



Laiye ADP: Document Processing in the Age of Agentic AI

CONTENTS

Chapter 1

Paradigm Shift in Document Processing4

- 1.1 The “Last Mile” of Enterprise Process Automation
- 1.2 Limitations and Bottlenecks of OCR and IDP
- 1.3 Large Language Models and Agents: The Paradigm Shift in Document Processing

Chapter 2

ADP - From Recognition to Understanding7

- 2.1 Defining ADP: The Birth of a New Paradigm
- 2.2 LLM ≠ ADP: The Relationship Between Operating Systems and Applications
- 2.3 Three Core Characteristics of ADP

Chapter 3

ADP: The First Step Towards Enterprise AI Agent14

- 3.1 Why Should Enterprise AI Agents Start with Document Processing?
- 3.2 From ADP to Enterprise AI Agents: A Progressive Evolution Path
- 3.3 ADP Suitability Self-Assessment: Is Your Enterprise Ready?
- 3.4 Organizational Change: Making Human-AI Collaboration Reality

Chapter 4

Architecture and Enterprise-Grade Capabilities of ADP23

4.1 Three-Tier Architecture Design

4.2 Performance Metric

4.3 Security and Compliance

4.4 Open Integration: Seamless Integration with Enterprise IT Ecosystem

Chapter 5

Industry Applications and Value Validation.....30

5.1 Manufacturing: Purchase Order Processing Automation

5.2 Multinational Enterprise: Global Invoice Processing Automation

5.3 Insurance: Claims Processing Automation

5.4 Banking: SME Loan Review Automation

5.5 Implementation Risks and Response Strategies

Chapter 6

Future Vision39

6.1 Short-Term Vision (2025-2026): Becoming Widely Adopted Enterprise Agentic AI Applications

6.2 Mid to Long-Term Vision (2027-2030): Becoming One of the Core Platforms for Enterprise Agentic AI

Begin Your ADP Journey41

Chapter 1 Paradigm Shift in Document Processing

1.1 The “Last Mile” of Enterprise Process Automation

Consider this scenario: a global manufacturing company’s finance department processes thousands of invoices every month from hundreds of suppliers around the world. These invoices arrive in multiple formats and languages. A team of more than ten finance professionals spends their days opening files one by one, extracting key data for verification, and manually entering it into enterprise systems.

Despite adopting various automation and AI tools, the company struggles with ever-changing supplier invoice templates that cause frequent recognition errors which ironically increasing the amount of manual rework required.

This situation is far from unique. Despite substantial investments in automation and AI, enterprises continue to face significant friction in what should be a straightforward task: processing documents.

According to IDC research, unstructured data now accounts for over 80% of all enterprise data and is growing at a 60% compound annual rate. By 2027, total global data volume is expected to reach 284.3 zettabytes, with unstructured data comprising 86.8%.

Document types such as contracts, invoices, purchase orders, and reports are the primary carriers of this unstructured information. Although these documents hold valuable business insights, most organizations lack the capability to efficiently extract and utilize them. This gap has become a major bottleneck in process automation, leaving substantial business value unrealized.

1.2 Limitations and Bottlenecks of OCR and IDP

Why do traditional technologies struggle to meet these demands? A closer look at OCR and IDP reveals structural limitations.

OCR (Optical Character Recognition) identifies text and its position with high accuracy on standard documents but stops at recognition. It reads words without understanding them. When faced with complex layouts such as multi-column text, graphics, or tables that span pages, OCR performance drops sharply.

IDP (Intelligent Document Processing) improves upon OCR by combining machine learning, natural language processing, and rule engines. This allows organizations to train models on their own document samples. However, as deployment scales, new challenges appear. Model training requires large volumes of labeled data, often dozens to hundreds of samples per document type. Any change in format requires extensive retraining and reannotation.

Enterprise environments add further complexity. A single company may handle hundreds of document types, and IDP's one-model-per-scenario approach increases project complexity, deployment time, and maintenance cost. As a result, most organizations limit IDP use to the top 10 to 20 percent of high-volume cases, leaving the rest untouched.

Multilingual content introduces additional cost because each language needs its own model. System fragmentation also reduces efficiency, as separate OCR or IDP systems are often deployed for invoices, identification documents, purchase orders, and contracts.

In the end, both OCR and IDP rely on memorization. They map fields to fixed templates instead of understanding the meaning, logic, and context of documents. Just as a student who memorizes answers cannot solve new problems, these systems cannot keep up with the complexity of enterprise document processing.

1.3 Large Language Models and Agents: The Paradigm Shift in Document Processing

Recent advances in artificial intelligence have redefined how organizations process documents.

Vision Language Models (VLM) allow AI to read complex layouts, tables, charts, and multilingual text while ignoring visual noise such as stamps or watermarks.

Large Language Models (LLM) build on this by understanding context, performing reasoning, and handling document tasks without custom training.

Agent technologies operationalize these models. They plan actions, select tools, and run document workflows automatically.

Together, these capabilities form **Agentic Document Processing (ADP)**, which eliminates templates and manual setup. ADP understands user intent and executes tasks independently.

Instead of building separate models for each document type, users can simply request, “Extract supplier, amount, tax, and line items.” The system adapts to any format or language.

Gartner estimates the intelligent document processing market will grow from 2 billion USD in 2024 to 5.2 billion USD by 2029, at nearly 29 percent annual growth. Organizations no longer need another OCR or IDP tool. They need a platform that truly understands documents and delivers complete automation.

Chapter Summary

Document processing is undergoing a major shift. It is moving from rule-based templates and sample-trained models to intelligent understanding powered by large language models. Traditional OCR and IDP rely on memorization. They follow templates and rules but cannot grasp document meaning or business logic.

Agentic Document Processing (ADP) changes this. AI becomes an agent that understands business context, learns autonomously, and executes tasks independently. This is more than technological improvement. It represents a step change in enterprise automation. Conquering the document “last mile” makes true end-to-end automation possible.

