




Laiye ADP: Document Processing in the Age of Agentic AI

CONTENTS

Chapter 1:

Paradigm Shift in Document Processing 4

- 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 Understanding 8

- 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:

The First Step Towards Enterprise AI Agent 15

- 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 ADP

..... 25

4.1 Three-Tier Architecture Design

4.2 Performance Metrics

4.3 Security and Compliance

4.4 Open Integration: Seamless Integration with Enterprise IT Ecosystem

Chapter 5:

Industry Applications and Value Validation 33

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 Vision 43

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 Journey 45

Chapter 1:

Paradigm Shift in Document Processing

1.1 The "Last Mile" of Enterprise Process Automation

Consider this scenario: a multinational manufacturing company's finance department processes thousands of invoices monthly from hundreds of global suppliers. These invoices arrive in different languages and formats. A team of ten-plus finance professionals spends their days opening these documents one by one, locating key information for verification, and manually entering data into the system. Despite trying various automation and AI tools, the company finds that constantly changing supplier invoice templates lead to frequent recognition errors, actually increasing the manual correction workload.

This is far from an isolated case. Despite massive investments in automation and AI, enterprises remain hindered by what appears to be a fundamental task: document processing. According to IDC research, unstructured data now comprises over 80% of total enterprise data and is expanding at a 60% compound annual growth rate. By 2027, global data volume will reach 284.30 ZB, with unstructured data accounting for 86.8%. Documents—contracts, invoices, orders, reports—serve as the primary carriers of this unstructured data. While these documents contain critical business intelligence, the lack of effective processing capabilities has created a major bottleneck in business process automation, leaving significant value untapped.

1.2 Limitations and Bottlenecks of OCR and IDP

Why can't traditional technologies meet these demands? A deeper analysis of OCR and IDP reveals critical limitations and bottlenecks.

OCR (Optical Character Recognition) technology identifies text and its position in documents, achieving accuracy rates above 90% on standard printed materials. However, OCR's fundamental limitation is clear: it can only "recognize text," not "understand documents." When confronted with complex layouts—multi-column formatting, mixed text and graphics, borderless tables, or cross-page tables—OCR becomes ineffective.

IDP (Intelligent Document Processing) builds upon OCR by incorporating machine learning, NLP, and rule engines, enabling customers to train their own models by

labelling document samples. However, when enterprises attempt to scale IDP deployments, they quickly encounter new bottlenecks.

First, IDP model training costs are prohibitively high. Training IDP models demands substantial annotated data: each document type requires dozens to hundreds of annotated samples, with significant time and cost implications. When document formats change, extensive re-annotation is needed.

Second, enterprise document environments are far more complex than anticipated: a single organization may encounter hundreds of document types. IDP's "one model per scenario" approach results in high project complexity, extended implementation cycles, and substantial maintenance overhead. From an ROI standpoint, enterprises typically limit IDP deployment to only the top 10-20% of high-frequency scenarios, leaving the majority of use cases unaddressed.

Third, multilingual support remains inadequate. Different languages require separate IDP models. For the multilingual supplier invoices mentioned earlier, IDP necessitates one model per language, significantly increasing both costs of ownership and maintenance.

Fourth, system fragmentation creates interoperability challenges. To address diverse requirements, enterprises often deploy multiple document processing systems: OCR system A for invoices, OCR system B for identification documents, IDP system C for purchase orders, IDP system D for contracts. This creates substantial challenges for system integration, compatibility and maintenance.

Whether OCR or IDP, both face a fundamental common limitation: they rely on "memorization"—remembering that specific fields appear in specific positions, that certain templates correspond to particular extraction rules. What enterprises truly need is "understanding"—comprehending document semantics, grasping business logic, and recognizing contextual relationships. Just as we cannot expect a student who merely memorizes answers to solve problems, OCR and IDP systems cannot adequately address enterprises' dynamic and complex document processing requirements.

1.3 Large Language Models and Agents:

The Paradigm Shift in Document Processing

Over the past few years, several critical technological inflection points have converged, catalyzing a paradigm-level transformation in document processing.

Vision Language Models (VLM) represent a breakthrough, enabling AI to truly "comprehend" complex documents for the first time. VLMs can understand visual document structure (layout, tables, charts), recognize multiple languages, handle complex formatting such as cross-page and multi-column layouts, effectively filter interference factors like stamps and watermarks, and support borderless table recognition. **Large Language Models (LLM)** demonstrate remarkable general-purpose capabilities. Pre-trained on massive text corpora, LLMs possess broad understanding of world knowledge and language structure, can perform reasoning and logical inference, and handle diverse document processing tasks without task-specific training. **Agentic Technologies** have transformed LLMs capabilities into production-ready solutions. Through task planning, tool invocation, and multi-step reasoning, agents can orchestrate complex end-to-end document processing workflows.

The convergence of these technologies has spawned an entirely new document processing paradigm—**Agentic Document Processing (ADP)**. ADP is a next-generation solution built on LLMs and agentic technologies. Rather than being a "tool" requiring template configuration or sample annotation, ADP functions as an "agent" capable of understanding business requirements and autonomously planning and executing tasks. Traditional document processing systems are "tools": users must explicitly specify "how to do it" (How). ADP is an "agent": users simply specify "what to do" (What), and the system autonomously understands, plans, and executes.

Consider invoice processing: the traditional IDP approach requires preparing annotated samples for each invoice type, training multiple models to recognize formats from different suppliers, configuring extraction rules for fields like amounts, tax rates, and line items, then repeating this entire process when suppliers change their templates. With ADP's agentic approach, users simply articulate requirements in natural language: "Extract the supplier, amount, tax, and line item details from this batch of international invoices." The system automatically recognizes document types and languages, extracts required information, and requires no configuration adjustments when new formats emerge.

