BRIGHTGROVE

Fast, safe, reliable and adaptable
business partner to help your business evolve


**CASE STUDY**

# Modernizing a mission-critical ERP

with a Secure Azure AI Assistant for Manufacturing and Data Center Operations

***Location***
*France*

***Industry***
*HVAC, Refrigeration & Data Center Equipment Manufacturing*

***Duration***
*~12 weeks (MVP delivery)*

***Team Size***
*2–3 core specialists (Senior DevOps Engineer, Senior AI/ML Engineer, PM oversight)*

***Engagement Scope***
*AI Knowledge Management, Azure Cloud Architecture, LLM + RAG System Design, Legacy ERP Modernization Support*

## CLIENT OVERVIEW

LFB Group operates a mission-critical ERP system. The system supports core operations across HVAC manufacturing and Data Center infrastructure—from inventory management and production control to sales order processing and business accounting.

While the system was critical to operations, keeping it running long-term had become a challenge. The goal was to update the technology to meet current needs. This involved organizing the team's internal guides and moving toward a setup where technical knowledge was shared across the whole group.

The organization needed a way to safeguard institutional knowledge and modernize knowledge access—all without disrupting daily operations.

## CHALLENGE

LFB Group sought to address the technical complexities common in enterprises that maintain long-standing, heavily customized systems. The goal was to refine several intersecting factors that arise as these systems evolve.

Together with PlatinumProd, a technical consultancy partner on the project, we identified these optimization opportunities and began working toward a modernized framework:

**Consolidating Knowledge Assets:** Years of ERP operation had generated a vast wealth of documentation, code history, and technical manuals. The goal was to transform these valuable resources from a collection of individual assets into a unified intelligence ecosystem.

**Institutionalizing Technical Expertise:** As the organization prepared for its next strategic phase, a key priority was the formalization of its deep internal expertise. This involved capturing the specialized logic and customization history of the ERP system to ensure long-term operational resilience.

**Enhancing Troubleshooting & Onboarding:** With a highly complex environment, the organization sought to further streamline its incident resolution and onboarding processes. The objective was to transition from manual cross-referencing of logs and code to a more integrated, high-speed discovery model.

**Security & Governance Standards:** The organization followed a strict internal policy to keep all data and operations fully under their own control. This meant any project had to run entirely within their existing Azure environment, without using outside APIs or third-party platforms to handle internal documents. Every part of the system also had to work directly with their established security rules and user access levels.

After the discovery phase, the focus shifted to building a secure intelligence layer around the ERP. The goal was to consolidate the organization's deep technical heritage and ensure that years of institutional knowledge were captured and easily accessible for the future.

### Our Approach: Structure First, Then Intelligence

Before the current generation of AI tools became available, addressing LFB's challenges would have meant a traditional knowledge management initiative: months of manual documentation audits, structured interviews with the SME to extract and formalize tribal knowledge, building a conventional search portal or wiki, and hiring additional engineers to reduce key-persons dependency through headcount. It would have been expensive, slow, and static. The resulting knowledge base would have started decaying the moment it was published, as the live ERP system continued to evolve.

With enterprise-grade LLM and RAG capabilities now mature enough for production use, a fundamentally different approach became possible—one that could ingest, structure, and make searchable the full breadth of LFB's institutional knowledge in weeks rather than quarters, and keep it continuously updated. But the availability of AI alone was not the solution. The question was how to apply it responsibly.

A common mistake in enterprise AI adoption is to focus on the AI itself, selecting a model, fine-tuning prompts, generating output, while neglecting the structural foundation that determines whether that output will be useful. Brightgrove took the opposite approach.

The core principle behind GuardianLLM is that AI should generate value. An LLM connected to a disorganized knowledge base will produce confident-sounding answers that are incomplete, outdated, or wrong. The same LLM connected to a well-structured, continuously updated, properly indexed knowledge base becomes a reliable operational tool.

This meant that before any AI component was introduced, the engagement focused on:

**Knowledge architecture:** Mapping, categorizing, and structuring LFB's scattered documentation, ERP source code, ticketing history, and operational logs into a coherent, searchable corpus.