Chapter 2 ADP - From Recognition to Understanding

2.1 Defining ADP: The Birth of a New Paradigm

What is ADP?

Agentic Document Processing (ADP) is a next-generation platform that uses large language models, vision-language models, and agent technologies to deliver fully automated document processing.

ADP has three core capabilities:

General understanding

Powered by VLMs and LLMs, it can interpret multiple languages, formats, and scenarios without rules or annotations.

Autonomous planning and execution

It understands task objectives, plans steps, invokes tools, and completes complex tasks independently.

End-to-end business automation

It supports a complete workflow from document ingestion to business decisions, enabling human-AI collaboration.

ADP introduces two major shifts in document processing:

- **From recognition to understanding:** OCR reads characters, IDP extracts fields, ADP understands meaning and business context and executes autonomously.
- **From tools to agents:** Traditional systems are passive tools requiring manual annotation and repeated model training. ADP acts as a proactive agent that completes tasks based on objectives.

In short: **OCR recognizes characters, IDP extracts fields, ADP understands business.**

2.2 LLM $\neq$ ADP: The Relationship Between Operating Systems and Applications

Large language models provide the foundation for ADP, but they are not sufficient on their own. A useful analogy is that **LLMs act as an operating system** while **ADP is the application** built on top.

LLMs face three fundamental limitations in enterprise document processing:

- **Stateless operation:** LLMs do not retain information about documents they have processed, the data extracted, or exceptions encountered. Enterprises process varying document types and formats daily. Direct use of LLMs requires new instructions each time, which prevents contextual continuity.
- **No continuous learning:** LLMs cannot improve from real-world experience. Consistent errors or misidentified formats cannot be corrected automatically. Each task is processed independently, so LLMs alone cannot generate a data-driven improvement cycle.
- **Limited multi-step workflow handling:** Document processing often requires classification, content extraction, verification, calculation, and business execution. LLMs are effective at single-instance reasoning but cannot reliably execute such multi-step workflows.

ADP addresses these gaps. It builds on the understanding capabilities of LLMs while adding memory, learning, and orchestration. This transforms the raw understanding engine into a practical, enterprise-ready system capable of end-to-end document processing.

[Table 2-1: Essential Differences Between LLMs and ADP]

Dimension	LLMs	ADP
Positioning	General understanding engine	Dedicated application system
Memory	Stateless, independent processing each time	Stateful, maintains processing context
Learning	Fixed capabilities, cannot optimize	Continuous learning, improves with use
Orchestration	Single-step reasoning, lacks composability	Multi-step orchestration, end-to-end processes
Tools	Cannot proactively invoke external tools	Autonomously invokes APIs, systems, tools

The next section explores these capabilities in detail and highlights the core characteristics that make ADP a practical and scalable solution for enterprise document processing.

2.3 Three Core Characteristics of ADP

Characteristic 1: Zero-Shot Learning, Support 100+ Languages

Traditional IDP requires extensive annotated samples for each document type. ADP achieves zero-shot learning. Users simply describe requirements in natural language, for example, “Extract supplier name, invoice number, line items, and total amount from this batch of invoices.” The system executes immediately without training. When templates change, no reconfiguration is needed.

ADP supports over 100 languages, including English, Simplified and Traditional Chinese, Japanese, multiple European languages, and several Southeast Asian languages. It can process documents containing mixed languages correctly. This capability is essential for multinational enterprises and cross-border operations.

ADP also generalizes across multiple dimensions:

- **Format generalization:** Supports paragraphs, multi-column layouts, and tables.
- **Scenario generalization:** Adapts from a few examples to previously unseen cases.
- **Task generalization:** Handles not only extraction but also classification, comparison, review, and summarization.

Characteristic 2: Human-AI Collaboration Forms Data Flywheel

ADP creates a data flywheel through human-AI collaboration, enabling continuous improvement in document processing. The system uses confidence levels and business rules to determine when human intervention is needed, ensuring intelligent routing.

- **Confidence rules:** The system scores recognition results automatically. High confidence results (above 95 percent) pass without review. Medium confidence (85 to 95 percent) is flagged for quick human verification. Low confidence (below 85 percent) is routed for manual processing.
- **Business rules:** Users can define validation conditions, such as “order number must include letters and numbers.” Results that fail these rules are sent for review, even if the confidence score is high.

Human corrections feed back into the system automatically. ADP adjusts prompts or fine-tunes models based on this feedback, continuously improving performance. Unlike traditional IDP, which remains static after deployment and requires retraining to adapt to business changes, ADP evolves over time. This creates a self-reinforcing data flywheel and establishes a durable “data moat.” The table below illustrates the technological evolution from OCR to IDP to ADP.

[Table 2-2: OCR → IDP → ADP Technology Evolution Comparison]

Dimension	OCR [1st Generation]	IDP [2nd Generation]	ADP [3rd Generation]
Core Technology	Pattern recognition	Machine learning + rules	LLMs + agents
Understanding Level	Character recognition	Field extraction	Semantic understanding + reasoning + execution
Configuration Method	Template configuration	Sample annotation + training	Natural language instructions
Complex Layout Understanding Capability	No	Weak	Strong
Multilingual Support	Requires separate training	Supports few mainstream languages	100+ languages
Handwriting Recognition Accuracy	Weak	Weak	Strong
Annotation Cost	High (template configuration)	Very high (hundreds of samples per model)	Zero annotation
Continuous Improvement Capability	Weak (relies on manual template configuration)	Medium (requires retraining)	Strong (human-AI collaboration data flywheel)
Business Value	Digitization	Structuring	Automation + intelligence

The core differences between OCR, IDP, and ADP can be summarized in three levels:

OCR: Seeing characters

Recognizes text without understanding document structure or meaning.

IDP: Understanding words

Extracts fields using annotations and training, with limited ability to generalize to new formats or scenarios.

ADP: Understanding intent

Delivers end-to-end comprehension and execution without training, enabling zero-shot transfer across documents and tasks.