According to Gartner, the global intelligent document processing market will grow from approximately \$2 billion in 2024 to \$5.2 billion in 2029, representing a compound annual growth rate approaching 29%. While over 100 vendors have entered the IDP market globally, what enterprises truly need is not merely "better OCR" or "smarter IDP," but a next-generation solution capable of truly understanding documents and delivering end-to-end automation—precisely where ADP's value proposition lies.

Chapter Summary:

Document processing is undergoing a paradigm shift—from rule-based template recognition, through sample-based model training, to intelligent understanding powered by LLMs. Traditional OCR and IDP technologies both face a fundamental chasm "from memorization to understanding": they can only remember templates and rules but cannot truly comprehend document semantics and business logic. ADP's emergence marks document processing's entry into the agent era: AI is no longer merely a tool requiring careful training and configuration, but an agent capable of understanding business context, learning autonomously, and executing collaboratively. This represents not just technological advancement, but a quantum leap in enterprise automation capabilities. When documents—this "last mile"—are finally conquered, true end-to-end enterprise automation becomes achievable.

Chapter 2: ADP - From Recognition to Understanding

2.1 Defining ADP: The Birth of a New Paradigm

What is ADP?

ADP (Agentic Document Processing) is a next-generation platform that leverages large language models and vision-language models, combined with agentic technologies, to achieve end-to-end automated document processing.

This definition encompasses three key elements:

General understanding powered by VLMs and LLMs

Without relying on rules or annotations, it possesses general understanding capabilities across multiple languages, formats, and scenarios

Autonomous planning and execution through agents

Capable of understanding task objectives, autonomously planning steps, invoking tools, and completing complex tasks

End-to-end business automation with human-in-the-loop

Forming a complete closed loop from document injection through business decisions to human-AI collaboration

ADP represents two fundamental shifts in the document processing field. The first shift is from recognition to understanding. OCR and IDP approaches recognize characters, extract fields, and output data - a mechanical processing method; whereas ADP's approach understands semantics, understands business, and executes autonomously - an intelligent processing method. The second shift is from tools to agents. Traditional systems are passive tools that require extensive manual annotation and repetitive model training; whereas ADP is a proactive agent that can autonomously complete tasks based on objectives.

"OCR recognizes characters, IDP extracts fields, ADP understands business."

2.2 LLM ≠ ADP: The Relationship Between Operating Systems and Applications

The universal understanding capability provided by LLMs is the foundation of ADP, but this foundation alone is insufficient. The best analogy for understanding the relationship between LLMs and ADP is the relationship between operating systems and applications: LLMs are the underlying "understanding engine," while ADP is the "application" built on top of it.

Three Fundamental Limitations of LLMs

First, LLMs are essentially stateless - they don't remember what documents they've processed before, what information they've extracted, or what exceptions they've encountered. The reality of enterprise document processing is: today processing 100 invoices from Supplier A, tomorrow processing 100 invoices from Supplier B, and the day after that comes a new format from Supplier C. If using LLMs directly, different instructions are needed each time. This "amnesia" makes LLMs unsuitable for enterprise application scenarios that require contextual continuity.

Second, LLMs cannot continuously learn and optimize. The capabilities of LLMs are fixed after training is complete, and they cannot learn and improve from actual business operations. When enterprises use LLMs to process documents, even if certain fields are consistently extracted incorrectly or certain formats are consistently misidentified, the model cannot automatically optimize. Human feedback and correction data cannot be absorbed by the system; each processing instance starts "from scratch." This means that using LLMs alone for document processing cannot create the "better with use" data flywheel effect.

Third, LLMs cannot handle multi-step complex tasks. Real-world document processing often requires multiple steps: first classify the document, then read its content, extract business-relevant information, perform calculations and verification, and finally execute business actions. LLMs excel at single-instance reasoning but cannot reliably execute such multi-step workflows. This "solo operation" characteristic makes LLMs unsuitable for complex end-to-end business processes.

These three fundamental limitations determine that enterprises cannot directly use LLMs to process documents, but need to build an application layer on top of LLMs. ADP is precisely such an application system. It not only leverages the understanding capabilities of LLMs, but also supplements memory, learning, and orchestration capabilities on this foundation, transforming the "understanding engine" into a "usable product."

[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

So how exactly does ADP address these three limitations of LLMs? What are its core characteristics? We elaborate in detail in the next section.

2.3 Three Core Characteristics of ADP

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

Traditional IDP requires preparing large volumes of annotated samples for each document type, whereas ADP achieves zero-shot learning: simply describe requirements in natural language, such as "extract supplier name, invoice number, line items, and total amount from this batch of invoices," and it takes effect immediately

without waiting for training. Even when suppliers change templates, no reconfiguration is needed.

ADP supports 100+ languages, including English, Simplified and Traditional Chinese, Japanese, multiple European languages, multiple Southeast Asian languages, etc. Multilingual mixed documents can also be processed correctly. This multilingual support capability is crucial for multinational enterprises and cross-border businesses.

ADP's generalization capability manifests across multiple dimensions: format generalization, supporting parsing of various layouts including paragraphs, multi-column formats, and tables; scenario generalization, where capabilities adapted on a few samples can transfer to previously unseen samples; and task generalization, not only document extraction but also performing tasks such as classification, comparison, review, and summarization.

Characteristic 2: Human-AI Collaboration Forms Data Flywheel

ADP's unique advantage lies in forming a data flywheel through human-AI collaboration, making the system smarter with use and continuously improving document processing effectiveness. ADP's human-AI collaboration can be based on both confidence levels and business rules to determine whether human intervention is needed, achieving intelligent routing. This specifically includes two aspects:

- Confidence rules: The system automatically scores the confidence of recognition results; high confidence (e.g., >95%) passes directly without review, medium confidence (e.g., 85-95%) marks questionable points for quick human verification, and low confidence (<85%) is routed to manual processing;
- Business rules: Users can customize business validation conditions, for example, "order number must consist of numbers and letters." Results that don't comply with rules will be automatically routed to manual review even if confidence is high.

