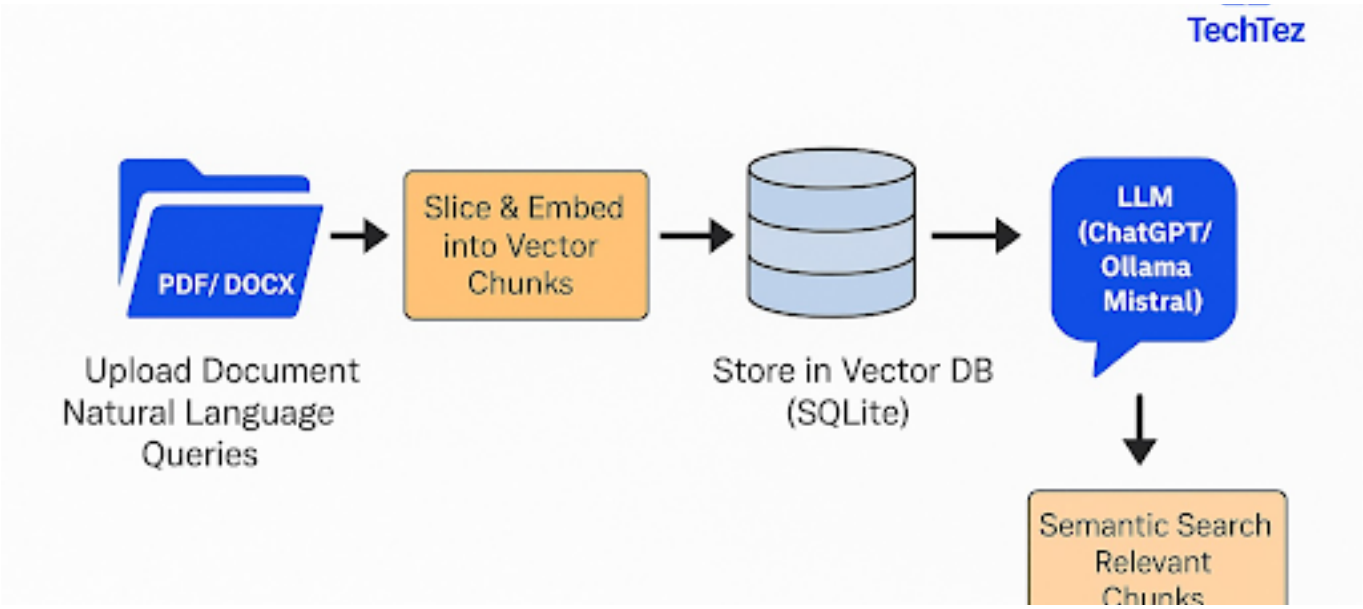

Conversational Chatbot Powered by Document Intelligence

Posted on September 7, 2025 by Admin

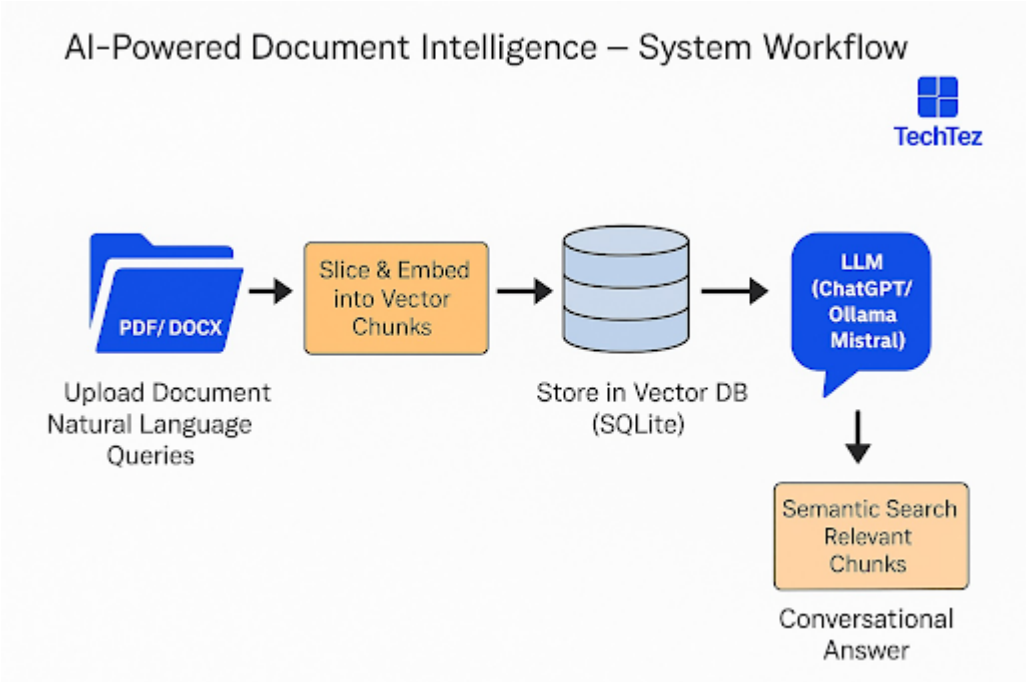

Case Study

Linkedin X-twitter

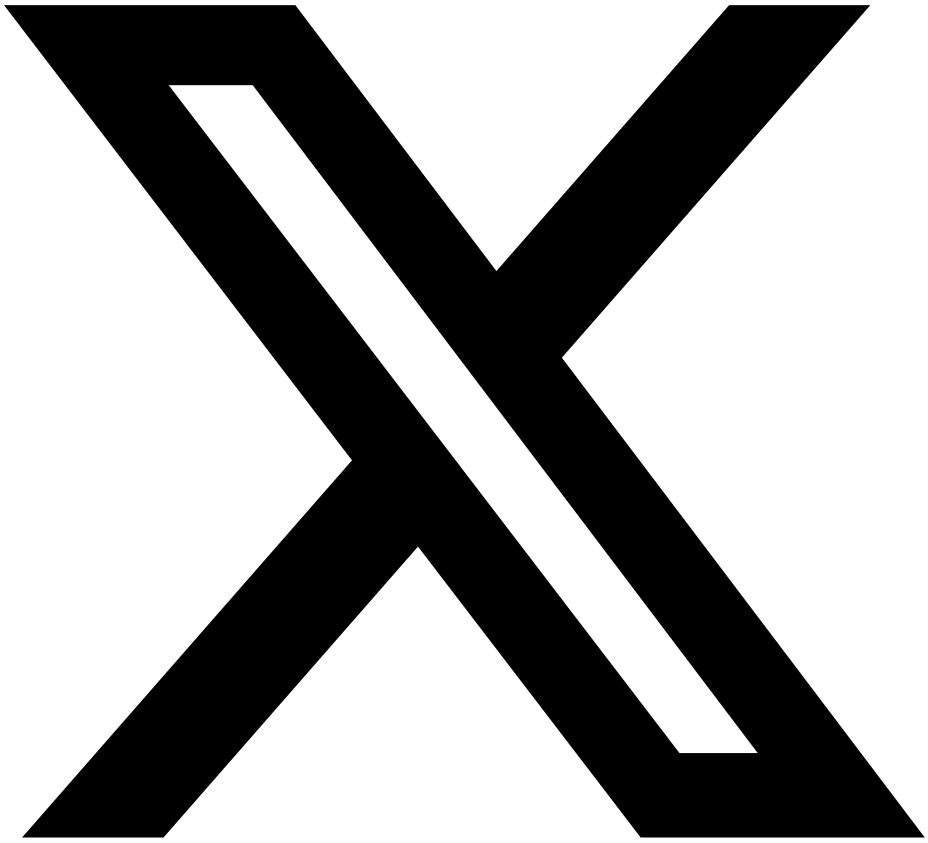

Facebook

Conversational Chatbot Powered by Document Intelligence

Client

A mid-sized SaaS company specializing in HR and payroll management faced a growing barrier: their support and operations teams struggled to quickly extract key information from a maze of internal documents: policy PDFs, compliance guides, and technical manuals. Manual searching and frequent escalations to subject matter experts slowed down resolutions and increased workloads.

TechTez was brought in to develop an AI-powered chatbot that could extract context-aware answers from unstructured documents using natural language. The solution now powers internal support workflows across multiple departments.

The Challenge

Enterprise teams are often buried under static, unstructured documents in a variety of formats (PDF, DOCX, etc.). Retrieving accurate, contextual answers is slow, labor-intensive, and not scalable, especially when answers are scattered across numerous files.

Key obstacles included:

- **No Contextual Search:** Existing document systems lacked the intelligence for natural, context-driven queries.

-  **Manual Reviews:** Employees wasted valuable hours scanning lengthy documents.

- **Disconnected Answers:** Traditional tools made it difficult to link user questions with the right portions of content.

