ABBYY

PLAYBOOK

Next-Generation Document Automation: Combining Document AI and Generative AI

A strategic playbook for combining IDP and LLMs to build reliable, scalable workflows that deliver value, not risk

Table of Contents

AI is reshaping how companies manage information, decisions, and risk. Yet, applying the wrong kind of AI to document processing — particularly for business-critical workflows — can create more problems than it solves.

This playbook explores how business and operations leaders can combine **purpose-built Document AI with large language models (LLMs)** to drive higher accuracy, smarter automation, and measurable business value.

By designing workflows that **use the right AI at the right time**, organizations can optimize efficiency, ensure compliance, and enable future-ready operations.

ABBYY | Next-Generation Document Automation: Combining Document AI and Generative AI

The new AI landscape for documents

Leveraging the full potential of AI-powered business process automation and expanding beyond what was possible with rule-based automation means document processing needs two complementary types of intelligence: **intelligent document processing (IDP) and large language models (LLMs)**.

AI Type	Strengths	Limitations
Document AI (IDP)	Extracts structured, reliable, consistent data with high accuracy; understands document layouts (tables, forms).	Limited in higher-order reasoning, summarization.
Gen AI & Large Language Models (LLMs)	Excellent for interpreting, summarizing, or reasoning based on text.	Prone to hallucinations; struggles without structured inputs.

Smart leaders recognize:

LLMs are not replacements for structured extraction — they are enhancers when used correctly.

Why applying LLMs directly to documents fails

Common pitfalls when using LLMs alone for documents include:

Unstructured inputs ⟶ Chaotic, unpredictable outputs.

Raw documents (full of tables, mixed layouts, handwritten notes) confuse LLMs, leading to erratic or incorrect outputs.

Hallucinated facts ⟶ Risk of compliance and business errors.

LLMs are trained to generate plausible-sounding text, not to extract precise, verifiable data — risking major errors in business processes.

Inconsistent performance ⟶ Difficulty scaling with trust.

Small layout or wording changes can cause unpredictable results, making LLM outputs unreliable for at-scale operations.

Why applying LLMs directly to documents fails (cont.)

Common pitfalls when using LLMs alone for documents include:

Lack of explainability ⟶ Weakness in audit and regulatory environments.

LLMs often act as black boxes, making it difficult to trace how outputs are generated — a critical flaw for business-critical workflows, especially in regulated industries like finance, healthcare, and insurance.

Believing LLMs can replace OCR and IDP ⟶ Misplaced trust.

LLMs are not designed to handle document layout parsing (e.g., forms, tables, signatures) or precision data extraction tasks.

LLMs are powerful for content generation and reasoning— but they're not designed to extract structured facts with accuracy or consistency.

Building strong foundations with Document AI

Purpose-built Document AI specializes in:

Reading and parsing complex document layouts (tables, nested forms).

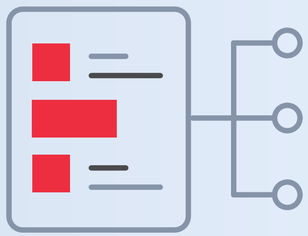

Extracting structured fields (amounts, dates, names, classifications).

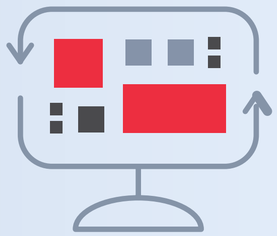

Validating data with business rules.