Human corrections are automatically fed back to the system, which adjusts prompts or fine-tunes models based on this feedback. This represents ADP's fundamental advantage over traditional IDP: traditional IDP maintains fixed performance after deployment, with business changes requiring retraining; whereas ADP continuously evolves post-deployment, with business data automatically feeding back for optimization, creating a "data moat."

The table below provides an intuitive comparison of 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 manifest at three levels:

OCR achieved "seeing characters"

recognizing characters without understanding document structure and semantics

IDP achieved "understanding words"

extracting fields while relying on annotation and training, with limited transfer capability

ADP achieved "understanding intent"

delivering end-to-end understanding and execution with no training required and zero-shot transfer capability

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

In an era where models evolve rapidly, it's difficult for a single model to suit all types of document processing tasks. A simple invoice processing workflow might require: a specialized document classification model to identify invoice types, a vision-language model to understand complex layouts and table structures, and a large language model to extract fields and perform semantic reasoning. Different documents, different scenarios, and different tasks require different models to complete - this is the reality of document processing.

More importantly, real-world enterprise document processing workflows require not just models, but various tools: querying ERP systems to verify inventory, calling currency exchange APIs for conversions, using code for numerical calculations, sending emails to notify relevant personnel, etc. Traditionally, connecting these models and tools relies on manually writing rules or code, which not only has high development costs but also lacks flexibility - any business process change requires recoding.

ADP's core advantage lies in its agents' ability to automatically orchestrate multiple models and tools. In complex document processing scenarios, ADP agents can understand task objectives, autonomously determine which models and tools to invoke and in what sequence, handle intermediate results, monitor execution processes, and adaptively adjust when encountering exceptions, ultimately delivering business-required results. This agent orchestration delivers three fundamental advantages. First, automatic multi-model collaboration: different models perform their specialized functions - VLMs handle visual understanding, LLMs handle semantic understanding and reasoning, with the system automatically selecting the most appropriate model for each subtask. Second, automatic chaining of models and tools: no manual coding required - agents autonomously invoke APIs, query ERP systems, send emails, and manage data flow between components. Third, adaptive execution: the system can independently assess and adjust workflows when encountering exceptions, rather than simply reporting errors and halting.

Chapter Summary:

The introduction of ADP marks document processing's entry into the agent era. This is not merely an upgraded version of OCR or IDP, but an entirely new paradigm - leveraging the universal understanding capabilities of LLMs combined with the autonomous reasoning capabilities of agents to achieve complex workflow automation through orchestrating multiple models and tools, while forming a data flywheel through human-AI collaboration that enables continuous system evolution. ADP's three core characteristics (zero-annotation capabilities, multi-model collaborative agent orchestration, and a continuously evolving data flywheel) together constitute its competitive advantages, enabling it to genuinely solve enterprise document processing pain points and achieve the leap from "recognition" to "understanding," from "tools" to "agents." ADP brings not merely efficiency improvements, but 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 capable of understanding business, making autonomous decisions, and working collaboratively. However, moving from concept to implementation requires a pragmatic entry point. This chapter explains why ADP is the ideal first step for deploying enterprise AI agents.

3.1 Why Should Enterprise AI Agents Start with Document Processing?

Universality: Every Enterprise Has Documents

Regardless of enterprise size, industry, or geography, all organizations process substantial volumes of documents. From business documents such as contracts, orders, invoices, and reports, to internal documents like workflow approvals, meeting minutes, and project documentation, to external documents including email attachments, customer materials, and supplier files—documents are omnipresent. Research indicates that unstructured data accounts for over 80% of total enterprise data, with documents serving as the primary carrier of unstructured information. This universality provides a solid foundation for large-scale ADP adoption.

Clear Pain Points: Labor-Intensive and Error-Prone

Document processing has long been a critical pain point in enterprise operations. Finance teams process invoices, expense reports, and reconciliation statements daily; administrative staff handle contracts, approval workflows, and archive management; audit teams process quality inspection reports, compliance documents, and risk control materials. These tasks are highly repetitive, low in value density, result in poor employee satisfaction, and are prone to errors. This labor-intensive, inefficient, and error-prone reality is precisely what ADP addresses.

Clear ROI: Quantifiable Value

Unlike many AI applications, document processing delivers clearly quantifiable value. Labor cost savings, processing efficiency improvements, and error cost reductions can all be measured with concrete metrics (Chapter 5 details ADP's typical application scenarios and business value). Unlike many AI projects requiring long-term investment

before demonstrating results, ADP delivers immediate value, significantly reducing enterprise decision-making barriers and investment risk.

Controllable Risk: Progressive Deployment

ADP supports progressive deployment strategies that allow enterprises to adapt flexibly to their circumstances. Phase one positions ADP as an assistive tool, with the system extracting data while humans review outputs to verify accuracy and build trust. Phase two introduces intelligent collaboration, where high-confidence documents are automatically processed while medium and low-confidence documents route to manual review, progressively reducing manual intervention. Phase three achieves high automation, with over 90% of documents processed automatically, manual handling limited to complex exceptions, and a data flywheel enabling continuous optimization. This progressive approach makes ADP implementation risk substantially lower than AI projects requiring all-or-nothing deployment.

