



# AI KNOWLEDGE BASE FOR ENTERPRISE

*Solution Portfolio*

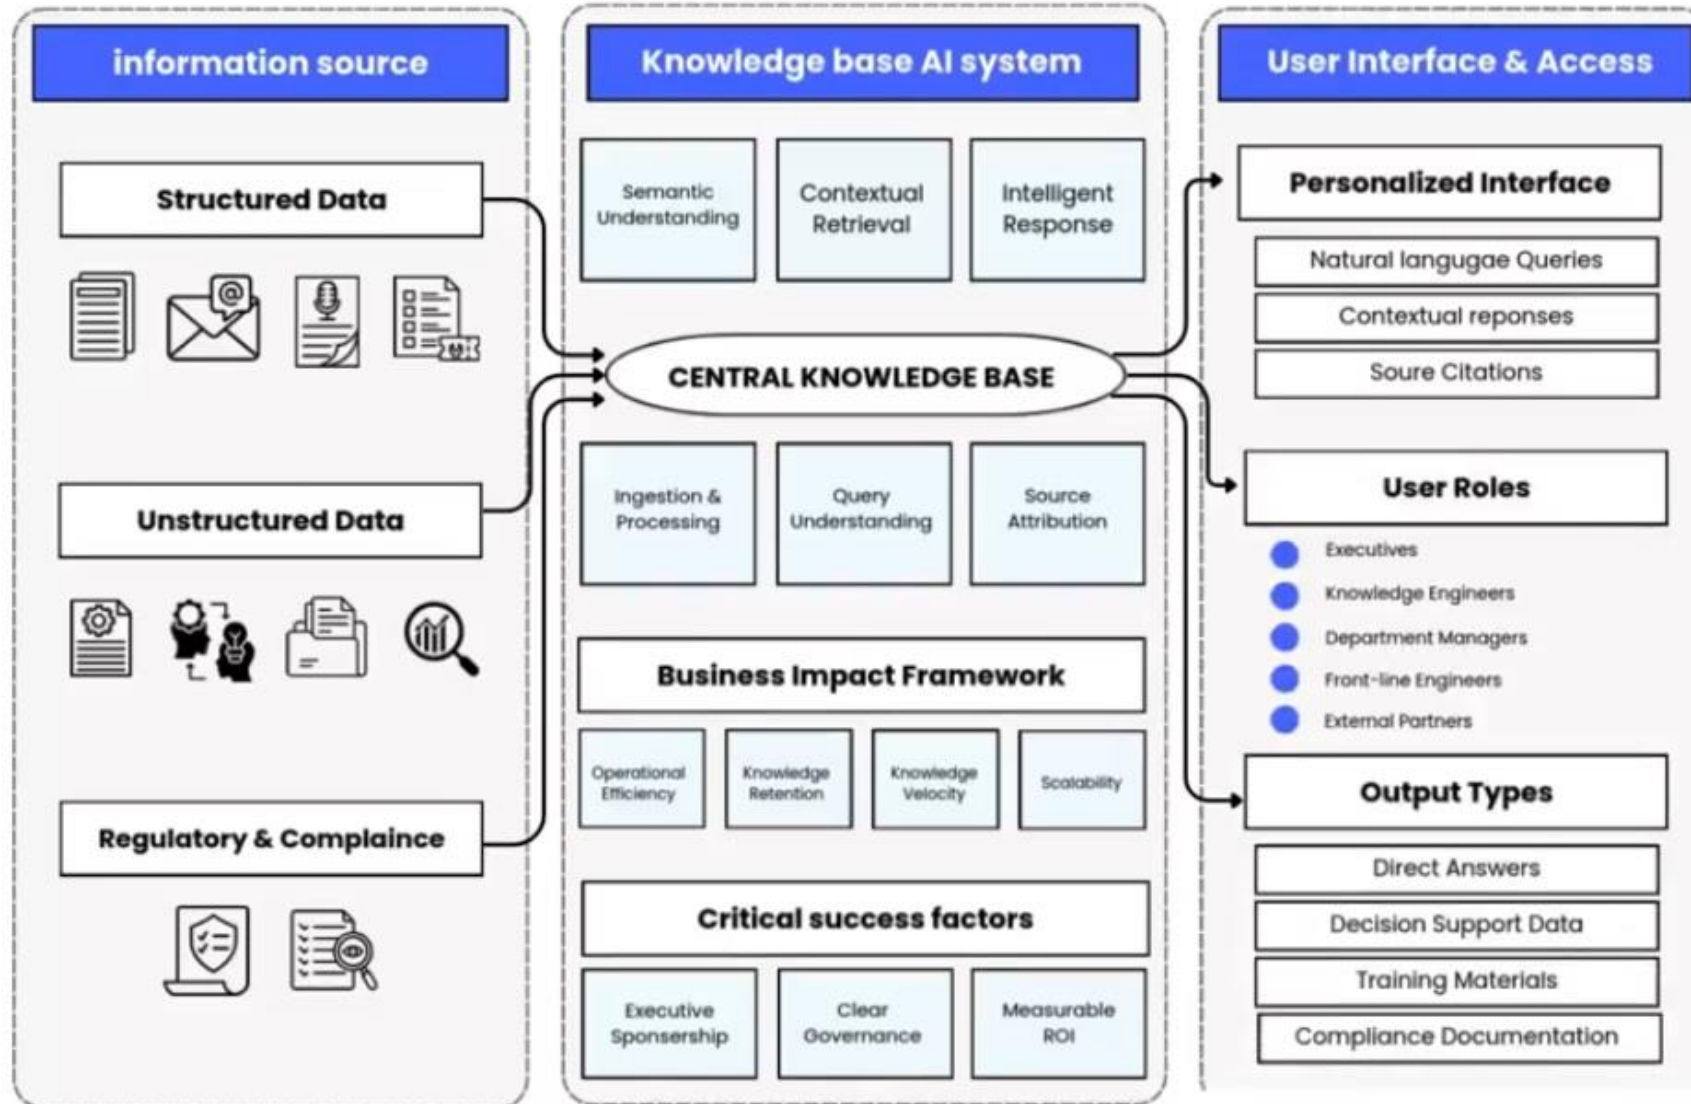


## SOLUTION OVERVIEW

*This solution combine Optical Character Recognition (OCR), cost-optimized Large Language Models (LLM), and intelligent document categorization into a single, secure platform*