**Source quality governance:** Establishing ingestion rules that filter, version, and validate content before it enters the knowledge base, ensuring the AI only works with authoritative, current materials.

**Retrieval design over generation design:** Investing in Azure AI Search's vector and semantic indexing to ensure that the right information is surfaced before the LLM ever attempts to generate an answer.

**Continuous refresh mechanisms:** Building automated pipelines (Git sync, log ingestion, ticketing integration) so the knowledge base stays current without manual intervention.

## SOLUTION

Brightgrove designed and delivered GuardianLLM—a secure, Azure-hosted Retrieval-Augmented Generation (RAG) knowledge assistant built exclusively on LFB's internal materials. The system was designed to operate entirely within the client's-controlled cloud environment with zero external data exposure.

### Secure Enterprise Architecture

| Component | Role |
|---|---|
| Azure OpenAI | Core LLM engine for natural language understanding and generation |
| Azure AI Search | Vector and semantic indexing for relevance-ranked document retrieval |
| Azure Blob Storage | Centralized document repository for ingested materials |
| Azure Entra SSO | Enterprise identity management with role-based access control |
| Monitoring & Logging | Usage analytics, cost governance, and audit trail capabilities |

### Phased MVP Delivery

With a 12-week MVP deadline and a requirement for measurable business validation, the engagement followed a three-phase delivery model that balanced speed with quality, enabling early value delivery while progressively expanding capabilities:

**Phase 1. Foundational Platform:** Manual document ingestion, ERP codebase upload, RAG pipeline implementation, secure web interface with SSO integration, and initial user acceptance testing.

**Delivered:** A functional internal AI assistant capable of answering natural-language questions against uploaded ERP documentation and code.