Technology Maturity: Production-Ready Today

Unlike many AI concepts, ADP technology is production-ready. Rapidly advancing Large Language Models (LLMs) and Vision Language Models (VLMs) demonstrate robust document understanding capabilities meeting production requirements. Agentic technologies have matured for enterprise deployment, supports on-premises installation for data security, and provides standard APIs for seamless integration. Gartner identifies intelligent document processing as a top 10 enterprise AI implementation scenario, with major enterprises prioritizing document automation in their digital transformation initiatives. Multiple enterprise success stories prove this is not proof-of-concept territory but scalable production deployment. Technology maturity, market validation, and ecosystem development make ADP an AI application enterprises can implement today with immediate results.

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

Phase 1: Validation (1-6 months)

The first phase objective is validating ADP capabilities in a high-value scenario and establishing a foundation of trust.

01

Phase 2: Expansion (6-12 months)

Following successful first-phase validation, phase two expands ADP capabilities across multiple scenarios and departments.

02

Phase 3: Deep Integration (12-24 months)

Phase three embeds ADP deeply into business processes, making it a core IT architecture component.

03

Phase 1: Validation (1-6 months)

The first phase objective is validating ADP capabilities in a high-value scenario and establishing a foundation of trust. Enterprises typically select pilot scenarios such as invoice processing in finance, order or contract processing in procurement, or resume and onboarding material processing in HR. These scenarios offer clear business value, relatively standardized processes, and sufficient data volume.

Success criteria include accuracy exceeding 90%, manual review ratio below 20%, and return on investment period under 12 months. Through this phase, enterprises validate ADP capabilities, build confidence in AI technology, and establish the foundation for subsequent expansion.

Phase 2: Expansion (6-12 months)

Following successful first-phase validation, phase two expands ADP capabilities across multiple scenarios and departments. Expansion takes various forms: processing different document types within the same department, such as extending from invoices to orders and contracts; or processing similar scenarios across departments, such as extending from finance invoices to procurement and sales invoices.

Key initiatives in this phase include building an enterprise-level ADP platform, establishing document processing standards and specifications, and developing replicable implementation methodologies. Through horizontal expansion, ADP evolves from an isolated pilot into enterprise-level foundational capability, processing more document types, covering more departments, and generating greater value. Enterprises accumulate document processing best practices and develop their own AI application methodologies.

Phase 3: Deep Integration (12-24 months)

Phase three embeds ADP deeply into business processes, making it a core IT architecture component. At this stage, ADP transitions from standalone tool to deeply integrated system, connecting with core platforms like ERP and CRM, embedding into end-to-end processes such as P2P (Procure-to-Pay) and O2C (Order-to-Cash), and becoming an indispensable intelligent node in business workflows.

During deep integration, enterprises build document knowledge bases and accumulate enterprise-specific document models. The data flywheel established through human-AI collaboration activates, continuously optimizing system performance and strengthening adaptability to enterprise-specific scenarios. This phase's hallmark is ADP's transformation from "optional tool" to "essential infrastructure," with its processing capabilities becoming a core 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	Score
Document Processing Volume (Monthly)	<500 documents: 5 points 500-2,000 documents: 10 points 2,000-5,000 documents: 15 points 5,000-10,000 documents: 18 points >10,000 documents: 20 points	____/20 points
Business Value of Document Processing	- Non-core business, minimal error impact: 3 points - Affects business efficiency, moderate error impact: 6 points - Linked to key business processes, significant error impact: 9 points - Directly affects customer experience or compliance risk: 12 points - Business bottleneck, errors may cause major losses: 15 points	____/15 points
Document Processing as % of Employee Work Time	<10%: 2 points 10%-30%: 4 points 30%-50%: 6 points 50%-70%: 8 points >70%: 10 points	____/10 points
Document Diversity	- Single format, highly structured: 1 point 2-5 formats, semi-structured: 2 points 6-10 formats, significant format variation: 3 points 10-20 formats, including unstructured documents: 4 points >20 formats, including handwritten, images, multilingual: 5 points	____/5 points

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

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

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

Technology implementation success depends on organizational readiness. ADP is not merely a technology project but an organizational transformation. This section provides concise change guidance for successful human-AI collaboration.

Job Transformation: From Operation to Judgment

Post-ADP implementation, employee work focus shifts fundamentally. For example, a finance department's 15-person invoice processing team might reorganize into 2 AI trainers (managing annotation feedback and system optimization), 3 data reviewers (handling low-confidence documents and exceptional cases), and 2 business analysts (conducting data analysis and process optimization with liberated capacity). The key is transitioning employees from "repetitive labor" to "high-value work," from "operational" to "judgment-based" roles.

Addressing Change Resistance

Employees may fear "AI will eliminate my job." Enterprises must clearly communicate: ADP is an augmentation tool liberating employees from repetitive labor to perform more valuable work. Case studies and data should demonstrate how ADP makes work more meaningful and skills more valuable.

Employees may worry "I can't use AI." Enterprises should provide comprehensive training, design intuitive user interfaces accessible to non-technical staff, and establish "AI Ambassador" roles where technology-savvy employees assist colleagues with one-on-one guidance.

Employees may fear "My professional skills will depreciate." Enterprises should redefine professional competency, elevating from "can operate" to "can judge." For example, advancing from "can enter invoice data" to "can detect invoice anomalies and trace root causes" represents higher-level capability.

CEO-Driven Change: Critical Actions

CEO or business leader support is paramount. Critical actions include: articulating AI strategy and vision at company-wide meetings, explaining ADP implementation rationale and implications for enterprise and employees; regularly reviewing project progress, engaging with frontline employees, understanding practical challenges and resolving them promptly; establishing "AI Innovation Awards" recognizing outstanding

employees and teams in AI adoption; incorporating AI application effectiveness into departmental KPIs, ensuring genuine organizational commitment.

Chapter Summary:

ADP is not merely a document processing tool but **the ideal first step for enterprise agentic AI**. Its entry point is precisely targeted: document processing offers universality, clear pain points, quantifiable ROI, controllable risk, scalability, and technology maturity—making it the optimal starting point for enterprise AI deployment. Its evolution path is well-defined: from single-point breakthrough to horizontal expansion, from deep integration to agent networks, with clear objectives and success criteria at each phase.