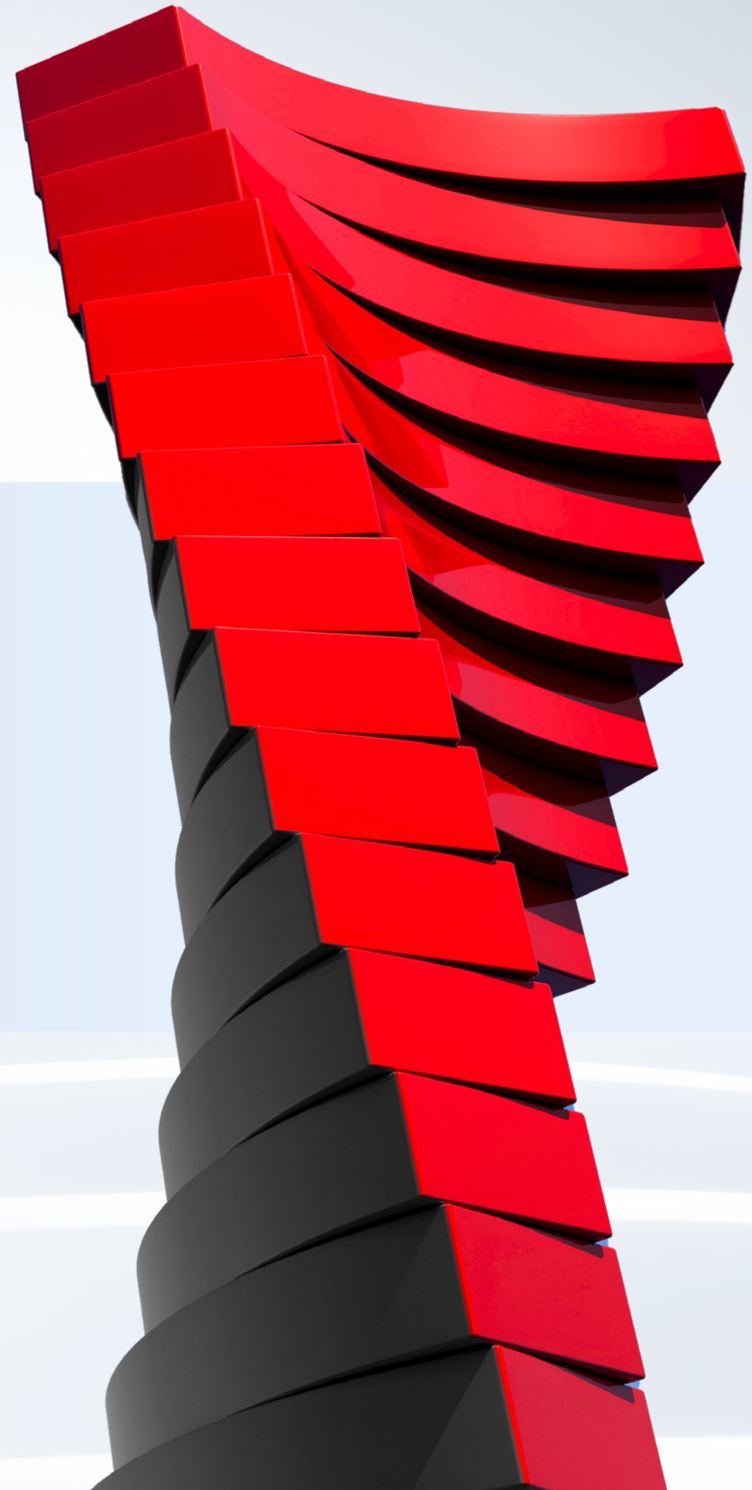

By using Document AI first, companies create a reliable, fact-based layer — the necessary foundation for any subsequent AI reasoning, summarization, or decision-making.

The power of pre-trained document processing models

When accuracy, speed, and trust matter, **pre-trained Document AI models** that are purpose-built for real-world business documents offer unmatched advantages over generic AI systems.

Optimized for business contexts

- Trained specifically on contracts, invoices, claims, forms, and other real-world documents — not internet text or general web content

End-to-end document processing

- Including classification, document splitting and assembly, data extraction, validation, human-in-the-loop, integration via API

Out-of-the-box intelligence

- Immediate availability of key capabilities:
  - Data normalization (e.g., date formats, currency)
  - Confidence scoring for extracted fields
  - Pre-defined validation rules for critical fields

Adaptable to business needs

- Easily configure and extend pre-built models:
  - Expand validation logic to meet industry-specific or company-specific compliance needs

Efficiency and speed

- Skip months of training or tuning
- Plug-and-play integration into automation workflows
- Accelerate time-to-value with minimal customization

Foundation for trusted automation

- Reliable, explainable outputs designed for high-stakes decision-making
- Perfect fit for agentic automation, RAG-based architectures, and scalable document workflows

Pre-trained Document AI models don’t just extract information — they deliver structured, validated, business-ready data that accelerates automation while minimizing risk.