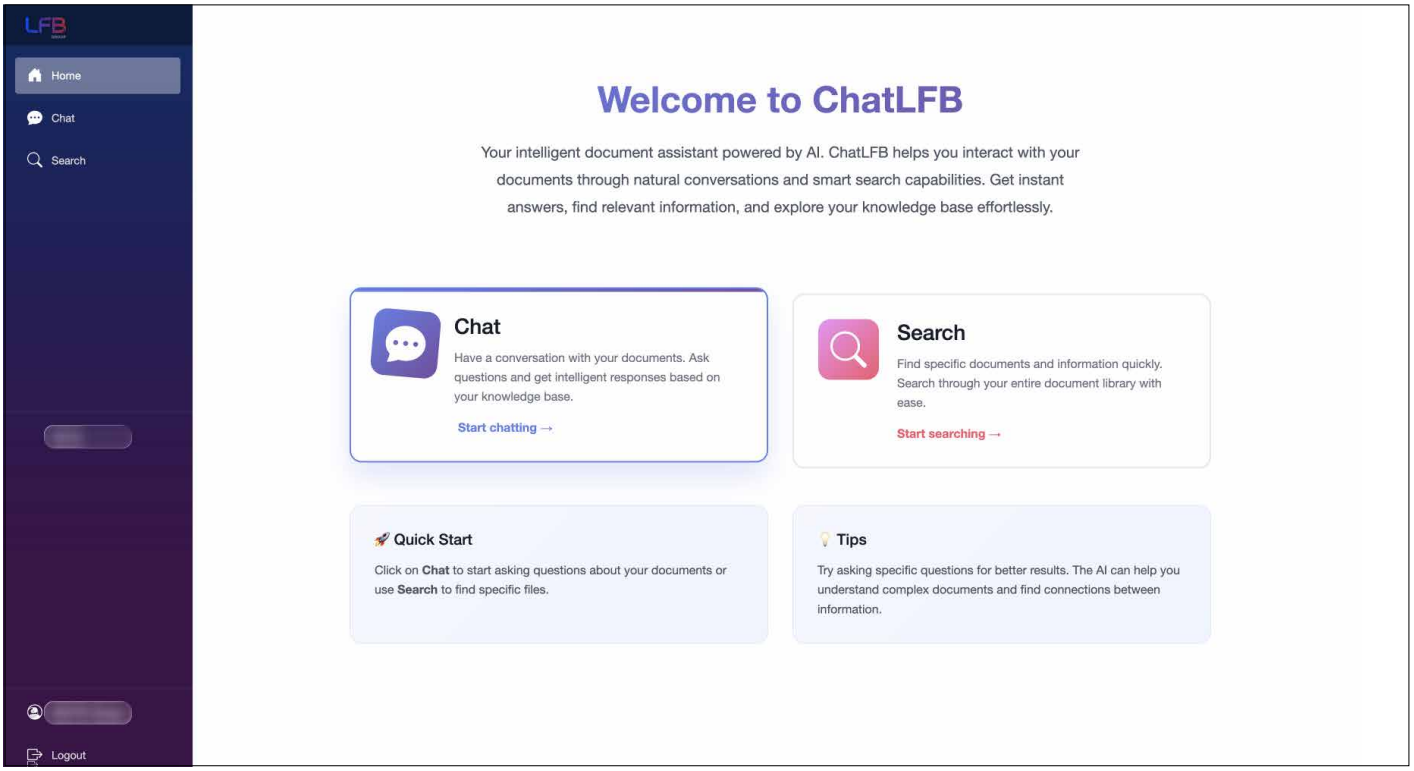

**Phase 2.** Advanced Intelligence: OCR and document cracking for scanned materials, image-aware response generation (schematics, screenshots, diagrams), enhanced ingestion pipeline, and prompt optimization.

**Delivered:** A multimodal intelligence layer that could interpret visual documentation alongside text.

**Phase 3.** Full Integration & Automation: Git repository integration, ticketing system synchronization, automated log ingestion, knowledge base continuous refresh.