Characteristic 3: Agentic Orchestration - Enabling Multiple Models and Tools to Work Together

In modern enterprise workflows, no single model can handle all document processing tasks. For example, an invoice workflow may require a document classification model to identify invoice types, a vision-language model to interpret complex layouts, and a large language model to extract fields and reason semantically. Different documents and tasks demand different models.

Real-world workflows also rely on external tools. This can include querying ERP systems, accessing currency exchange APIs, performing calculations, or sending notifications. Traditionally, connecting models and tools requires manual coding or rules, which is costly and inflexible. Any process change often demands re-coding as well.

ADP overcomes this through agentic orchestration. ADP agents understand task objectives and autonomously select which models and tools to use, in which sequence. They manage intermediate results, monitor execution, and adapt when exceptions occur, delivering the required business outcomes.

Agent orchestration provides three key advantages:

- **Automatic multi-model collaboration:** Each model performs its specialized function. Vision-language models handle visual content, LLMs manage semantic reasoning, and the system selects the right model for each subtask.
- **Automatic chaining of models and tools:** Agents invoke APIs, query ERP systems, send notifications, and manage data flow without manual coding.
- **Adaptive execution:** Workflows adjust independently when exceptions arise, avoiding halted processes or errors.

Chapter Summary

The introduction of ADP signals the beginning of the agent era in document processing. It is not a simple upgrade of OCR or IDP, but a new paradigm that combines the understanding capabilities of LLMs with the autonomous reasoning of agents. ADP orchestrates multiple models and tools to handle complex workflows while forming a human-AI data flywheel that continuously improves performance.

Its three core characteristics, zero annotation learning, multi model agent orchestration, and a continuously evolving data flywheel, provide distinct competitive advantages. Together, they enable enterprises to move from recognition to understanding and from passive tools to proactive agents. ADP delivers more than efficiency gains. It represents a fundamental transformation in enterprise document processing capabilities.

Chapter 3 ADP: The First Step Towards Enterprise AI Agent

An enterprise AI agent is an intelligent system that understands business, makes autonomous decisions, and collaborates with humans. Implementing such agents requires a pragmatic entry point. ADP provides that entry point by delivering a focused, high-value application of enterprise AI.

3.1 Why Should Enterprise AI Agents Start with Document Processing?

Universality: Every Enterprise Has Documents

All organizations process large volumes of documents, including contracts, orders, invoices, approvals, and email attachments. Documents are the primary carrier of unstructured data, which accounts for over 80 percent of enterprise data, making document processing a strong foundation for ADP adoption.

Clear Pain Points: Labor-Intensive and Error-Prone

Document processing is repetitive, low-value, and error prone. Finance, administrative, and audit teams spend significant time on tasks such as invoice handling, approvals, and compliance reports. ADP addresses these inefficiencies directly.

Clear ROI: Quantifiable Value

Labor cost savings, efficiency gains, and error reduction are measurable. ADP delivers immediate value, unlike many AI projects that require long-term investment to show results. Chapter 5 details typical scenarios and business impact.

Controllable Risk: Progressive Deployment

ADP can be deployed in phases. Initially, humans review outputs. Next, high-confidence documents are processed automatically while others are routed for review. Eventually, over 90 percent of documents are processed autonomously, with a data flywheel continuously improving performance.

Technology Maturity: Production-Ready Today

ADP leverages mature LLMs, VLMs, and agentic technologies, supports on-premises deployment, and integrates via standard APIs. Gartner highlights document automation as a top enterprise AI application, and multiple successful deployments confirm its scalability and immediate impact.

3.2 From ADP to Enterprise AI Agents: A Progressive Evolution Path

Phase 1: Validation (1-6 months)

The first phase focuses on validating ADP in high-value scenarios and building trust.

01

Phase 2: Expansion (6-12 months)

Phase two extends ADP across scenarios and departments.

02

Phase 3: Deep Integration (12-24 months)

Phase three embeds ADP into core business processes and IT architecture.

03

Phase 1: Validation (1-6 months)

The first phase focuses on validating ADP in high-value scenarios and building trust. Typical pilots include invoice processing in finance, order or contract handling in procurement, or resume processing in HR. These scenarios offer clear value, standardized processes, and sufficient data volume.

Success is measured by accuracy above 90 percent, manual review below 20 percent, and ROI within 12 months. Enterprises confirm ADP capabilities, build confidence in AI, and lay the foundation for broader adoption.

Phase 2: Expansion (6-12 months)

Phase two extends ADP across scenarios and departments. This may involve processing additional document types within a department or replicating successful workflows in other departments.

Key initiatives include establishing an enterprise ADP platform, creating standards, and developing repeatable implementation methodologies. Expansion transforms ADP from a pilot into an enterprise capability, increasing coverage, efficiency, and value while accumulating best practices.

Phase 3: Deep Integration (12-24 months)

Phase three embeds ADP into core business processes and IT architecture. ADP integrates with ERP and CRM systems, supports end-to-end workflows such as Procure-to-Pay and Order-to-Cash, and becomes an essential intelligent node in operations.

Enterprises build document knowledge bases and accumulate enterprise-specific models. The human-AI data flywheel continuously optimizes performance, enhancing adaptability. ADP evolves from an optional tool to core infrastructure, providing a sustained competitive advantage.

3.3 ADP Suitability Self-Assessment: Is Your Enterprise Ready?

Before implementing ADP, enterprises should objectively assess their readiness and fit. This self-assessment tool evaluates three dimensions: business scenarios, technical environment, and organizational readiness (total: 100 points).

Assessment Dimensions and Scoring Criteria

I. Business Scenario Assessment (Maximum 50 points)

Business scenario assessment examines enterprise document processing scale, complexity, and business value, directly determining ADP implementation necessity and value potential.