Visit the ABBYY Marketplace to find pre-built models for your business process.

Combining forces — from extraction to reasoning

To fully leverage AI in document workflows, structure and reliability must come first.

This best-practice architecture ensures that automation is built on verified, business-ready data — not assumptions. By combining the precision of Document AI with the reasoning capabilities of LLMs, organizations can create smarter, safer workflows that scale reliably and deliver measurable business value.

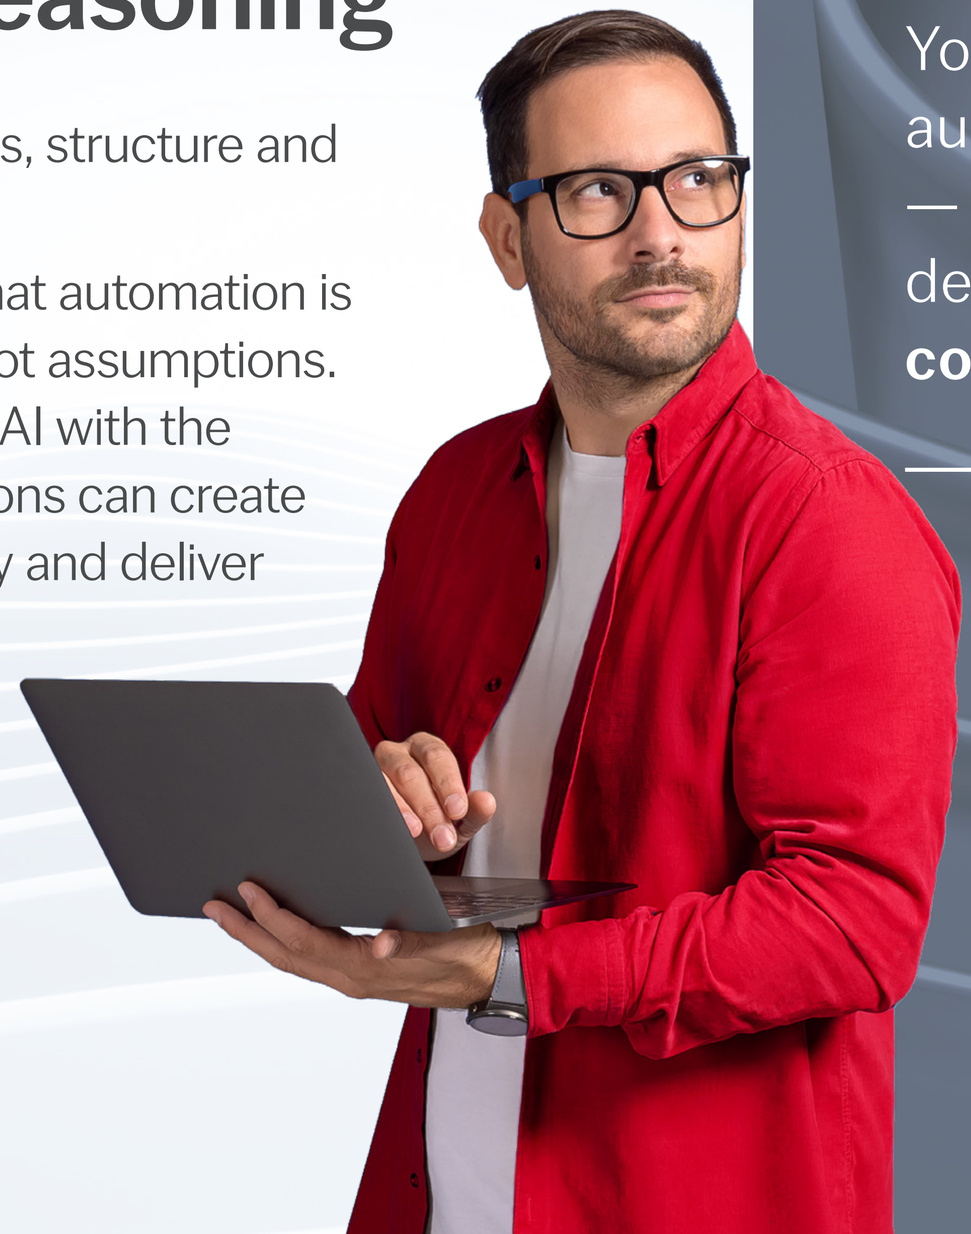

You don’t just automate tasks, — you automate decisions with **confidence**.

Best-practice architecture

Document ingestion

- Mobile
- eMail
- Shared folders (SFTP)
- API

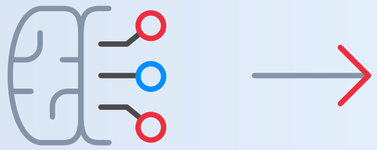

Document AI (intelligent document processing)

- Structured data extraction
- Validation and normalization
- Human-in-the-loop verification (if required based on confidence thresholds and business rules)

Validated data layer (acts as a “grounding” layer)

LLM reasoning (optionally via RAG)

- Summarization
- Interpretation
- Decision support

Business system integration

- ERP, CRM, case management
- Automated workflows
- Agentic systems

IDP + LLMs: The multiplier effect

Document processing is not a one-size-fits-all challenge. Purpose-built Document AI and general-purpose LLMs each excel in different areas — and when combined, they deliver far greater value than either could alone.

While Document AI offers precision, structure, and reliability, LLMs bring powerful reasoning and summarization capabilities. Together, they create smarter, more resilient workflows that balance accuracy with flexibility across the full document lifecycle.

	Purpose-built Document AI	General Purpose LLMs
Pros	• Highly optimized for document processing • Reliable, hallucination-free output • Faster and more efficient • Easier compliance and auditability • Pre-trained models for specific document types • Robust field-level validation and classification	• General model, no specific training required • Can reason and summarize complex documents • Fast setup for prototypes • Understands free-text and natural language