**Delivered:** A fully integrated, continuously updated AI system that stays current with evolving ERP operations.

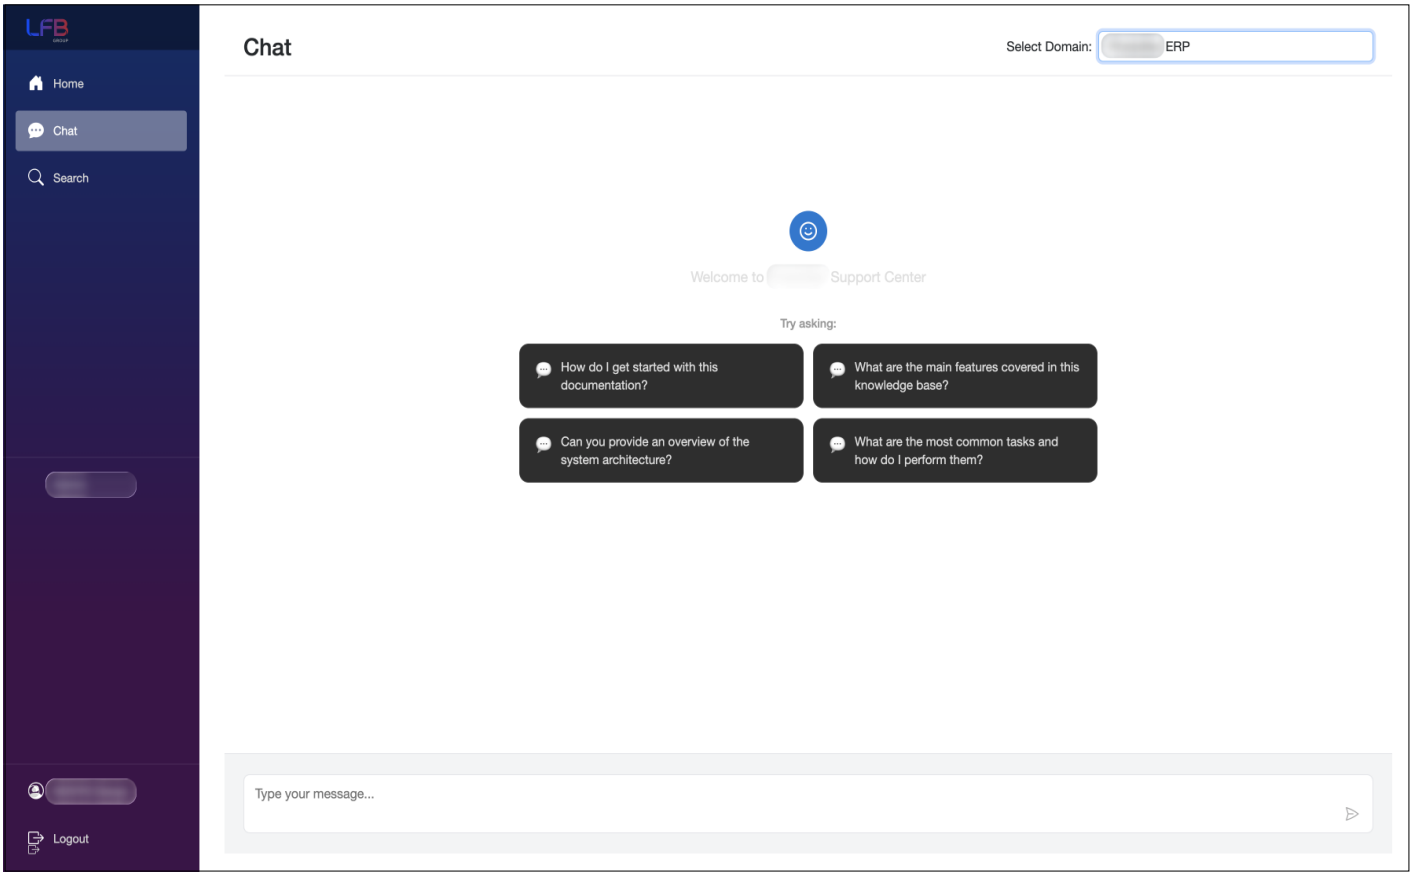


**Functional Capabilities**

Upon delivery, GuardianLLM provided the following production-ready capabilities:

- Natural-language Q&A across ERP source code, technical documentation, and operational logs.
- Image and schematic recognition – the system interprets diagrams, screenshots, and technical drawings embedded in documentation and incorporates them into responses.
- Internal search engine with relevance-based document ranking, functioning as an internal Google for the organization's knowledge base.
- Automated log summarization and correlation for faster incident investigation.
- ERP function tracing – e.g., tracing VAT calculation logic, BOM assembly rules, and custom business procedures across the codebase.
- Context-aware troubleshooting support that combines code, documentation, and historical incident data.
- Role-based access control ensuring that users only see information authorized for their department.
- Multilingual output (English, French, and extensible to additional languages).
- Internal code and SQL guidance aligned with LFB's coding standards and conventions.