[Table 3-1: Business Scenario Scoring Criteria]

Assessment Item	Scoring Criteria	Points	Score
Document Processing Volume (Monthly)	<500 documents	5	____/20 points
	500-2,000 documents	10	
	2,000-5,000 documents	15	
	5,000-10,000 documents	18	
	>10,000 documents	20	
Business Value of Document Processing	Non-core business, minimal error impact	3	____/15 points
	Affects business efficiency, moderate error impact	6	
	Linked to key business processes, significant error impact	9	
	Directly affects customer experience or compliance risk	12	
	Business bottleneck, errors may cause major losses	15	
Document Processing as % of Employee Work Time	<10%	2	____/10 points
	10%-30%	4	
	30%-50%	6	
	50%-70%	8	
	>70%	10	
Document Diversity	Single format, highly structured	1	____/5 points
	2-5 formats, semi-structured	2	
	6-10 formats, significant format variation	3	
	10-20 formats, including unstructured documents	4	
	>20 formats, including handwritten, images, multilingual	5	

Business Scenario Assessment Subtotal: ____/50 points

II. Technical Readiness Assessment (Maximum 25 points)

Technical readiness examines IT infrastructure maturity and data security compliance requirements, determining ADP implementation difficulty and deployment approach.

[Table 3-2: Technical Environment Scoring Criteria]

Assessment Item	Scoring Criteria	Points	Score
IT Infrastructure	Weak foundation, no professional IT team	0	____/10 points
	Basic IT environment, limited computing resources	4	
	Basic IT environment with adequate computing resources	6	
	Robust IT foundation with AI application experience	8	
	Dedicated AI team and comprehensive infrastructure	10	
System Integration Complexity	2 or fewer systems to integrate	5	____/5 points
	3-5 systems to integrate	3	
	More than 5 systems to integrate	1	
Data Security and Compliance	Low data security requirements, public cloud acceptable	6	____/10 points
	Moderate security requirements, hybrid cloud needed	8	
	High security requirements, on-premises deployment required	10	

Technical Environment Assessment Subtotal: ____/25 points

III. Organizational Readiness Assessment (Maximum 25 points)

Organizational readiness assessment examines executive support, employee acceptance, and change management capabilities, determining organizational resistance and implementation success probability.

[Table 3-3: Organizational Readiness Scoring Criteria]

Assessment Item	Scoring Criteria	Points	Score
Executive Support	Executives observing, no clear support	2	____/10 points
	Executives acknowledge but haven't allocated resources	4	
	Executives actively support with budget and team allocation	7	
	Executives deeply engaged, designated strategic priority	10	
Employee Acceptance	Strong employee resistance	0	____/10 points
	Employees neutral, willing to try with reservations	4	
	Employees positive, anticipating efficiency improvements	8	
	Employees proactively embrace AI, have prior experience	10	
Change Management Capability	Limited change management experience, weak communication	1	____/5 points
	Basic communication mechanisms, incomplete training systems	3	
	Established change management processes and training systems	5	

Organizational Readiness Assessment Subtotal: ____/25 points

[Total Self-Assessment Score: ____/100 points]

Assessment Results Interpretation

Based on total score, enterprises can determine ADP implementation readiness and recommended actions:

Total Score 80-100 points: Highly Recommended, Launch Immediately

Your enterprise is ideally suited for ADP implementation with all conditions for rapid success. Strong business demand, solid technical foundation, and robust organizational preparation indicate significant value generation in a short timeframe.

Action Recommendation: Launch immediately, selecting 2-3 high-value scenarios as initial pilots. Recommend a “parallel pilot + rapid expansion” strategy, launching simultaneously across multiple departments, validating quickly, then rolling out company-wide.

Total Score 60-79 points: Recommended, Start with Pilot

Your enterprise has a solid foundation for ADP implementation, though some aspects need strengthening.

Action Recommendation: Adopt a “single-point breakthrough + progressive expansion” strategy to minimize risk while validating value. Select one high-value scenario for pilot (such as finance invoice processing or procurement order processing), prioritizing high document volume, clear business value, and strong departmental support. Strengthen weak areas during the pilot, then expand progressively after demonstrating success.

Total Score Below 60 points: Not Yet Recommended, Strengthen Foundation First

Your enterprise may not be ready for ADP implementation currently. Recommend strengthening foundational conditions before reconsidering.

Action Recommendation: Identify weak areas from scoring and develop targeted improvement plans. If business scenario score is low (<25 points), map document processing workflows, identify core pain points, and evaluate business value. If technical readiness score is low (<12 points), upgrade IT infrastructure, provision GPU servers, or consider cloud deployment to lower barriers. If organizational readiness score is low (<12 points), strengthen organizational communication, secure clear executive support, and conduct employee AI awareness training. After completing preparation, start with small-scale pilots and progressively validate and adjust.

3.4 Organizational Change: Making Human-AI Collaboration Reality

Successful ADP implementation requires organizational readiness. ADP is not just a technology deployment but an enterprise transformation, shifting human work from repetitive tasks to high-value judgment-based roles.

Job Transformation: From Operation to Judgment

Post-ADP, employees focus on oversight, exception handling, and analysis. For example, a 15-person finance invoice team might reorganize into 2 AI trainers managing feedback and system optimization, 3 data reviewers handling low-confidence or exceptional documents, and 2 business analysts conducting process improvement using liberated capacity. The objective is to transition employees from operational work to judgment and decision-making roles.

Addressing Change Resistance

Employees may fear job loss. Enterprises must communicate that ADP augments work, freeing employees for more valuable tasks. Demonstrating case studies and benefits helps build trust.

Employees may worry about using AI. Provide comprehensive training, intuitive interfaces, and designate technology-savvy “AI Ambassadors” to assist colleagues.

Employees may fear skill depreciation. Redefine competencies from operational to judgment-based capabilities. For example, moving from entering invoice data to detecting anomalies and tracing root causes represents higher-level skill development.

CEO-Driven Change: Critical Actions

Leadership commitment is essential. Actions include: clearly articulating AI strategy and vision across the organization, explaining ADP implementation and its impact on employees, regularly reviewing progress and engaging frontline staff to address challenges, recognizing outstanding AI adoption through awards, and incorporating AI application performance into departmental KPIs. These steps ensure organizational alignment and sustained adoption.