Through suitability self-assessment, enterprises objectively evaluate ADP implementation readiness. Through organizational change guidance, enterprises systematically address personnel and cultural challenges during implementation. For enterprises exploring AI deployment, **ADP provides a pragmatic, feasible, high-return starting point**. It enables enterprises to validate AI value with minimal risk, cultivate AI capabilities through real-world scenarios, and progress toward an intelligent future incrementally. Starting with ADP, enterprises not only solve document processing challenges but, more importantly, embark on the evolutionary path toward enterprise AI agents. From ADP, the journey to enterprise intelligence begins.

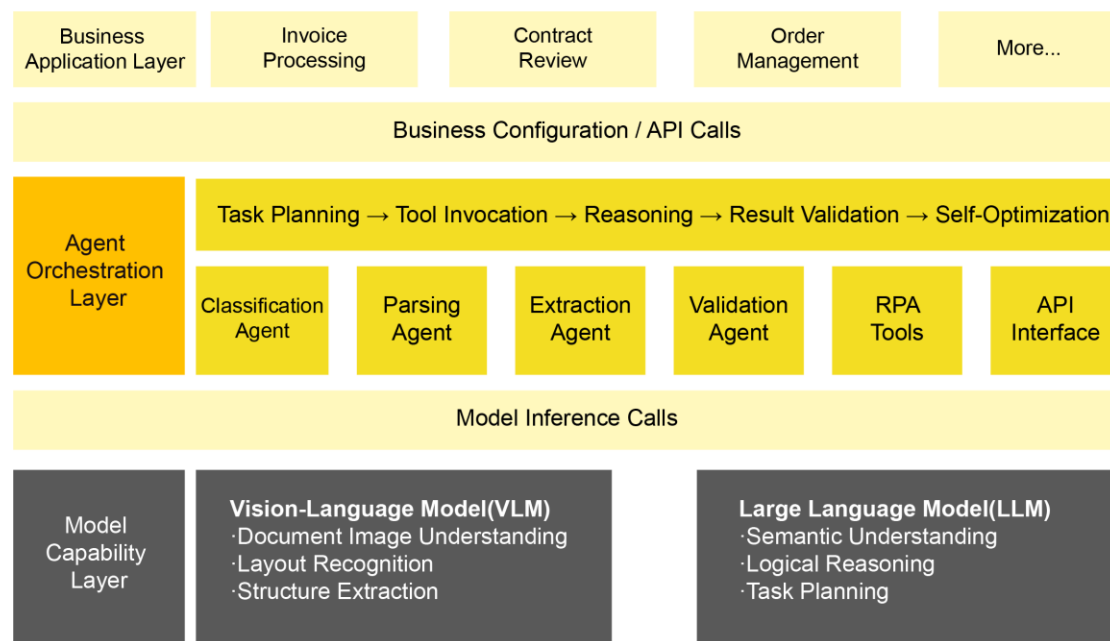
Chapter 4: Architecture and Enterprise-Grade Capabilities of ADP

In previous chapters, we discussed "why ADP is needed" and introduced "what ADP is". This chapter delves into the technical aspects, demonstrating that ADP is not only a product innovation but also an enterprise-grade platform that has been engineering-validated, deployable, and trustworthy.

4.1 Three-Tier Architecture Design

ADP adopts a three-tier architecture design that achieves decoupling of models, agents, and business scenarios, ensuring system scalability and maintainability.

Architecture Overview



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

The model layer is the core capability source of ADP. Its design follows two key principles: first, supporting the integration of multiple VLMs and LLMs with different characteristics, allowing upper-layer agents to flexibly select the most suitable model based on task requirements; second, fine-tuning for specific document processing tasks to achieve better performance and effectiveness.

Multi-Model Integration and Flexible Orchestration

In an era of rapidly evolving models, no single model performs optimally across all document processing tasks: some models excel at handling complex tables, some perform better in multilingual scenarios, and some deliver similar results at lower costs. ADP's model layer is designed as an extensible model pool architecture that supports simultaneous integration of multiple VLMs and LLMs, including general-purpose large models, enterprise proprietary models, and custom models fine-tuned for specific scenarios.

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

Although pre-trained VLMs and LLMs possess powerful general capabilities, fine-tuning for enterprise-specific document processing tasks can significantly improve accuracy and processing efficiency. ADP supports optimization at two levels.

The first level is prompt optimization. By analyzing actual enterprise document samples and business requirements, the system automatically optimizes prompt strategies by clarifying instruction formats, adding domain knowledge hints, and refining output structure definitions. This optimization approach is low-cost and quick to implement, typically improving accuracy by 5-10 percent.

The second level is model parameter fine-tuning. ADP supports fine-tuning of both VLMs and LLMs based on actual enterprise data. Fine-tuned models achieve significantly improved accuracy and performance compared to general-purpose large models when processing specific documents, substantially reducing operational costs.

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 user's document processing request, the agent orchestration system first understands the task objectives and constraints, identifying the task type (classification, parsing, extraction, comparison, etc.). It then decomposes

complex tasks into subtask sequences, selects appropriate models and tools based on each subtask's characteristics, and determines execution order and parallelization strategies. For example, agents can invoke VLMs to complete document parsing tasks, call LLMs for document extraction, use external APIs to retrieve information (such as exchange rate queries, address verification), and employ code execution and other tools for calculations. After completing an end-to-end document processing workflow, agents can also autonomously optimize based on feedback from human-machine collaboration or business rules, implementing a data flywheel.