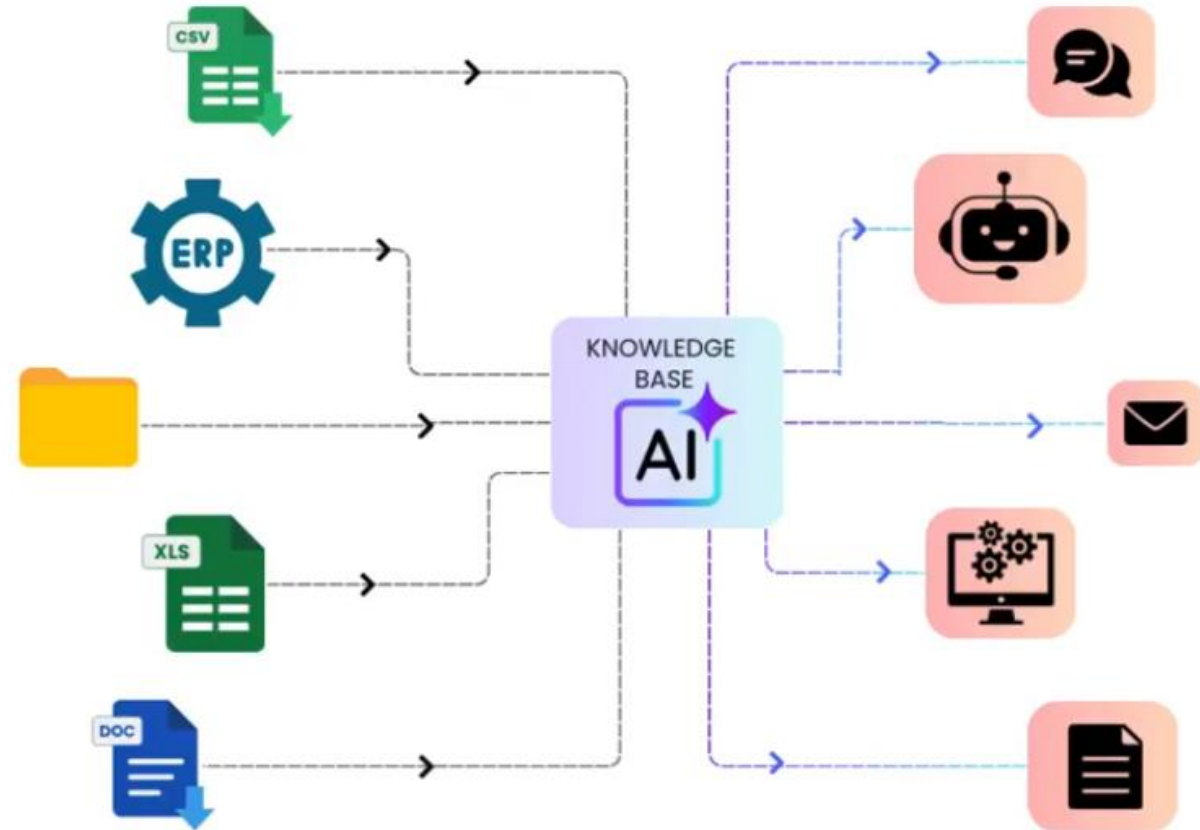
# OUR SOLUTION

## IS PLATFORM AI KNOWLEDGE BASE



# WHAT IS AI KNOWLEDGE BASE

Knowledge Base AI is an intelligent system that **ingests, comprehends, and retrieves** information from your organization's collective knowledge repositories



# WHAT IS AI KNOWLEDGE BASE

## Pillar 1: Document Ingestion (OCR Extract):

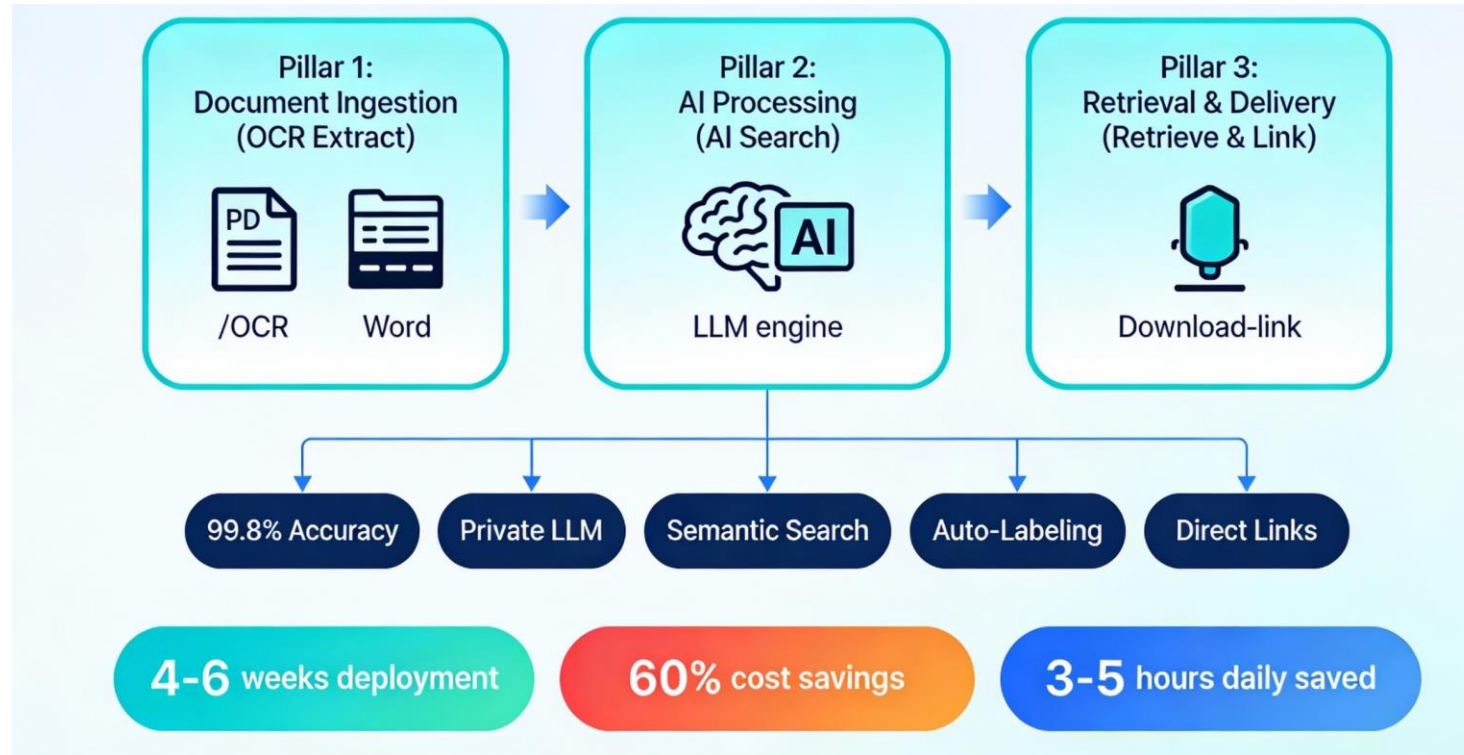
Multi-format processing with 99.8% accuracy, Vietnamese language support, batch processing