Chapter Summary

ADP is more than a document processing tool. It is the optimal first step toward enterprise agentic AI, offering universality, clear pain points, measurable ROI, controllable risk, and mature technology.

Its evolution is structured: start with pilots, expand across departments, integrate deeply into workflows, and build agent networks, with clear objectives and success criteria at each stage.

Suitability assessment and organizational change guidance ensure readiness and smooth adoption. ADP lets enterprises validate AI value, develop capabilities, and advance incrementally toward intelligent operations. It addresses document processing challenges while initiating the journey to enterprise AI agents.


Chapter 4 Architecture and Enterprise-Grade Capabilities of ADP

In previous chapters, we explained why ADP is needed and what it is. This chapter focuses on technical aspects, showing that ADP is not only a product innovation but also an enterprise-grade platform that is engineering-validated, deployable, and trustworthy.

4.1 Three-Tier Architecture Design

ADP uses a three-tier architecture that separates models, agents, and business scenarios. This decoupling ensures scalability, maintainability, and flexibility. Models handle core understanding tasks, agents orchestrate workflows and tool usage, and business scenarios define objectives and rules. This design allows enterprises to expand capabilities, integrate with existing systems, and adapt to evolving document processing requirements without disrupting operations.

Architecture Overview



First Tier: Model Capability Layer - Multi-Model Integration and Optimization

The model layer forms ADP's core capabilities. It follows two principles: supporting integration of multiple VLMs and LLMs with different strengths and fine-tuning models for specific document processing tasks to improve performance.

Multi-Model Integration and Flexible Orchestration

No single model performs optimally across all document tasks. Some excel at complex tables, others in multilingual processing, or deliver efficiency at lower cost. ADP's model layer is an extensible pool integrating multiple VLMs and LLMs, including general-purpose, enterprise-proprietary, and custom fine-tuned models.

Upper-layer agents automatically select the most suitable model based on task requirements. For example, complex financial statements use table-optimized models, handwritten forms use handwriting-capable models, and multilingual contracts use cross-lingual models. This multi-model architecture allows ADP to process diverse and complex documents without being constrained by a single model's limitations.

The upper-layer agent orchestration system can automatically select the most appropriate model based on task characteristics. For example, when processing complex financial statements, it invokes models that excel at table understanding; when processing handwritten application forms, it calls models with strong handwriting recognition capabilities; and when handling multilingual contracts, it uses models with excellent cross-lingual abilities. This multi-model collaborative architecture enables ADP to handle various complex document processing scenarios without being limited by the capability boundaries of any single model.

Task-Specific Fine-Tuning and Performance Optimization

Pre-trained VLMs and LLMs provide strong general capabilities, but fine-tuning for enterprise-specific tasks improves accuracy and efficiency. ADP supports two optimization levels:

Prompt Optimization: The system analyzes enterprise documents and requirements to refine instructions, add domain knowledge hints, and structure outputs. This low-cost, quick method typically improves accuracy by 5–10 percent.

Model Parameter Fine-Tuning: VLMs and LLMs can be fine-tuned on enterprise data to significantly enhance performance for specific document types, reducing operational costs compared to general-purpose models.

Second Tier: Agent Orchestration Layer - Task Planning and Tool Invocation

The agent layer is the brain of ADP, responsible for task understanding, planning, and execution. Upon receiving a request, agents identify task objectives and constraints, decompose complex tasks into subtasks, select models and tools, and determine execution order.

For example, agents can invoke VLMs for parsing, LLMs for extraction, APIs for verification or calculation, and code for workflow execution. Agents continuously optimize performance using feedback from human-AI collaboration or business rules, creating a self-improving data flywheel.

Complex workflows leverage multiple collaborating agents. A classification agent identifies document types, a parsing agent extracts structure, an extraction agent retrieves key fields, a validation agent ensures data integrity, a business review agent executes logic and provides recommendations, and a coordination agent orchestrates the workflow. Together, these agents autonomously complete end-to-end document processing.

Third Tier: Business Application Layer - Scenario-Based Configuration and API Integration

The business layer delivers pre-configured, scenario-based applications while supporting flexible configuration and extension. ADP includes ready-to-use scenarios across multiple domains:

- **Finance:** invoice processing, expense review, bill reconciliation
- **Procurement:** order processing, receipt verification, supplier management
- **Contracts:** contract review, clause extraction, risk identification
- **Logistics:** waybills, warehouse and shipping documents
- **Financial Services:** loan applications, credit reports, bank statements
- **Healthcare:** medical prescriptions, lab reports, expense documents

The layer also provides RESTful APIs, webhook callbacks, supports batch and real-time processing, and integrates seamlessly with ERP, RPA, BPM, and other enterprise systems.

4.2 Performance Metrics

In practical enterprise deployments, ADP demonstrates strong performance even on a single on-premises node. Key metrics include document throughput, processing latency, multi-language support, and accuracy across various document types. The following benchmarks illustrate typical on-premises performance under real-world conditions.

[Table 4-1: Performance of Typical Document Processing Tasks (On-Prem Deployment)]

Processing Type	Per-Page Latency	Optimal Concurrency	Throughput
Document Parsing (Image → Markdown)	8-12 seconds	20	~7,000 pages/hour
Document Extraction (Image → Structured Data)	20-30 seconds	15	~2,500 pages/hour

ADP performance adapts to business scenarios. Batch processing prioritizes throughput with higher concurrency, while real-time processing prioritizes latency with lower concurrency. Multi-GPU deployments scale throughput linearly. On standard GPUs such as NVIDIA L20 or RTX 4090, a single server can support medium-sized enterprises.