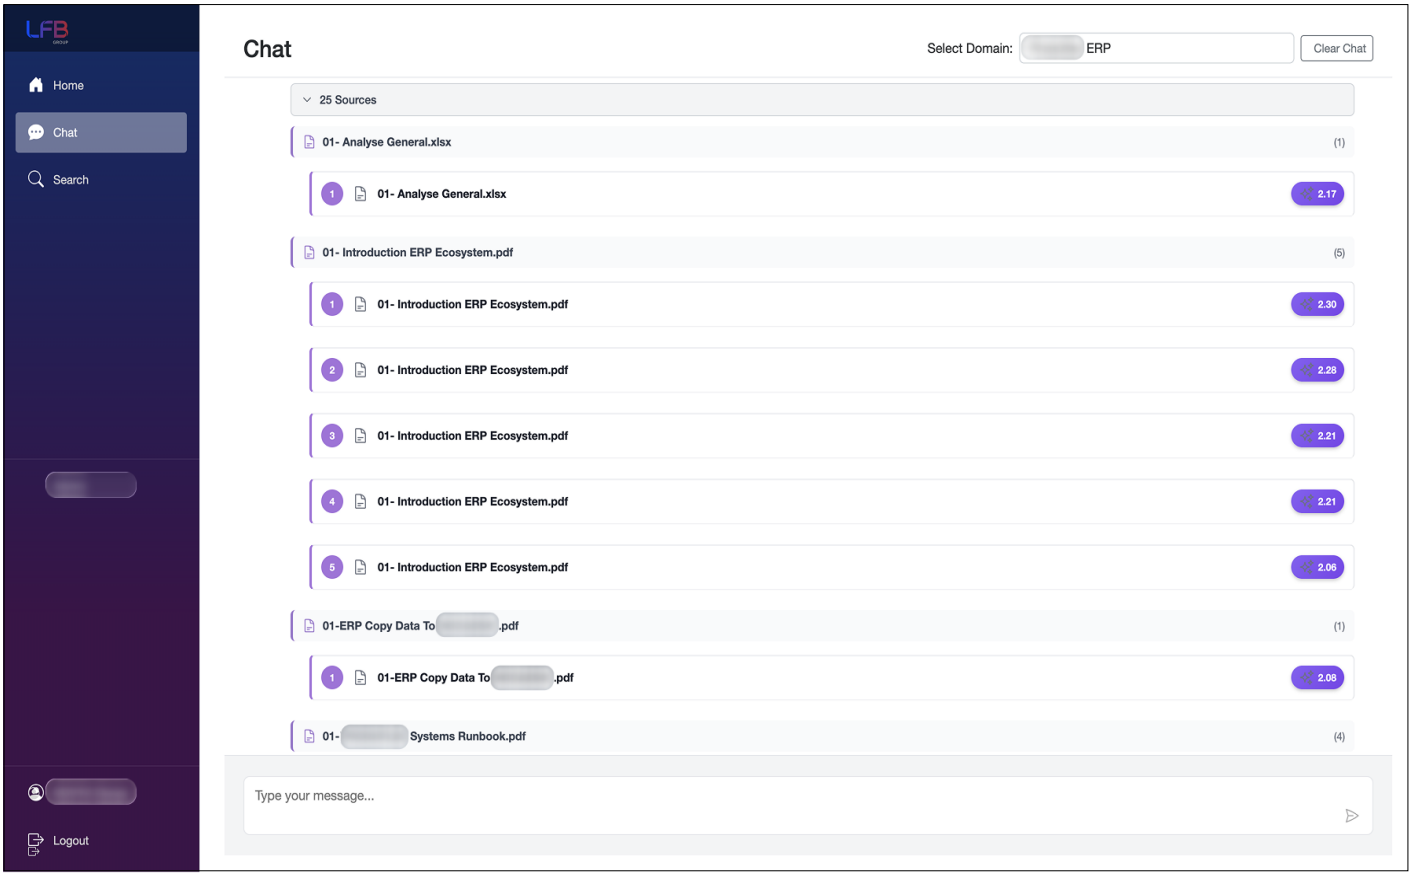
LFB
Home
Chat
Search
Logout
Chat
Select Domain: ERP
Clear Chat
25 Sources
01- Analyse General.xlsx (1)
1 01- Analyse General.xlsx 2.17
01- Introduction ERP Ecosystem.pdf (5)
1 01- Introduction ERP Ecosystem.pdf 2.30
2 01- Introduction ERP Ecosystem.pdf 2.28
3 01- Introduction ERP Ecosystem.pdf 2.21
4 01- Introduction ERP Ecosystem.pdf 2.21
5 01- Introduction ERP Ecosystem.pdf 2.06
01-ERP Copy Data To .pdf (1)
1 01-ERP Copy Data To .pdf 2.08
01- Systems Runbook.pdf (4)
Type your message...