- **Limited Accessibility:** No intuitive interface for asking plain-language questions; only tech-savvy users could dig deep.

Our Strategy

TechTez developed an advanced Conversational AI Chatbot using Retrieval-Augmented Generation (RAG) to transform how teams interact with company knowledge.

Architecture & Workflow

Solution Highlights:

- **Document Ingestion & Structuring:** All internal documents (PDF, DOCX) are automatically divided into logical, searchable blocks for optimal retrieval.

- **Vectorization & Storage:** Each content block is embedded and stored in a lightweight vector database (SQLite) for high-speed, accurate semantic search.

- 
Semantic Matching: When users ask a question, the chatbot semantically matches the query to the most relevant sections of content—regardless of wording.

- **AI-Powered Answer Generation:** Relevant sections are passed to a Large Language Model (ChatGPT or Ollama Mistral) to generate clear, human-like responses.

Technology Architecture

- LLMs: ChatGPT, Ollama Mistral

- Vector Storage: SQLite

- Programming: Python, LangChain, Retrieval-Augmented Generation (RAG)

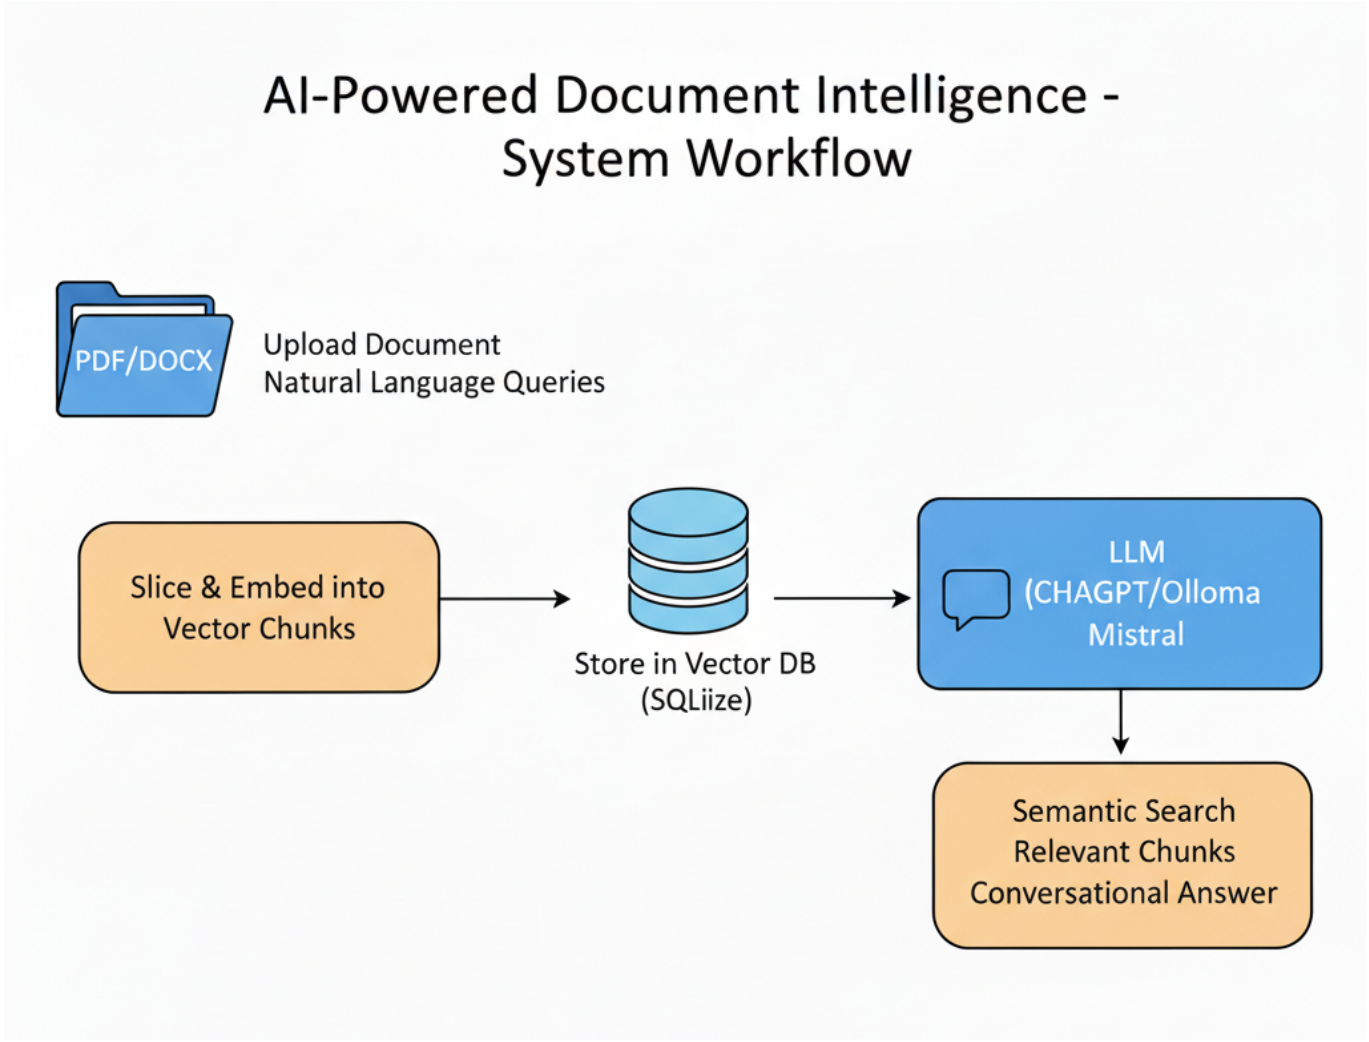
AI-Powered Document Intelligence -
System Workflow
PDF/DOCX
Upload Document
Natural Language Queries
Slice & Embed into
Vector Chunks
Store in Vector DB
(SQLiize)
LLM
(CHAGPT/Olloma
Mistral
Semantic Search
Relevant Chunks
Conversational Answer