[Table 4-2: Document Extraction Accuracy for Common Document Types (On-Prem Deployment)]

Document Type	Sample Size	Accuracy	Recall	F1-Score
Invoice Extraction	1,000 documents	92.3%	94.1%	93.2%
Purchase Order Extraction	800 documents	91.7%	93.5%	92.6%
Bank Statement Extraction	600 documents	94.2%	95.8%	95.0%
Contract Clause Extraction	500 documents	89.6%	91.2%	90.4%

Field extraction accuracy improves 10–20 percentage points over traditional IDP, driven by cross-format generalization for zero-shot transfer, support for 100+ languages, and robust handling of complex layouts including borderless tables, multi-column formats, and handwriting recognition.

4.3 Security and Compliance

Enterprise applications require strict security and compliance. ADP embeds security into its design, with **on-premises deployment** as a key feature. Both VLMs and LLMs can be deployed in customer data centers, ensuring all document processing and storage occurs locally. This approach fully protects enterprise data and eliminates external security concerns.

[Table 4-3: On-Prem Deployment Configuration Options]

Configuration	GPU Requirements	Performance	Applicable Scenarios
Standard Configuration	48GB (VLM) + 48GB×2 (LLM)	Optimal Performance	Large-scale Usage
Compact Configuration	48GB (VLM) + 48GB (LLM)	Balanced Performance	Medium-scale Usage
Minimal Configuration	48GB (VLM) + Customer's LLM	Basic Performance	Bring-your-own-LLM

Hardware Requirements: ADP runs on standard GPUs such as NVIDIA L20, RTX 4090, and RTX 3090, with CPUs of 48+ cores (>2GHz), 128GB+ memory, and 500GB+ storage. Deployment can be single-server with horizontal scaling and supports domestic operating systems and databases.

Data Security: Customer data is fully isolated with independent storage and access permissions. Development, testing, and production environments are separated with no cross-flow. All network transmissions use TLS and API calls enforce HTTPS.

Access Control: ADP supports multiple authentication methods including username/password, SSO, LDAP, and OAuth2.0, with optional MFA. API access uses token authentication with rotation. Role-based access control (RBAC) provides pre-configured and customizable roles, ensuring least-privilege access. All operations are logged, sensitive actions require secondary confirmation, and audit logs are immutable and regularly archived.

4.4 Open Integration: Seamless Integration with Enterprise IT Ecosystem

As an enterprise-grade platform, ADP possesses excellent openness and integration capabilities, enabling seamless collaboration with existing IT systems.

Bring-Your-Own-Models: ADP supports integration of enterprise-owned or third-party LLMs and VLMs alongside pre-configured models. The agent orchestration layer flexibly selects the most suitable model per task, allowing enterprises to leverage existing investments and optimize performance.

API-First Design: All core capabilities are exposed via RESTful APIs. Document parsing converts images to structured text, extraction pulls business data, classification identifies document types, and workflow APIs handle complex batch scenarios. APIs use JSON, provide multi-language code examples, support webhook callbacks, and SDKs for Python, Java, Node.js, and C# reduce integration effort.

Native RPA Integration: ADP integrates directly with Laiye RPA. Components can be dragged into workflows, automatically passing results to subsequent RPA tasks, supporting conditional logic, exception handling, and unified monitoring.

Enterprise System Integration: Standardized methods enable seamless ERP, CRM, and OA integration. Examples include automated invoice posting, order verification, accounts receivable reconciliation, contract archiving, sales order entry, expense claim review, and workflow routing. ADP can also act as a Model Context Protocol (MCP) Server for inter-agent collaboration, supporting enterprise-scale AI agent networks.

Chapter Summary

ADP's technical architecture balances advanced capabilities with engineering reliability. The three-tier design separates model capabilities, agent orchestration, and business applications, ensuring scalability. The model layer supports multi-model integration and task-specific fine-tuning, the agent layer enables autonomous workflow orchestration, and the business layer allows rapid scenario-based deployment. Performance and accuracy demonstrate effectiveness in real-world use, while comprehensive reliability measures ensure enterprise-grade stability. On-prem deployment safeguards data security and compliance. Open architecture and integration capabilities allow ADP to seamlessly connect with existing IT systems, functioning as an intelligent enhancement layer rather than an isolated platform.

Chapter 5 Industry Applications and Value Validation

Technology proves its value through measurable business outcomes. This chapter highlights four real-world industry scenarios showing how ADP drives process automation, boosts efficiency, reduces costs, and mitigates risks, closing the loop from AI capabilities to tangible business value.

5.1 Manufacturing: Purchase Order Processing Automation

Customer Pain Points

A global manufacturing enterprise processes 5,000 purchase orders monthly from multiple channels including email, chat, and customer systems. Orders arrive in Excel, PDF, and image formats with over 4,000 layout variations across eight languages. Key challenges include:

- **Efficiency bottlenecks:** Manual entry takes 5–10 minutes per order, requiring a 5-person team.
- **Quality risks:** Human errors in data entry, material code verification, and customer information cause inventory anomalies and customer complaints.
- **Resource misallocation:** 80% of the team's effort is spent on low-value data entry rather than analysis or customer development.
- **Expansion obstacles:** Multilingual orders, new customers, and new formats require retraining, and peak periods strain staffing capacity.

Solution

ADP delivers end-to-end automation across all order channels. Key fields such as order number, customer name, order date, product code, quantity, and amount are extracted through semantic understanding.

- **Human-AI collaboration:** Fields with low confidence trigger human verification for accuracy, forming a continuous data flywheel.
- **Automated review and validation:** Historical comparison, code conversion, and rule checks ensure correctness. Anomalies are routed for manual processing.

- **System integration:** Verified orders are automatically entered into order management systems via RPA and pushed to CRM for delivery.

This solution improves efficiency, reduces errors, and breaks down data silos, enabling the team to focus on high-value activities while continuously optimizing processing performance.

Value and Benefits

[Table 5-1: Manufacturing Order Processing Automation Comparison]

Metric	Before ADP	After ADP	Improvement
Order Processing Time	5-10 minutes/order	30 seconds/order	90% improvement
Daily Processing Capacity	60-80 orders/person	400-500 orders/person	6x increase
Entry Accuracy Rate	92%	98.5%	+6.5%
Required Personnel	5 people	1 person + system	80% reduction
Order Response Time	Average 4 hours	Average 15 minutes	94% improvement

“Previously, during peak periods, orders would pile up like mountains, with customers constantly calling to follow up. Now the system processes almost in real-time, and the business team can focus on customer communication and value creation.”