## Pillar 2: AI Processing (AI Search) :

Private LLM engine, semantic search, automatic classification, real-time response

## Pillar 3: Retrieve & Delivery :

Exact section identification, hyperlinked references, direct downloads, audit trails



# Three-Pillar Architecture

## Pillar 1: Document Ingestion (OCR Extract)

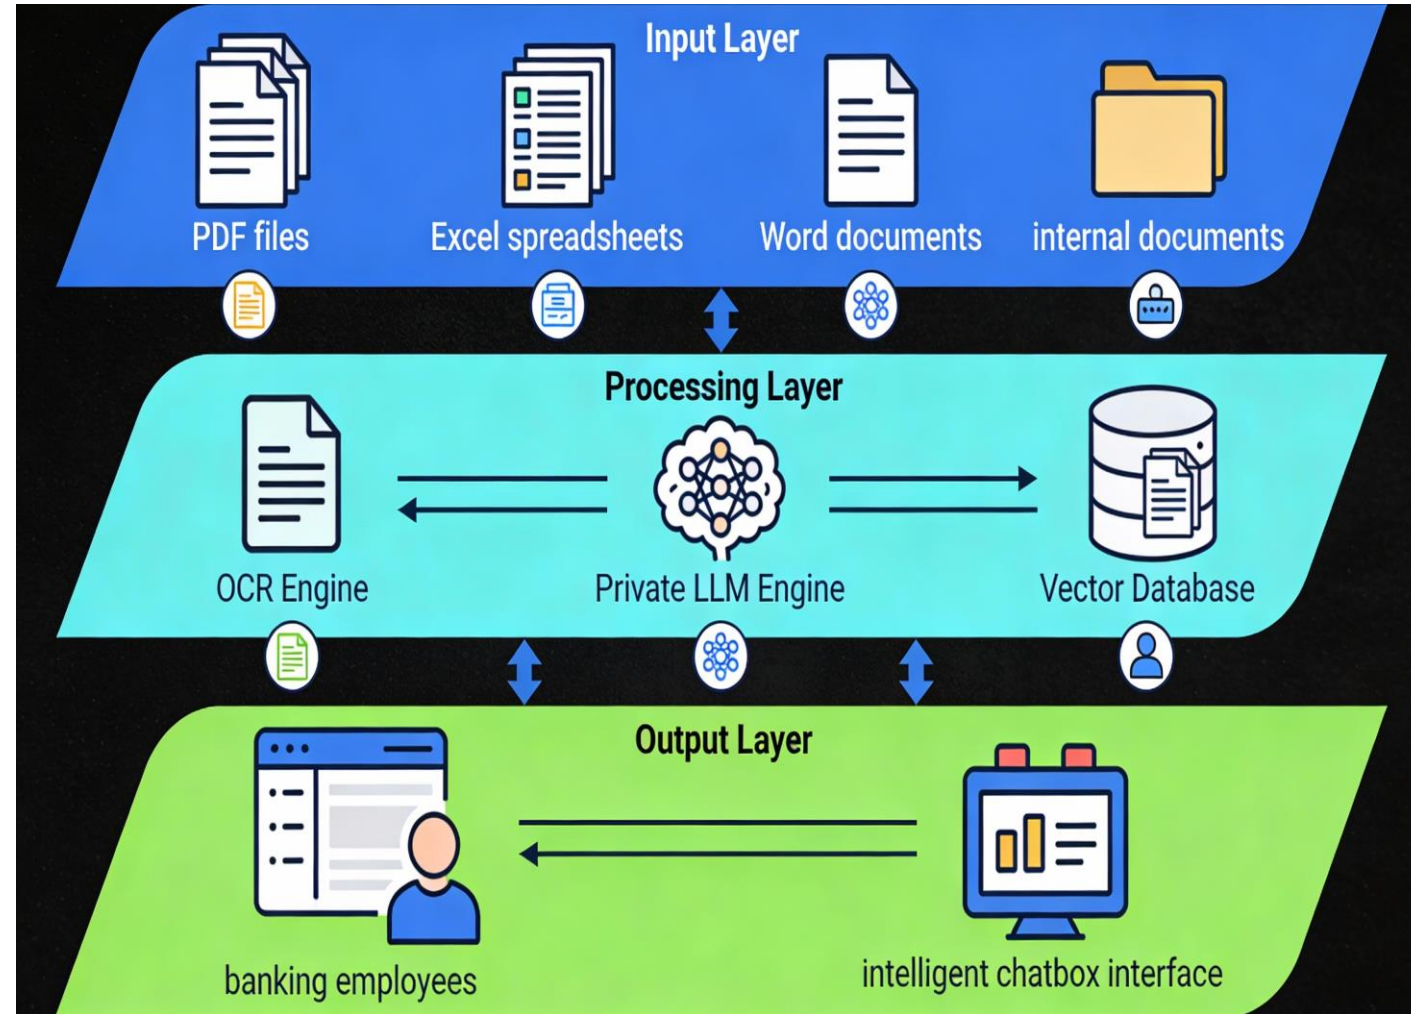
Multi-format processing with 99.8% accuracy, Vietnamese language support, batch processing