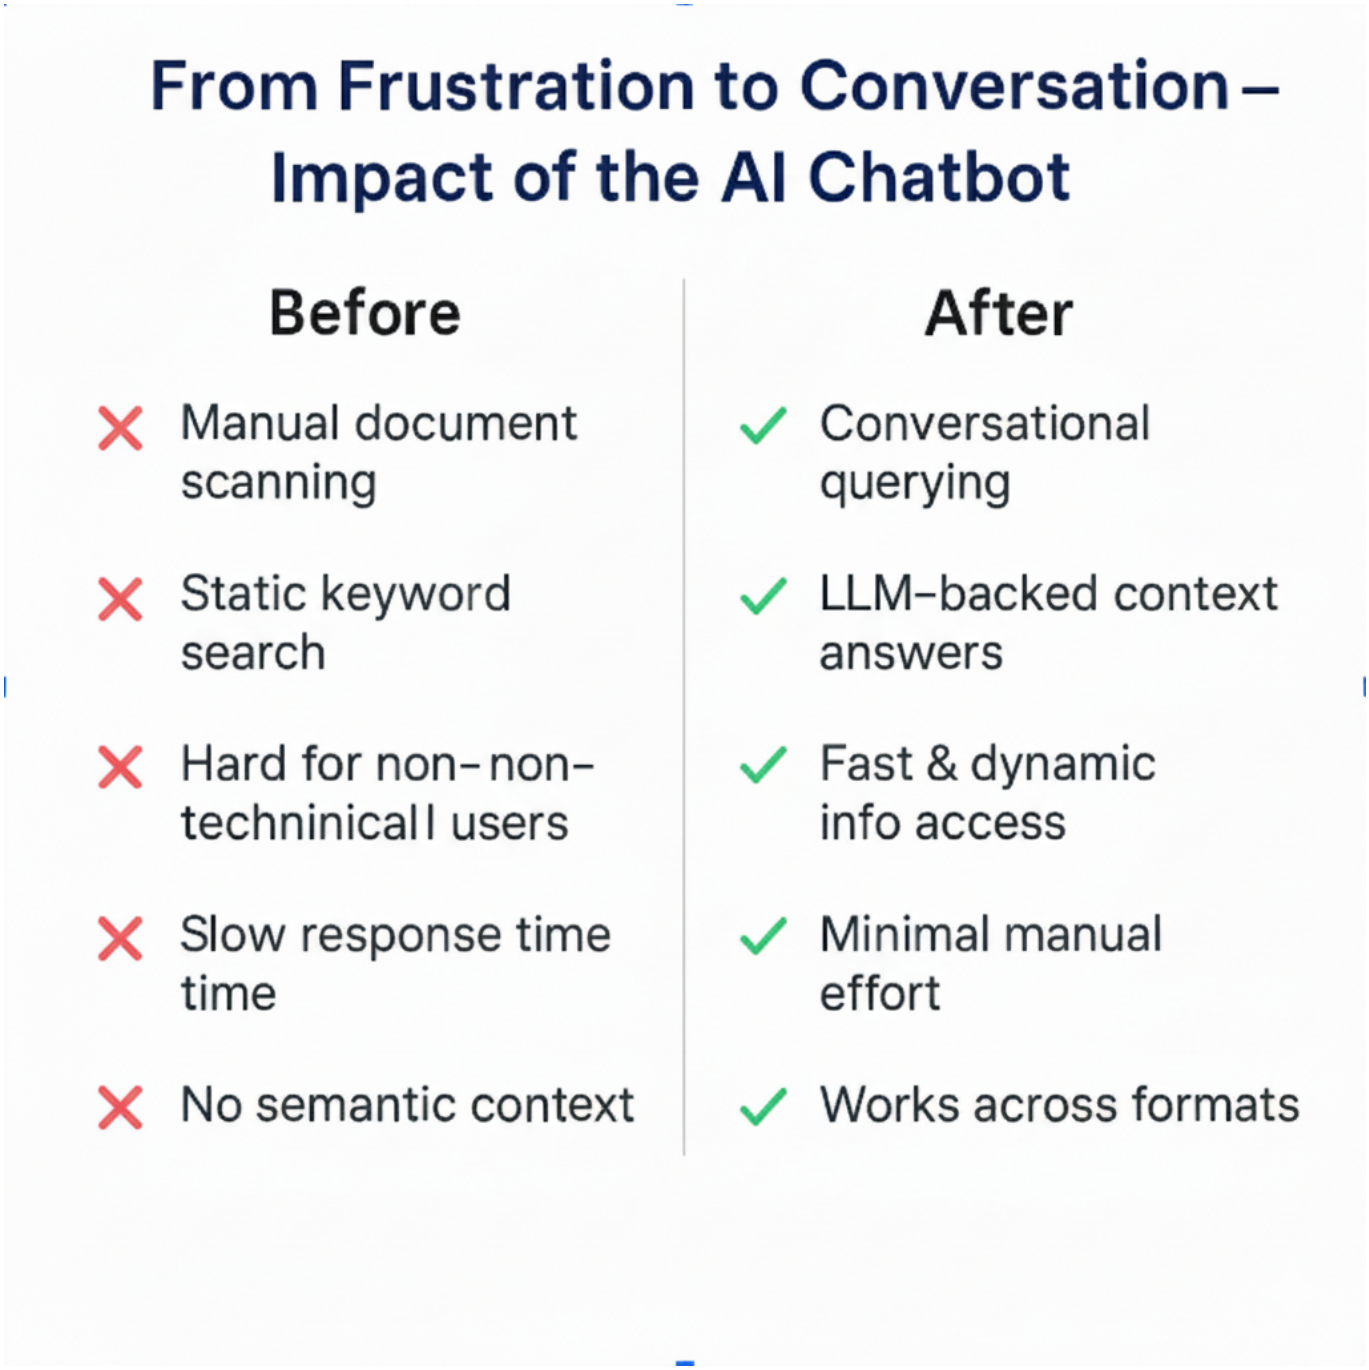

Results & Impact

-  Instant, conversational access to any internal document

- Search time reduced from minutes to seconds

- No need for manual tagging or indexing

- User-friendly: usable by non-technical staff

- Multi-format support (PDF, DOCX, etc.)

- Scalable, model-agnostic deployment

Outcomes:

- **40–50%** reduction in average document lookup time within the first two weeks

- **30%** decrease in support case resolution times

- Significant drop in escalations to SMEs, freeing experts for higher-value work

-  Improved response accuracy, reducing follow-up queries

Our Thought Leadership Guides

- Case Study

Simplifying Complex Telecom Integrations Using a Scalable Numbering Platform

A mid-sized SaaS company specializing in HR and payroll management faced a growing barrier:

- Case Study

Agentic AI in Telecom & Healthcare: The Platform Engineering Playbook

- Case Study

RAG Done Right: How to Build Enterprise-Grade Knowledge Assistants


TechTez
RAG
Done
Right
Enterprise Sources
Normalize & Enrich
Structure Aware Chunking
Hybrid Retrieve + Rerank
Security & Policies
Trusted Answers
How to Build Enterprise-Grade Knowledge Assistants | Beyond a Simple Vector Store