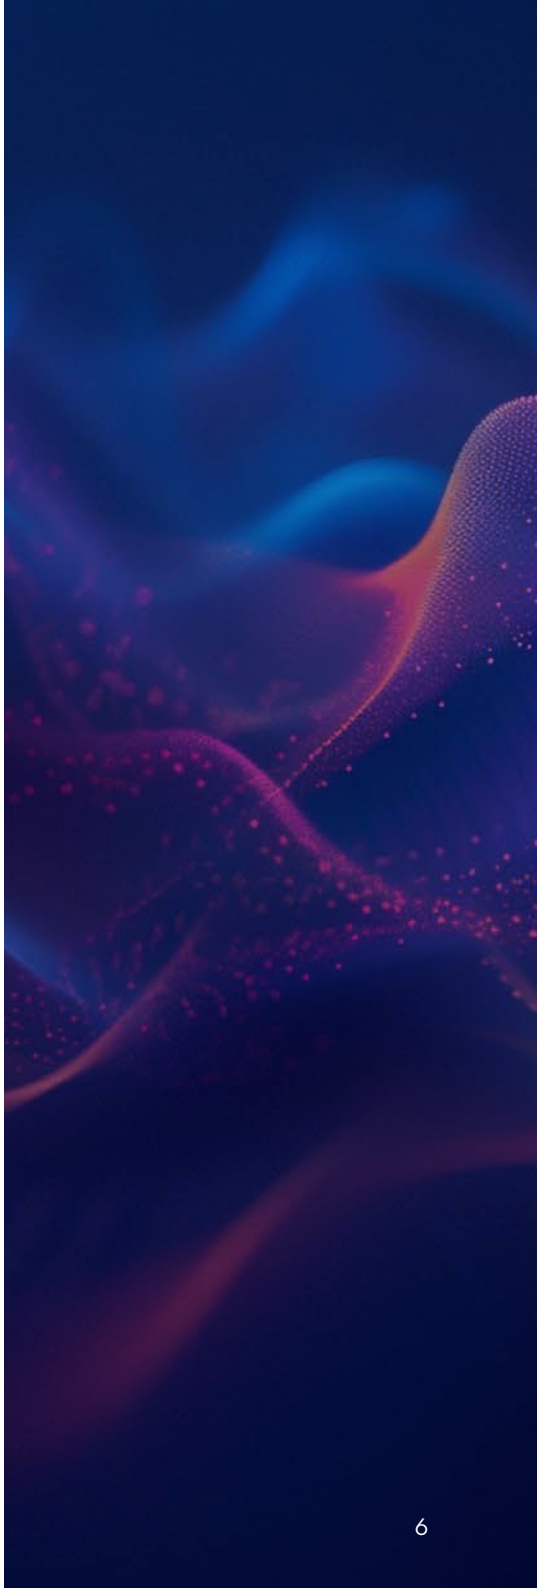

ABOUT BRILLIO

Brillio is The Enterprise AI Accelerator helping Fortune 1000 companies move from AI ambition to scaled impact, faster. Powered by our AI accelerator platform – Agentic Data and Application Management (ADAM), Brillio is one of the fastest-growing digital technology service providers, delivering transformation across five core workstreams: business-led transformation, customer experience transformation, AI and data engineering, digital engineering, and infrastructure engineering.

With 14 delivery locations across North America, Europe, and Asia and a team of over 6,000 customer-obsessed professionals, Brillio combines deep industry expertise, modern engineering, and accelerators to deliver measurable outcomes. Headquartered in Dallas, Texas, Brillio serves clients globally with a commitment to speed, scale, and measurable impact.

https://www.brillio.com/
Contact Us: info@brillio.com