## Pillar 2: Document Ingestion (OCR Extract)

Private LLM engine, semantic search, automatic classification, real-time response

## Pillar 3: Retrieve & Delivery

Exact section identification, hyperlinked references, direct downloads, audit trails



# TECHNICAL SPECIFICATIONS

**OCR Extract:** 99.8% accuracy, multi-format, Vietnamese language, batch processing

**AI Search:** Private LLM, cost-optimized, semantic understanding, real-time

**Retrieve & Link:** Exact sections, hyperlinks, downloads, source attribution

**Model Tuning:** Banking-specific, SBV Circular trained, continuous improvement, zero additional cost

Query Response Time

**< 2 seconds**

Search Accuracy

**99.8%**

Concurrent Users

**500+**

System Uptime

**99.9%**

## OCR Extract



- ◇ 99.8% Accuracy
- ◇ Multi-format support (PDF, Excel, Word)
- ◇ Vietnamese language OCR
- ◇ Batch processing

## AI Search



- ◇ Private LLM Engine
- ◇ Cost-optimized model
- ◇ Semantic search
- ◇ Real-time processing

## Retrieve & Link



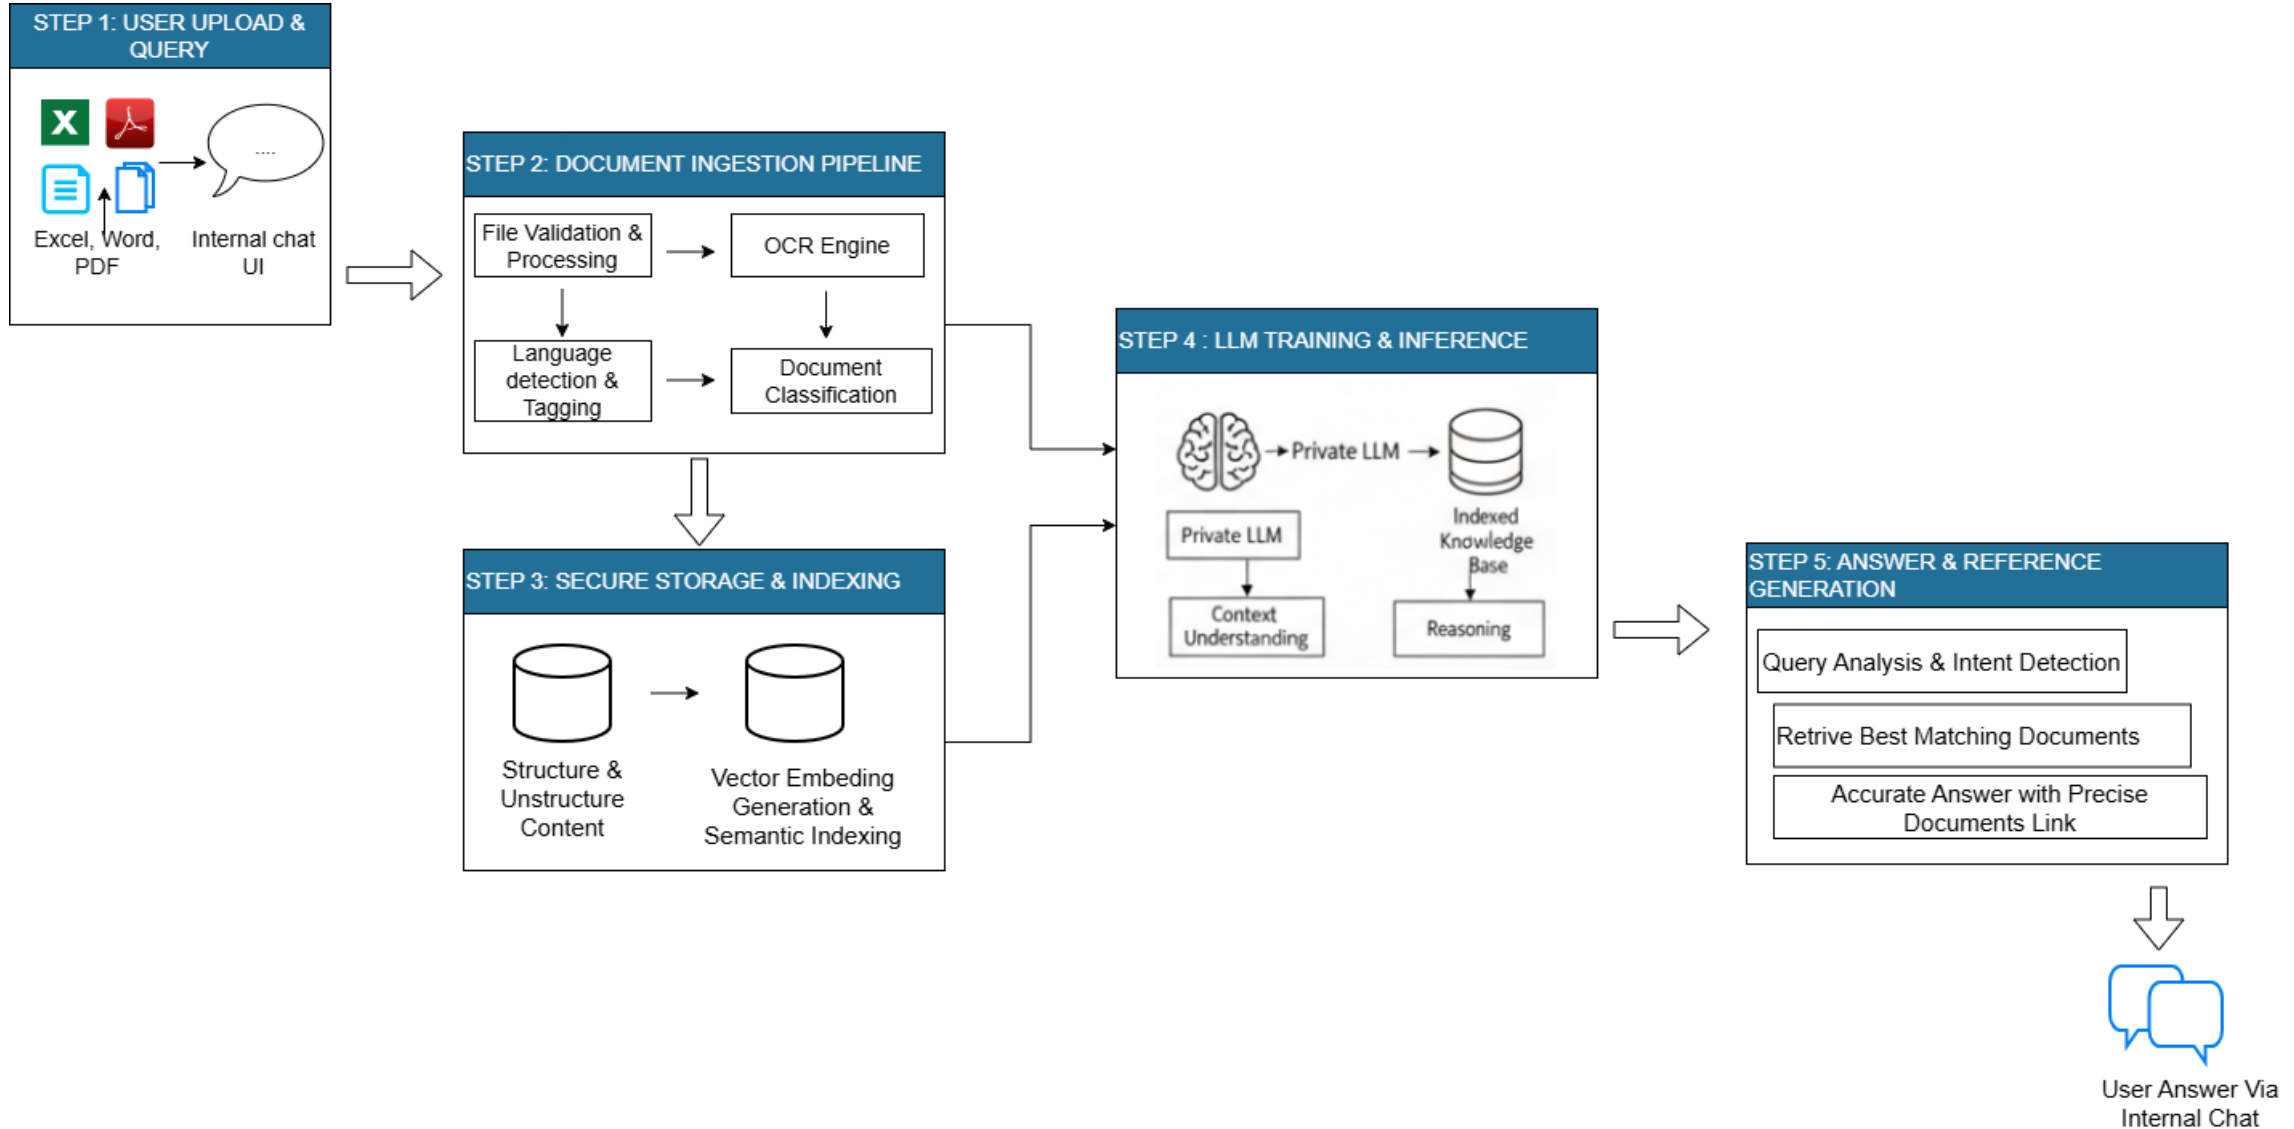
- ◇ Exact section identification
- ◇ Hyperlinked references
- ◇ Direct download links
- ◇ Source attribution

## Model Tuning



- ◇ Banking-specific training
- ◇ SBV Circular knowledge
- ◇ Continuous improvement
- ◇ Zero additional cost