Cons	• New document types require training a new model • May struggle with highly unstructured document types	• Trained to generate content, not extract facts • May hallucinate or make confident errors (prompt & pray) • Poor understanding of document structure or layout • Lack of output consistency • Slower and more resource intensive • Unpredictable in regulated environments • Poor explainability for extracted data • May require post-processing for reliability • Challenges in optimizing for high-volume production

Ideal Use Cases

Purpose-built Document AI	General Purpose LLMs
• Precise, reliable extraction of data from documents • Structured and semi-structured document types • Segmentation of unstructured content • High-volume, repetitive documents • Workflows that require data validation (e.g. finance, compliance) • Regulatory, auditable processes • Feeding structured data to downstream systems (RPA,BPA, analytics or GenAI agents) • Regulatory/compliance-heavy processes • Where precision is critical	• Exploratory document analysis (prompting) • Reasoning and summarization of documents, but less suited for high-stakes accuracy • Conversational interfaces based on extracted data • Augmenting document processing by automating tasks that require reasoning and content generation (drafting replies, executive summaries, dispute resolution summaries, contract risk assessment, etc.)

LLMs need structured, reliable data to produce outputs that are accurate, consistent, and auditable.

Retrieval-augmented generation (RAG) with IDP

When LLMs are used in document processes, retrieval-augmented generation (RAG) architectures ensure better performance by:

- Feeding LLMs only validated, structured data
- Preventing hallucinations by grounding outputs in real business facts
- Supporting compliance and explainability

Think of it as teaching AI to reason only within the trusted context of your data — not invent answers.

Document AI: The foundation for trustworthy agentic automation

As organizations move from simple task automation toward agentic automation—where intelligent systems can plan, decide, and act autonomously—reliability becomes non-negotiable.

Structured fact extraction

Agentic systems are only as good as the data they process. Purpose-built models extract precise, validated information from complex documents, creating a trustworthy data foundation.