For complex scenarios, ADP supports collaborative work among multiple agents. The document classification agent is responsible for identifying document types, the document parsing agent extracts document structure and converts it to Markdown, the information extraction agent extracts key fields, the data validation agent verifies data validity, the business review agent executes business logic and provides decision recommendations, and the orchestration coordination agent coordinates other agents to formulate complete execution plans. These agents perform their respective duties while collaborating with each other to complete complex document processing tasks.

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

The business layer provides out-of-the-box scenario-based applications as well as flexible configuration and extension capabilities. ADP pre-configures multiple business scenarios, including financial scenarios (invoice processing, expense reimbursement review, bill reconciliation), procurement scenarios (order processing, receipt verification, supplier management), contract scenarios (contract review, clause extraction, risk identification), logistics scenarios (waybill processing, warehouse documents, shipping documents), financial services scenarios (loan applications, credit reports, bank statements), and healthcare scenarios (medical prescriptions, lab reports, expense documents). Additionally, the business layer provides RESTful API interfaces, webhook callback mechanisms, supports both batch processing and real-time processing modes, and can seamlessly integrate with ERP, RPA, BPM, and other systems.

4.2 Performance Metrics

In real business scenarios, ADP's performance on a single node (on-prem deployment) is as follows:

[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

Performance can be optimized based on scenarios: batch processing scenarios prioritize throughput and can increase concurrency; real-time processing scenarios prioritize latency and reduce concurrency; multi-GPU deployments scale throughput linearly. Based on standard GPUs such as NVIDIA L20 or RTX 4090, a single server can meet the needs of 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%

In terms of field extraction accuracy, ADP achieves improvements of 10-20 percentage points compared to traditional IDP solutions, primarily due to the following capabilities: cross-format generalization enabling zero-shot transfer without retraining for new formats; multilingual support covering 100+ languages; and comprehensive handling of complex layouts including borderless tables, multi-column formats, and handwriting recognition.

4.3 Security and Compliance

Enterprise-grade applications must meet strict security and compliance requirements. ADP has incorporated security as a core consideration from the design phase, especially through on-prem deployment to thoroughly ensure data security.

Support for On-Prem Deployment

ADP supports fully localized deployment, with both VLMs and LLMs deployable in customer data centers. All document data is processed locally, and processing results are stored in the enterprise's own databases, completely eliminating enterprise data 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

On the hardware front, ADP supports standard GPUs such as NVIDIA L20, RTX 4090, and RTX 3090, requires CPUs with 48+ cores (clock speed >2GHz), 128GB+ memory, and 500GB+ hard disk storage. Deployment can be accomplished on a single server with support for horizontal scaling. The platform is compatible with domestic operating systems and databases, meeting domestic substitution requirements.

Data Security Assurance

Regarding data isolation, each enterprise customer's data is completely isolated with independent storage space and access permissions, preventing cross-tenant access. Development, testing, and production environments are fully isolated with no data flow between environments. For data encryption, all network transmissions use TLS encryption, and all API calls enforce HTTPS.

Access Control

For access control, ADP supports multiple authentication methods (username / password, SSO, LDAP, OAuth2.0) with optional multi-factor authentication (MFA). API access uses token authentication with regular rotation. Role-based access control (RBAC) is implemented with pre-configured roles including administrator, configurator, operator, and auditor, while also supporting custom roles and permission combinations following the principle of least privilege. All operations are logged, sensitive operations require secondary confirmation, and audit logs are immutable with regular archiving.

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.

Support for Bring-Your-Own-Models

ADP adopts an open architecture that allows customers to integrate their own LLMs and VLMs. Many large enterprises have already invested in building their own model capabilities or prefer to use specific commercial models. ADP's model layer supports standardized model integration protocols, enabling enterprises to integrate their proprietary models or selected third-party models into the ADP platform. These models work alongside the platform's pre-configured models and can be flexibly invoked by the agent orchestration layer based on task characteristics. This openness allows enterprises to fully leverage existing investments while selecting the most appropriate model combinations for specific tasks.

API-First: All Capabilities Programmable

ADP follows an API-first design philosophy, with all core capabilities exposed through RESTful APIs. The document parsing API converts images to structured text, the document extraction API extracts business data from documents, the document classification API identifies document types, and the workflow execution API supports complex batch processing scenarios. All APIs comply with REST standards, use JSON format for requests and responses, and provide complete API documentation with multi-language code examples. Webhook callback mechanisms are supported, with asynchronous processing completion actively notifying business systems. SDKs for mainstream languages such as Python, Java, Node.js, and C# are provided, automatically handling authentication and retries, supporting both synchronous and asynchronous calls, reducing integration development costs.

Native Integration with Laiye RPA

ADP achieves native integration with Laiye RPA, providing out-of-the-box functionality without additional development. ADP components can be directly dragged and dropped into RPA workflows to configure document processing tasks; ADP processing results are automatically passed to subsequent RPA processes, supporting conditional branching and exception handling, with a unified task monitoring and log viewing interface.

Easy Integration with Existing Enterprise Systems

ADP provides standardized integration methods for seamless incorporation into existing enterprise systems. Integration with ERP systems enables automatic supplier invoice posting, automatic purchase order verification, and automatic accounts payable/receivable matching. Integration with CRM systems enables automatic customer contract archiving, automatic sales order entry, and automatic customer information updates. Integration with OA systems enables automatic expense claim review, automatic approval workflow routing, and automatic document archiving. Furthermore, ADP can serve as an MCP (Model Context Protocol) Server that can be invoked by other agent systems, enabling inter-agent collaboration and building an enterprise-grade agent network.