# TECHNICAL DETAIL FLOW

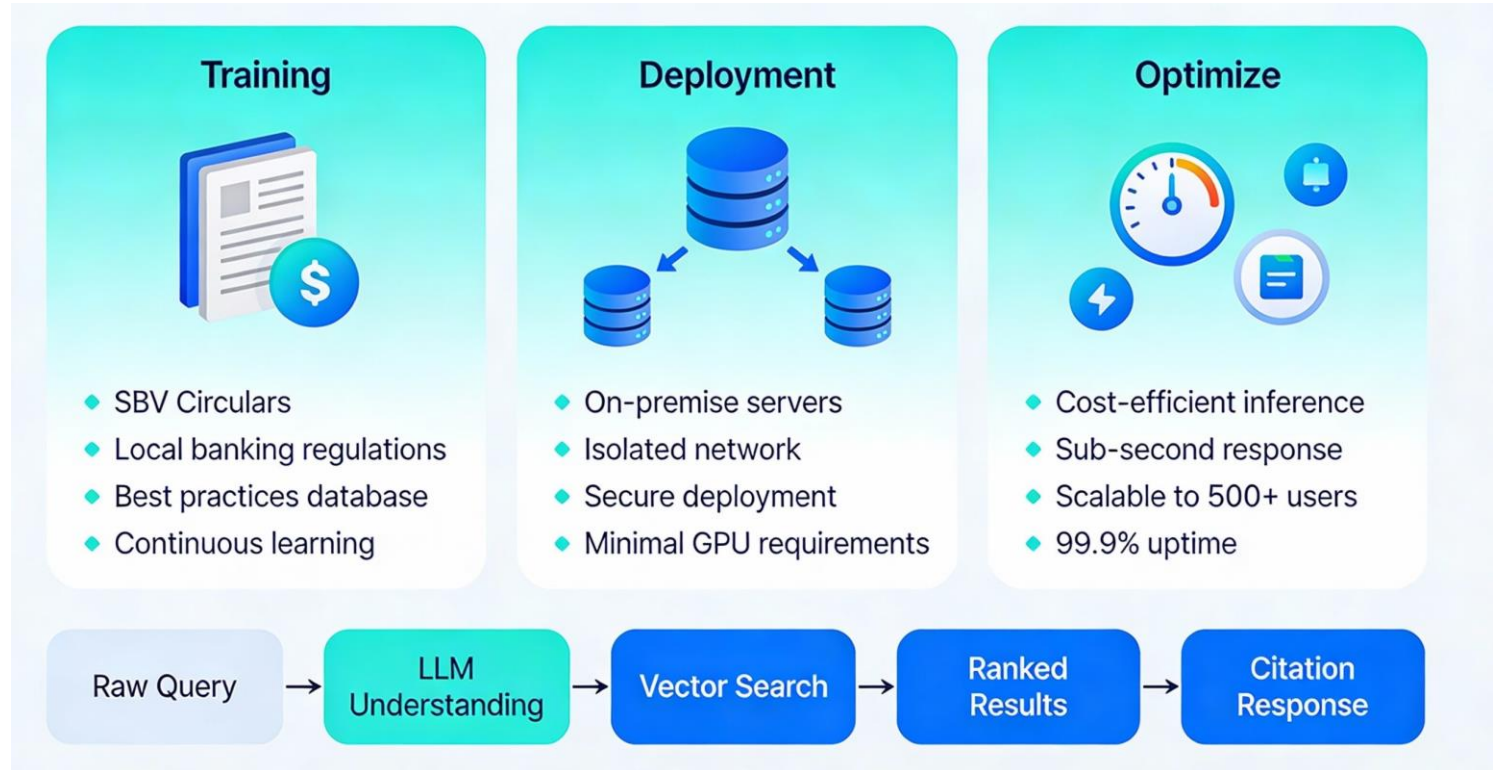


# LLM PROCESSING

**Training:** SBV Circulars, local banking regulations, best practices, continuous learning from new documents

**Deployment:** On-premise servers, secure isolated network, minimal GPU requirements, cost-efficient

**Optimize:** Sub-second response times, scalable to 500+ concurrent users, 99.9% uptime SLA



Processing Speed

**Sub-second**

Infrastructure Cost

**-60%**

# ROI & Financial Analysis

## Cost Reduction

### Infrastructure:

60% reduction in LLM-related costs

### Training:

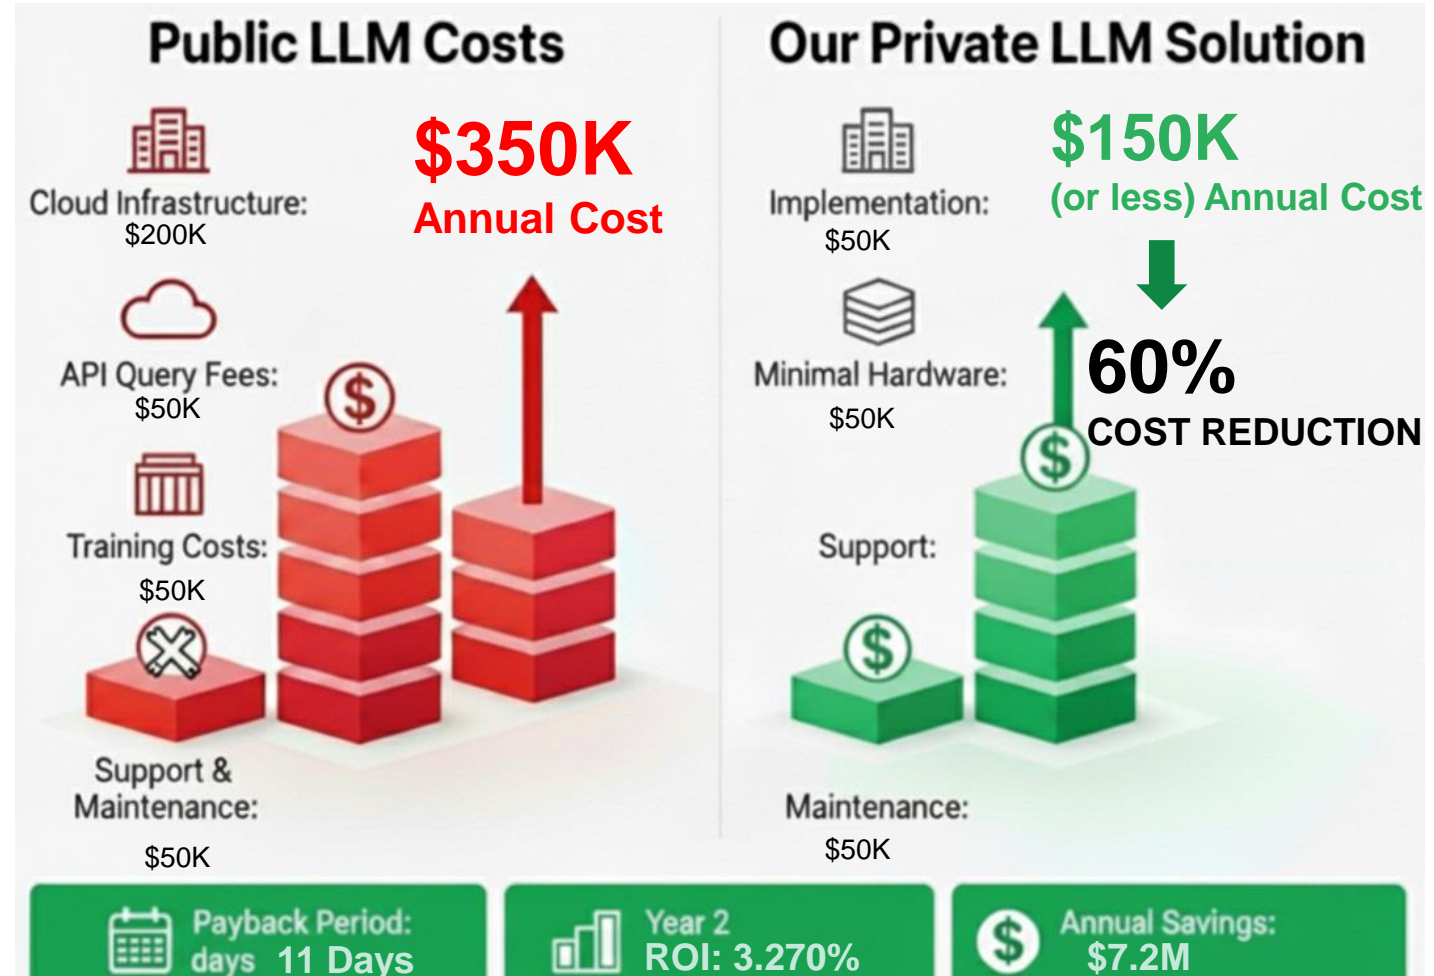
Eliminated need for external LLM subscriptions