— Procurement Manager

5.2 Multinational Enterprise: Global Invoice Processing

Automation

Customer Pain Points

A multinational enterprise with operations in over 20 countries processes 2,000+ supplier invoices monthly across 15+ languages. Challenges include multilingual processing, highly varied invoice formats, complex three-way matching with orders and receiving documents, and multi-currency settlements. Manual processes cause delays, errors, and long approval cycles, impacting cash flow and supplier relationships. Traditional IDP solutions struggle with diverse layouts and limited invoice volumes.

Solution

ADP enables end-to-end automation. It identifies invoice language and format, extracts key data (supplier, invoice number, line items, amounts, taxes, currency), and performs currency conversion and intelligent three-way matching. Semantic understanding handles inconsistencies in descriptions, partial deliveries, and special items like discounts or shipping fees. Automated tiered processing approves perfect matches, flags minor discrepancies for quick review, and escalates major discrepancies for manual handling, maximizing automation while maintaining accuracy.

Value and Benefits

[Table 5-2: Multinational Enterprise Invoice Processing Automation Comparison]

Metric	Before Implementation	After Implementation	Improvement
Invoice Processing Time	Average 45 minutes/invoice	Average 3 minutes/invoice	93% improvement
Three-way Matching Accuracy	85%	96%	+11%
Approval Cycle	Average 15 days	Average 3 days	80% reduction
End-to-End Automation Rate	0%	75%	+75%

Beyond direct efficiency improvements, the system brings additional value, the most prominent being more timely payments to suppliers improve supplier relationships and create opportunities to secure more favorable procurement terms.

“What surprises me most is the system's multilingual processing capability. Our finance team finally doesn't have to struggle with different languages. More importantly, reducing the approval cycle from 15 days to 3 days gives us the opportunity to obtain more favorable procurement terms, which directly translates into tangible financial benefits.”

— CFO

5.3 Insurance: Claims Processing Automation

Customer Pain Points

An insurance company processes over 300,000 paper documents annually across new policies, policy services, and claims. Challenges include:

- **High manual workload:** Staff scan, classify, verify, and enter documents manually, leading to repetitive labor, high costs, and error risks.
- **Time-consuming processing:** Key tasks like identity verification, document review, and field-by-field data entry slow overall operations.
- **Handwriting recognition challenges:** Multilingual handwritten content (simplified/traditional Chinese, English) increases processing time and error risk.

Solution

ADP implements phased automation starting with 50+ high-frequency document types, covering 80% of key scenarios. The platform integrates ADP as the “brain and eyes” with RPA as the “hands and feet,” coordinated via an intelligent scheduling center. ADP automatically classifies documents, extracts key fields (customer and insurance information), and interfaces with business systems. Built-in validation replaces manual verification, ensuring accuracy and forming a fully automated document processing loop.

Value and Benefits

[Table 5-3: Insurance Claims Processing Automation Assessment]

Metric	Before ADP	After ADP	Improvement
Document Processing Method	Purely manual	Primarily automated	Substantial automation
Document Recognition Accuracy	90% (manual)	95% (manual + AI)	+5%
Document Processing Time	5 days x 8 hours	7 days x 24 hours	Significant time reduction
Risk Management	Manual sampling	AI full coverage	Significant risk reduction

"Insurance document processing has always been our pain point, especially handwriting recognition. The ADP system's handwriting recognition capability surprised us. More importantly, the system can automatically verify information consistency across different documents, which is very important for risk management."

— Operations Supervisor

5.4 Banking: SME Loan Review Automation

Customer Pain Points

A city commercial bank processes 500+ SME loan applications monthly, each containing 10–20 documents. Challenges include:

- **High workload:** Manual review takes 3–5 minutes per document, 1–2 hours per application, overloading a 15-person team.
- **Unstable quality:** Reviewer experience and fatigue lead to missed risks.
- **Poor customer experience:** Approval cycles of 5–10 days require repeated document submissions.
- **Compliance risks:** Review traceability is difficult, increasing credit and regulatory risk.

Solution

ADP automates end-to-end loan document review. It classifies documents, extracts key fields (identity, business, financial, and banking information), and performs cross-validation and business rule checks. Intelligent risk control detects fraud, ensures data consistency, recognizes abnormal patterns, and queries external data sources. Based on risk scoring, low-risk applications are auto-approved, medium-risk cases flagged for manual review, and high-risk applications rejected or sent for supplementary documentation. This workflow reduces manual labor, ensures compliance, and accelerates loan approval while mitigating credit risk.

Value and Benefits

[Table 5-4: Bank Loan Review Automation Comparison]

Metric	Before ADP	After ADP	Improvement
Document Review Time	1-2 hours/application	15-20 minutes/application	85% reduction
Approval Cycle	5-10 days	1-2 days	80% reduction
Review Personnel	15 people	6 people	60% reduction

“ADP not only helps us improve review efficiency but also enhances review quality. Many risk points that are easily overlooked manually can be accurately identified by the system, avoiding potential losses. This value cannot be measured in monetary terms.”

— Risk Control Manager

5.5 Implementation Risks and Response Strategies

Enterprises may face challenges during ADP implementation. Key risks and responses include:

- **Diverse document types / accuracy concerns:** Enterprises with hundreds of document types may worry about processing accuracy. **Response:** Apply the Pareto principle, ensure focus first on the 20% of document types covering 80% of volume. Set initial accuracy targets at 80–85% with manual review. The data flywheel gradually improves accuracy to 90%+ over 3–6 months.
- **Data security concerns:** Sensitive documents in finance, healthcare, or government may raise compliance worries. **Response:** Use private deployment with all models and data running locally. Sensitive fields (e.g., ID numbers, bank accounts) are automatically encrypted.
- **Complex system integration / long implementation cycles:** Integrating ADP with ERP, CRM, or other systems may seem daunting. **Response:** Use phased integration via APIs, RPA, and “shadow mode” deployment to run alongside existing systems, ensuring continuity. RESTful APIs enable typical integration and testing within 1–2 weeks.