Chapter Summary:

ADP's technical architecture embodies a careful balance between "advanced capabilities" and "engineering reliability". The three-tier architecture achieves decoupling of model capabilities, agentic orchestration, and business applications, ensuring system scalability. The model layer supports multi-model integration and task-specific fine-tuning; the agent layer implements autonomous orchestration of complex workflows; and the business layer provides rapid scenario-based deployment. Performance and accuracy metrics demonstrate ADP's effectiveness in real-world scenarios, while multi-layered reliability assurance mechanisms ensure enterprise-grade stability. On-prem deployment thoroughly protects data security, meeting strict security and compliance requirements. The open architecture and rich integration capabilities enable ADP to seamlessly integrate into existing enterprise IT ecosystems, serving as an intelligent enhancement layer rather than an isolated system.

Chapter 5: Industry Applications and Value Validation

The value of technology is ultimately manifested through business outcomes. This chapter presents four real-world industry scenarios, demonstrating how ADP helps enterprises achieve process automation, improve efficiency, reduce costs, and mitigate risks, truly realizing the closed loop from "AI capabilities" to "business value."

5.1 Manufacturing: Purchase Order Processing

Automation

Customer Pain Points

A manufacturing enterprise serving global customers processes approximately 5,000 purchase orders monthly. Orders come from diverse channels, including email, chat and customer systems, in various formats such as Excel, PDF, and images, with over 4,000 layout variations, covering more than 8 languages including Chinese, English, Japanese, German, Thai, and Spanish. This fragmented order landscape creates severe operational challenges for the enterprise.

The core pain points manifest in four areas. First, there are severe efficiency bottlenecks: manual entry requires 5-10 minutes per order, plus additional time for verifying material codes and customer information, requiring a dedicated team of 5 people for order entry. Second, quality risks are prominent: frequent human errors lead to data entry mistakes and omissions, inadequate order verification causes inventory anomalies, untimely customer information maintenance affects delivery, and erroneous orders result in customer complaints and economic losses. Third, resource misallocation is problematic: 80% of the team's energy is consumed by basic data entry work, unable to free up personnel for high-value activities such as order analysis and customer development. Fourth, expansion obstacles are evident: processing multilingual orders is difficult, new customers and new formats require retraining, and staffing cannot be scaled up promptly during peak business periods.

Solution

The ADP platform addresses these pain points through end-to-end process automation. The system supports full-channel order file access, enabling unified management of orders from multiple sources including customer systems, email, and IM software. The platform automatically extracts key fields through semantic

understanding, including order number, customer name, order date, product code, name, quantity, and amount.

The system implements an intelligent confidence assessment mechanism: for fields with confidence below a certain threshold, human verification is triggered for secondary calibration, ensuring accuracy in the last mile. This human-machine collaboration model ensures both efficiency and quality. Extracted order information is automatically reviewed, including comparison with historical orders, code conversion, and rule verification. For orders with review anomalies, the system automatically notifies personnel for manual processing; orders that pass review are automatically entered into the order management system by RPA robots and can be pushed down to the CRM system for delivery. Throughout the process, feedback from human calibration feeds back optimization suggestions, continuously improving effectiveness and forming a data flywheel. This end-to-end automation solution not only addresses order processing efficiency but, more importantly, breaks down data silos.

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 factories and offices in over 20 countries processes 2,000+ overseas supplier invoices monthly, involving more than 15 languages, requiring three-way matching with purchase orders and receiving documents, and involving multi-currency settlement and exchange rate conversion. This complex global business scenario creates enormous challenges for the finance team.

The core pain points are concentrated in four areas. First, multilingual processing is difficult: finance personnel cannot read invoices in all languages, requiring translation assistance leads to lengthy processes, and misunderstandings easily occur. Second, formats vary tremendously: invoices from over 200 suppliers have different formats, traditional IDP requires separate configuration for each supplier, and the system fails when suppliers change invoice templates. Third, three-way matching is complex: information inconsistencies between invoices, orders, and receiving documents mean manual matching is extremely time-consuming, with frequent omissions and errors. Fourth, approval cycles are long: the average approval cycle reaches 15 days, affecting supplier relationships and missing early payment discount opportunities.

These issues not only affect financial operational efficiency but, more importantly, impact the enterprise's cash flow management and supplier relationships. Traditional IDP solutions, requiring extensive sample training, prove inadequate when facing diverse layouts and limited volumes of overseas invoices.

Solution

The ADP platform achieves end-to-end automation through an intelligent invoice processing workflow. After receiving scanned copies, emails, or electronic invoices, the system can automatically identify language and format, extracting key information such as supplier, invoice number, item details, amount, tax, and currency. The system supports unified conversion of multiple currencies to the base currency and

automatically matches purchase orders and receiving documents through intelligent matching algorithms.

The system's core advantage lies in its powerful intelligent matching capability. Even when product descriptions on invoices and orders are not entirely consistent (e.g., "Laptop Computer" vs. "笔记本电脑" [laptop in Chinese]), the system can correctly match through semantic understanding. The system can automatically handle partial delivery situations, i.e., scenarios where one order corresponds to multiple invoices, and can identify and process special items such as price adjustments, discounts, and shipping fees.

After three-way consistency verification, the system automatically categorizes processing based on matching results: perfect matches are automatically approved and enter the payment process; small discrepancies are flagged with explanations and transferred to expedited review; major discrepancies generate detailed reports and are transferred to manual review. This tiered processing mechanism maximizes automation percentage while ensuring 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. 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 faces significant pressure processing large volumes of paper documents (such as identity certificates, insurance proposal forms, medical documents, etc.) in three core business processes: new policies, policy services, and claims. With the continued development of insurance business, especially the growth of third-party channels, the volume of paper documents has increased year over year, exceeding 300,000 annually. The core pain points of claims document processing manifest in three areas.

First, manual workload is enormous and repetitive. The existing document volume is already substantial, with the entire process relying on manual operations. Sales and related roles need to scan various documents offline into the imaging system and manually complete document classification; operations and related roles need to view documents one by one in the imaging system, manually enter key information into business systems, while manually verifying information consistency and document completeness. Both types of roles bear substantial mechanical repetitive labor, with high labor costs and susceptibility to errors from fatigue.