### Maintenance:

80% reduction in document management staff workload

### Compliance:

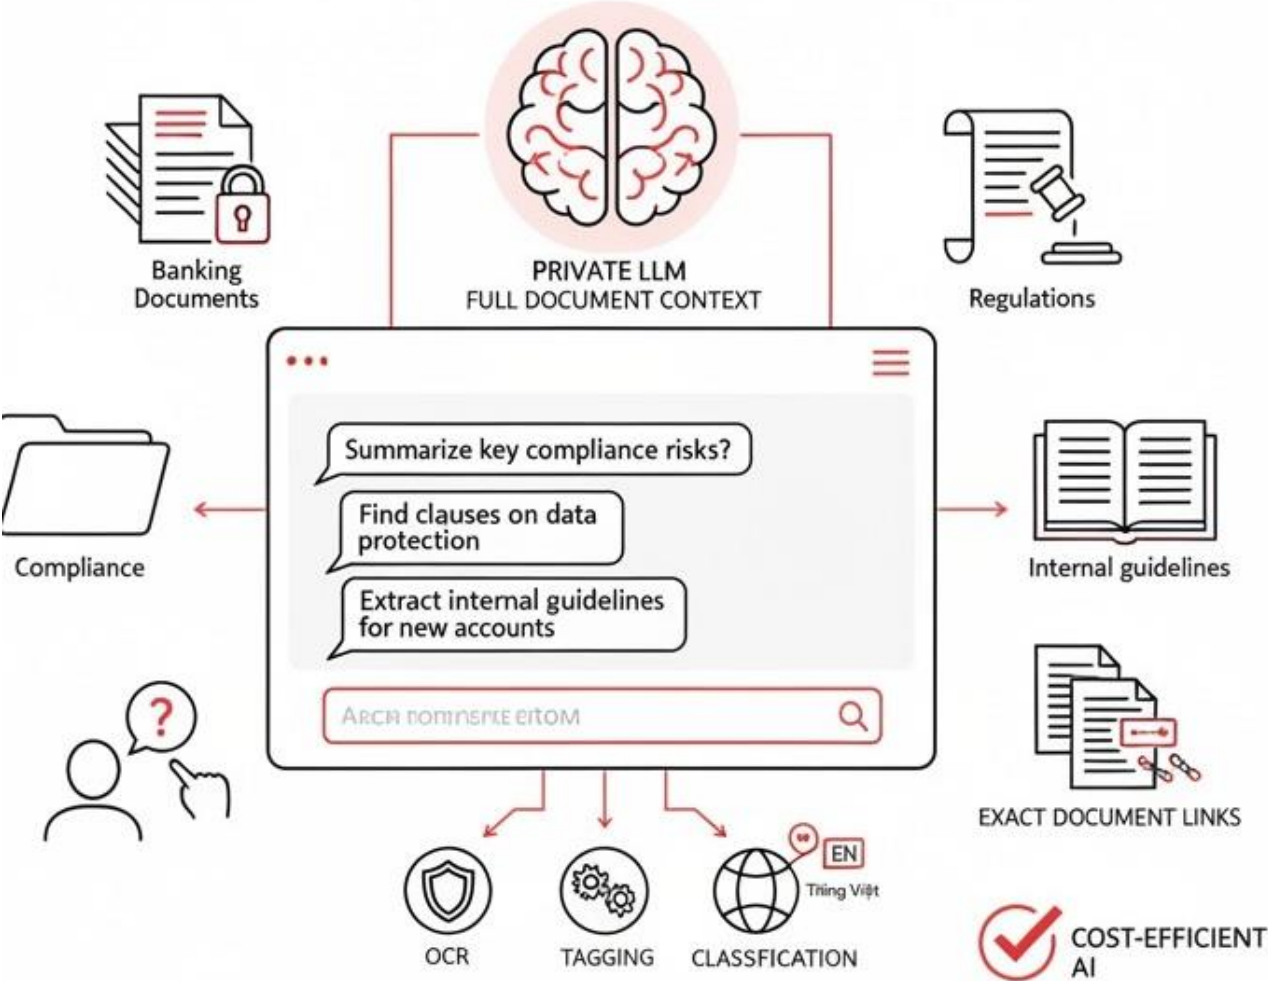
Reduced audit failures and regulatory fine risks