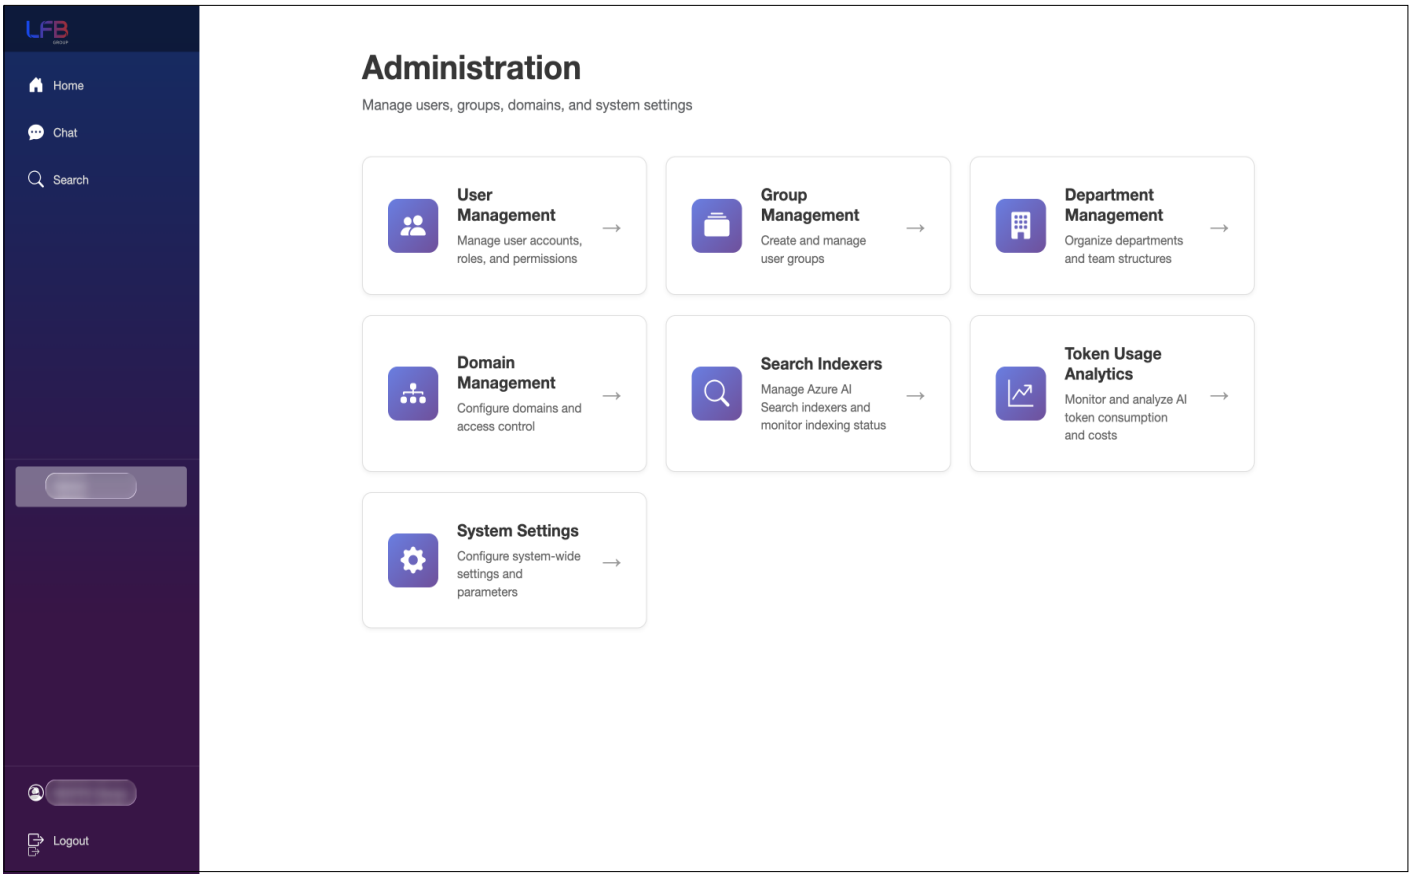
LFB
Home
Chat
Search
Logout
Administration
Manage users, groups, domains, and system settings
User Management
Manage user accounts, roles, and permissions
Group Management
Create and manage user groups
Department Management
Organize departments and team structures
Domain Management
Configure domains and access control
Search Indexers
Manage Azure AI Search indexers and monitor indexing status
Token Usage Analytics
Monitor and analyze AI token consumption and costs
System Settings
Configure system-wide settings and parameters

## RESULTS

GuardianLLM is now live in LFB's production environment and actively used by internal teams on a daily basis. What began as a tool scoped exclusively for the IT department quickly demonstrated broader applicability across the organization.

### Measurable Business Impact

**Reduced Key-Experts Dependency:** Institutional ERP knowledge, previously held by a few key experts, is now searchable, structured, and accessible to the entire team. The organization's operational continuity risk has been reduced. In the past, accessing details on the ERP's specialized logic relied on direct consultation with key experts. Today, GuardianLLM provides those same insights in seconds. By centralizing thousands of technical assets, the organization has transformed specialized individual expertise into a shared, high-availability organizational asset.