Reliable decision triggers

Agents rely on clear, structured outputs (not freeform text) to autonomously trigger business actions, decisions, or next steps.

Risk mitigation

When agents act autonomously, hallucinated or missing data can cause major operational or compliance failures. Document AI minimizes these risks by grounding agent decisions in validated facts.

In agentic automation, the quality of your decisions is only as good as the quality of your data.

Purpose-built Document AI ensures AI agents reason, decide, and act on a foundation of verified business facts.

Document AI: The foundation for trustworthy agentic automation (cont.)

Contextual understanding

Document AI captures the nuances of layout, relationships, and field-level meaning within complex documents, providing the rich context that agents require to reason effectively.

Compliance and auditability

In autonomous environments, traceability is essential. Document AI outputs are explainable and traceable, creating an auditable trail that supports governance and regulatory requirements.

Scaling safe autonomy

Without reliable document extraction, agentic automation remains risky and brittle. Document AI acts as the stabilizer, making it possible to scale autonomous document-centric workflows confidently. risks by grounding agent decisions in validated facts.

Get the Blueprint to AI Success: Building Intelligent and Agentic Automation That Works

Measuring and optimizing with process analytics

Operations leaders across industries face growing pressure to **prove the tangible value of automation initiatives—demonstrating measurable impact, optimizing continuously, and adapting quickly** to evolving business and regulatory demands.

To continuously improve, companies must see beyond outputs to operational performance.

Process Analytics for IDP enables:

- Visibility into throughput, accuracy, and bottlenecks
- ROI tracking on document automation initiatives
- Insights for process redesign and workforce optimization

You can’t improve what you can’t measure. Process Analytics for IDP turns automation into a measurable strategic advantage.

Real-world use cases

	Financial Services	Insurance	Healthcare	Legal
Purpose-built Document AI	Extract structured loan data	Capture structured claims information	Extract patient intake data	Parse contracts into structured fields
General-purpose LLMs	Summarize client risk profiles	Automate eligibility decisions with LLMs	Summarize treatment history for care teams	Flag risks and draft summaries with LLMs

Across industries, companies are not replacing processes — they are upgrading them by combining precision extraction with smart interpretation.

Customer Story | **RETAIL**

Generative AI + IDP extracts data from lease agreements with 82% accuracy

Challenge

30.000 lease agreements | **25** FTEs | **60%** Accuracy | **new** FASB* regulations

Keys to unlocking value

Combining IDP with Gen AI | Superior accuracy | Manual review feedback improving model

IMPACT

82%
Accuracy

20 FTE
savings through automating manual processes

350 DATA POINTS
compliance with FASB regulations

ABBYY | Next-Generation Document Automation: Combining Document AI and Generative AI

*Financial Accounting Standards Board

Strategic takeaways

Precision before interpretation:
Build workflows where factual extraction comes first.

Control reliability:
Use LLMs for what they do best — generating insights — but only on trusted foundations.

Optimize continuously:
Leverage analytics to refine, measure, and expand success.

Future-proof your operations:
Systems that combine reliability with flexibility will dominate in AI-driven markets.

The future of document automation isn't about picking between Document AI and LLMs — it's about strategically combining them.

By extracting facts first and using LLMs to reason based on those facts, businesses can drive smarter, safer, and faster outcomes — and capture true competitive advantage.

Slavena Hristova

Director of Product Marketing, Document AI

Slavena Hristova is a seasoned product marketing leader specializing in AI-powered intelligent document processing, OCR, and business process automation. As Director of Product Marketing at ABBYY, she drives the global strategy for the Document AI product line, shaping its market positioning, go-to-market execution, and customer adoption.

With deep expertise in product marketing and management, Hristova bridges the gap between technology and business needs, enabling organizations to harness AI-driven automation for smarter document workflows. Passionate about innovation and the evolving role of AI in enterprise automation, she brings a strategic and results-driven approach to transforming how businesses process and extract value from their data.

Discover what's possible with the industry's leading, purpose-built Document AI platform.

ABBYY

www.abbyy.com