# BUSINESS USE CASE

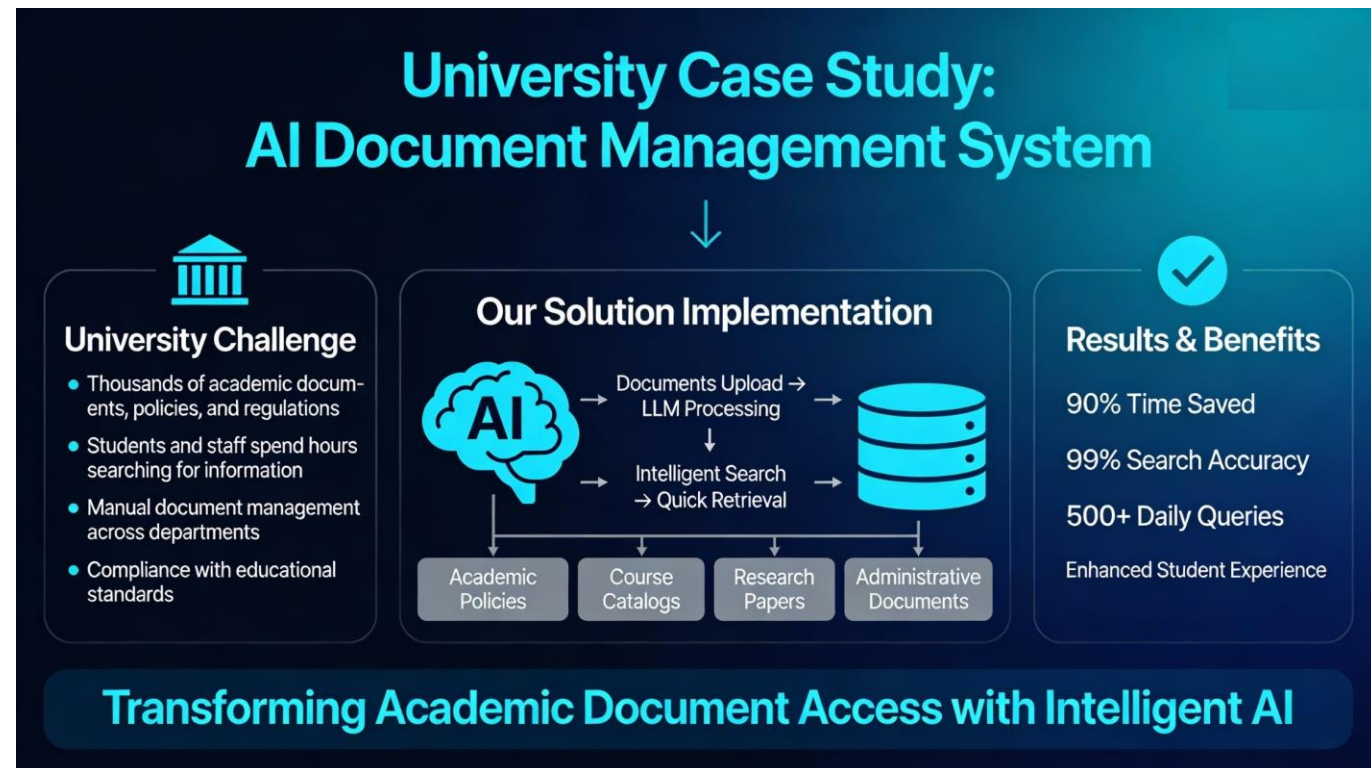


# AI ASSISTANT FOR BANKING



Transform scattered books, lecture notes, and research papers into a searchable AI Knowledge Base that supports lecturers, students, and administrators.

- ✓ Lecturers upload textbooks, slides, papers; AI summarizes and organizes content
- ✓ Students ask: "Please overview this chapter" or "Explain this theorem simply"
- ✓ AI provides context-aware summaries, definitions, and prerequisite concepts.
- ✓ Supports multi-language content (English/Vietnamese) and different faculties.



# WHY OUR SOLUTION

---



## Cost Efficiency

- *60% lower infrastructure costs compared to enterprise LLM solutions*
- *One-time implementation investment with minimal ongoing costs*



## Data Privacy

- *Private deployment -- all data remains on-premise*
- *Meets regulatory expectations for sensitive banking data*



## Multilingual

- *Understands Vietnamese banking terminology and context*
- *Optimized for Vietnamese language processing*



## Intelligent

- *Automatic document categorization across 10+ regulatory categories*
- *Reduces manual document organization workload*
- *Smart tagging system learns from user behavior*



## Auto Format

- *Remove unnecessary artifacts, normalizes spacing, indentation and alignment for readability*
- *Ensure compatibility with web publishing, database and CMS*



## Continuous Model Improvement

- *LLM updates based on new regulations automatically*
- *Learning system improves accuracy with usage*
- *Regular model retraining with latest banking guidelines*

## Our LLM Understands All Your Documents

### What Users Can Ask



**Please overview the content**  
LLM reads and summarizes full document context



**Which invoices do I need to pay?**  
LLM finds and summarizes relevant invoices instantly

### What Our LLM Does

1

Understands and stores all uploaded document content

2

Summarizes, searches, and reasons across regulations, contracts, and invoices

3

Always serves users with context-aware, accurate answers

### Why Our Solution is Different



Optimized model parameters



Editable token-level control



Hardware-optimized for cost and speed



Better revenue impact for banks

# THANK YOU!

---



*Tuan Hoang*



*tuanhx@hophan.com.vn*



*0902185151*