**Faster Onboarding:** New engineers can ask natural-language questions and receive context-rich, source-cited answers immediately. The system effectively serves as an always available senior colleague. What previously involved an extended period of hands-on mentorship and peer-to-peer knowledge sharing is now available through a direct conversation with GuardianLLM. New team members get context-rich, source-cited answers on their first day.

**Accelerated Troubleshooting:** Support engineers retrieve incident resolutions, log correlations, and relevant code references in minutes rather than hours. Cross-referencing that previously required manual search across multiple systems is now handled by a single query. Incident investigation that previously took hours of manual cross-referencing across logs, documentation, and code now takes seconds.

**Improved Knowledge Governance:** LFB now maintains a centralized, continuously indexed, and automatically updated knowledge repository, replacing fragmented documentation with a single source of truth. Documentation that was scattered across shared drives, email threads, ticketing systems, and one person's memory now lives in a single, repository. Every document is automatically refreshed – the knowledge base stays current without anyone manually maintaining it.

**Strategic Due Diligence Readiness:** Knowledge risk has been significantly reduced through structured capture and searchable intelligence, strengthening the organization's position for strategic evaluations and compliance reviews. Before GuardianLLM, a due-diligence review of ERP knowledge would have surfaced a single-point-of-failure risk that could affect valuation. Now, the organization can demonstrate structured, searchable, independently accessible institutional knowledge – a fundamentally stronger position for any strategic evaluation.

### Expansion Potential

Originally planned for IT usage only, the architecture was intentionally designed to scale across departments. LFB has identified the following areas for expansion:

**Finance:** Automated retrieval of accounting procedures and ERP financial logic.

**Sales:** Quick access to product configuration rules and pricing logic.

**Supply Chain**: Vendor documentation, procurement workflows, and BOM traceability.

**Manufacturing:** Production scheduling rules, quality control procedures.

**Human Resources:** Policy documentation and employee onboarding support.

**Cross-Department:** Organization-wide knowledge search and institutional memory.

## WHY IT WORKED

The success of GuardianLLM was not accidental. Several deliberate architectural and delivery decisions contributed to a smooth engagement and a result that exceeded initial expectations:

**Structure before intelligence:** The team invested in knowledge architecture and source quality before activating AI capabilities, ensuring the system generates value rather than volume.

**RAG-first architecture:** Every response was grounded in verified internal sources, reducing hallucination risk and increasing trust in enterprise use cases.

**Clean documentation corpus:** A well-prepared documentation base made reliable retrieval possible and strengthened overall system performance.

**Phased execution:** Allowed the team to deliver early value while incorporating real user feedback into subsequent phases, minimizing risk without sacrificing momentum.

**Full Azure isolation:** Satisfied enterprise governance requirements from day one, removing a common adoption blocker and accelerating stakeholder buy-in.

**Close SME collaboration:** Ensured that critical institutional knowledge was captured systematically rather than lost through attrition or organizational change.

**Agile delivery cadence:** With weekly sync meetings maintained alignment between Brightgrove and LFB, minimized feedback loops, and kept the project on track for its 12-week deadline.

## STRATEGIC VALUE

GuardianLLM evolved from a focused MVP into a scalable AI knowledge platform that now serves as the foundation for LFB's broader ERP modernization strategy. The system functions as a reusable blueprint for secure enterprise AI adoption, proving that legacy environments can benefit from modern AI without compromising governance, security, or operational stability.

Critically, the project demonstrated that the differentiator in enterprise AI is not the model – it's the methodology. The right AI tools, applied to a well-structured knowledge foundation with clear governance, produce measurable business outcomes. Without that foundation, even the most advanced model generates noise.

### What's Next

Planned evolution includes:

- Expansion to additional business units (Finance, Supply Chain, HR)
- Advanced analytics layer with usage insights and knowledge gap identification
- Deeper ERP automation through AI-assisted workflow triggers
- AI-supported modernization initiatives for legacy code migration planning
- Continuous model tuning based on production usage patterns and user feedback

**Ready to take the next step?**
**Reach us today at info@brightgrove.com**