Chapter Summary

Successful ADP application relies on three factors:

- **Select scenarios with clear pain points:** High document volumes, labor intensity, timeliness demands, and quality risks make scenarios ideal for ADP. All four case studies demonstrate these characteristics, enabling rapid impact.
- **Pursue end-to-end process automation:** ADP must integrate deeply with business workflows to form a complete automation loop. Manufacturing orders are recognized, reviewed, entered, and delivered automatically; multinational invoices undergo three-way matching and approval. Full automation maximizes business value.
- **Enable continuous optimization:** Models and rules are iteratively refined using real business data, forming a data flywheel. Feedback from manual calibration or phased document coverage ensures sustained accuracy and long-term value.

Chapter 6 Future Vision

The emergence of ADP marks the entry of document processing into the agentic era, but this is just the beginning. This chapter explores ADP's future evolution: expand applications in the short term, build platform in the mid-term, and build ecosystem in the long term.

6.1 Short-Term Vision (2025-2026): Becoming Widely Adopted Enterprise Agentic AI Applications

Product Vision: Focus on core enterprise business scenarios and establish replicable solutions.

At this stage, ADP will focus on **adoption of agentic document processing applications:**

Focus on High-Value Scenarios:

- **Accuracy Improvements:** Achieve 95%+ accuracy in scenarios such as invoices, orders, and contracts
- **Industry-Specific Solutions:** Develop specialized agentic document processing solutions for finance, manufacturing, healthcare, and other industries
- **Multilingual Globalization:** Zero-shot processing capability for 100+ languages, supporting global deployment for multinational enterprises

Validate Business Models:

- **Quantifiable ROI:** Achieve 6-12 months investment return in typical scenarios
- **Scalable Replication:** Establish standardized implementation methodology to rapidly replicate successful experiences
- **Initiate Data Flywheel:** Enable continuous system learning and optimization through human-AI collaboration

User Value: Make document processing agents a standard tool for enterprise cost reduction and efficiency improvement, achieving the leap from “pilot” to “scaled adoption.”

6.2 Mid to Long-Term Vision (2027-2030): Becoming One of the Core Platforms for Enterprise Agentic AI

Product Vision: As an agentic document processing platform, provide document processing capabilities for the enterprise agentic AI ecosystem.

At this stage, ADP will transform from a **tool** to a **platform**:

Focus on Building Platform and Ecosystem:

- **From Standalone Agent to Capability Foundation:** ADP becomes the “eyes” of enterprise agents, providing document understanding capabilities for sales, finance, procurement, and other business agents
- **Capability Openness:** Expose core capabilities through API and open source, enabling ISVs and system integrators to build industry solutions on ADP
- **Developer Community:** Provide documentation, examples, and best practices to lower integration and customization barriers

User Value: Equip every business agent with powerful document processing capabilities, accelerating enterprise intelligence evolution from “point applications” to “comprehensive collaboration.”

Conclusion

From tool to platform, from platform to ecosystem—this is ADP’s evolutionary path. In the short term, expand applications to make ADP widely used in core scenarios; in the mid to long term, build platform and ecosystem to make ADP the capability foundation for enterprise agents, driving the document intelligence industry toward maturity and standardization.

We look forward to creating this future together with our customers and partners.

Begin Your ADP Journey

If you wish to learn more about how ADP can help enterprises achieve intelligent transformation in document processing, or if you want to experience ADP's powerful capabilities firsthand, we cordially invite you to contact Laiye Technology.

What We Offer

Product Demonstration:

Our professional team will showcase ADP's effectiveness in real business scenarios, including comprehensive demonstrations of core use cases such as invoice processing, contract management, and order recognition.

01

Product Trial:

Based on your business needs, we will configure a dedicated trial environment, allowing your team to test ADP's processing capabilities and accuracy with actual documents.

02

Solution Consultation:

Senior solution experts will gain a deep understanding of your document processing challenges and business scenarios, then customize an ADP implementation plan tailored to your needs, including technical architecture, deployment strategy, and return on investment analysis.

03

Best Practice Sharing:

Share successful cases and implementation experiences from various industries to help you avoid common pitfalls and accelerate project delivery.

04

About Laiye

Laiye is leading the transformation of enterprise operations through intelligent automation and agentic AI. As a global innovator, we empower organizations to close the gap between strategy and execution by leveraging digital workers to deliver measurable business outcomes. With offices in Beijing, Shanghai, Shenzhen, Hong Kong, Singapore, and other locations, Laiye serves enterprises across finance, manufacturing, energy, retail, pharmaceuticals, and more, providing solutions that generate tangible value.

At the core of this transformation is ADP, Agentic Document Processing, a next-generation platform that combines large language models and agentic technology to turn labor-intensive and error-prone document workflows into intelligent and adaptive processes. Hundreds of enterprises have achieved substantial gains in efficiency, accuracy, and cost reduction while unlocking the potential of data-driven insights to support critical business decisions.

The future of agentic document processing extends beyond workflow automation. By integrating scalable AI agents with human-AI collaboration, enterprises can continuously optimize their processes, accelerate outcomes, strengthen operational resilience, and create lasting value at scale. Laiye invites enterprises, partners, and stakeholders to join in shaping this future, harnessing ADP to address operational challenges today and build the foundation for enterprise intelligence tomorrow.

The era of intelligent agentic automation is here. Together we can transform document processing from a routine operation into a strategic capability, driving operational excellence, unlocking new business value, and enabling enterprises to thrive in an increasingly digital and data-driven world.

LAIYE



Website:<https://laiye.com>

Email: mkt@laiye.com

Contact us:

Scan the QR code to start the conversation now.