Second, time-consuming steps are prominent, with low processing efficiency. Manual processing time is mainly concentrated in three key nodes: information confirmation of identity documents one by one; field-by-field review and manual entry of standardized documents such as insurance proposal forms and application forms; identity matching of policyholders/insured persons for non-standardized documents and document completeness checks.

Third, handwriting recognition difficulty exacerbates manual burden. Documents to be processed contain both printed and handwritten content, with handwritten content covering three languages: simplified Chinese, traditional Chinese, and English. Some

handwriting has poor standardization, significantly increasing the time required for manual recognition and verification.

Solution

The project adopts a phased implementation strategy, first focusing on over 50 types of high-frequency documents, covering 80% of high-frequency scenarios, ensuring quick project results while laying the foundation for subsequent expansion.

The project builds an end-to-end integrated document processing platform through ADP Agentic Document Processing platform + RPA. ADP serves as the platform's "brain" and "eyes," responsible for understanding and extracting information; RPA serves as the system's "hands" and "feet," responsible for executing repetitive entry operations. The two are seamlessly integrated through an "intelligent scheduling center," forming a complete automation closed loop. ADP first automatically classifies document types, then extracts key information (such as customer information and insurance fields), and automatically interfaces with the business system's new policy and policy service modules. The system also has automatic review capabilities, automatically verifying identity information consistency and key document completeness, replacing manual confirmation steps.

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 a large volume of SME loan business, handling 500+ loan applications monthly. Each application contains 10-20 documents, requiring review of identity documents, business licenses, financial reports, bank statements, asset certificates, and other document types. This massive volume of manual document review creates serious operational pressure and risk hazards for the bank.

The core pain points manifest in four areas. First, review workload is enormous: manual review requires 3-5 minutes per document, 1-2 hours per loan application, and the 15-person review team still frequently works overtime. Second, review quality is unstable: depending on reviewer experience, quality varies, and fatigued work easily misses risk points. Third, customer experience is poor: approval cycles last 5-10 days, customers need to supplement materials multiple times, affecting loan acquisition rates and customer satisfaction. Fourth, compliance risks are prominent: the review process is difficult to trace, post-event verification of review basis is difficult, and regulatory inspection pressure is high. These issues not only affect the bank's operational efficiency and customer satisfaction but, more importantly, increase credit risk and compliance risk.

Solution

The ADP intelligent document review system addresses these pain points through end-to-end automation workflow. After customers submit loan application documents, the system automatically classifies document types, extracts key information, and structures it. For identity cards, the system extracts name, ID number, address, and validity period; for business licenses, it extracts company name, unified social credit code, business scope, and registered capital; for financial reports, it extracts revenue, profit, assets, and liabilities; for bank statements, it extracts deposits, withdrawals, balance, and counterparties.

The system then performs cross-validation, checking whether names, ID numbers, and company names are consistent. It then conducts business rule checks, including

whether identity cards are expired, whether business licenses are valid, whether financial indicators meet standards (debt ratio, liquidity ratio, etc.), whether bank statements match declared income, and whether blacklisted customers exist.

The system's core advantage lies in its powerful intelligent risk control capability. OCR fraud detection can identify anomalies such as photocopies, re-photography, and Photoshop modifications; data consistency checks can verify logical consistency between financial report numbers and bank statements; abnormal pattern recognition can identify risk patterns such as batch company registration and related-party transactions; the system can also automatically query external data such as enterprise credit reports, court enforcement records, and public opinion information.

Finally, the system generates risk scores and recommendations. For low-risk applications, it recommends rapid approval; for medium-risk applications, it flags suspicious points and transfers to manual review; for high-risk applications, it details reasons and rejects or requests supplementary materials.

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

While seeing the significant value ADP brings, enterprises may encounter some challenges during implementation. Understanding these potential risks and preparing response strategies in advance is key to ensuring project success.

Risk One: Document types are overly diverse, accuracy does not meet expectations

This is the most common concern for enterprises. Some enterprises may have hundreds of different document types with extremely inconsistent formats, worrying that ADP cannot accurately process them all. In reality, this problem can be solved through pragmatic implementation strategies.

The response strategy is to adopt the "Pareto principle," first focusing on processing the 20% of document types with the largest volume, which typically covers 80% of processing volume. In the first phase, set the target at 80-85% accuracy, with manual review mechanisms to ensure quality. As the system operates, through the data flywheel effect over 3-6 months, accuracy can gradually improve to over 90%.

Risk Two: Data security concerns, especially when processing sensitive documents

Industries such as finance, healthcare, and government have extremely high data security requirements, worrying that uploading documents to the cloud or external processing brings compliance risks. This is an entirely reasonable concern that needs to be thoroughly addressed through technical solutions.

The response strategy is to adopt a private deployment solution. ADP supports fully localized deployment, with all models and data running in the enterprise intranet, with data not leaving the domain. It also provides data desensitization functionality, automatically encrypting sensitive fields such as ID numbers and bank account numbers.

Risk Three: Integration with existing systems is complex, implementation cycle is long

Enterprises may worry that integrating ADP with existing ERP, CRM, and other systems will be complex, affecting business continuity. In fact, ADP's design has considered integration friendliness from the beginning.

The response strategy is to integrate with systems using APIs, RPA, and other methods, and implement in phases. ADP provides RESTful APIs, supporting integration with mainstream ERP and CRM systems, with development and joint testing typically completed in 1-2 weeks. Adopt "shadow mode" implementation, where ADP runs in parallel with existing systems, ensuring business continuity.

Chapter Summary

From these case studies, we can summarize three key factors for successful ADP application:

First, select scenarios with clear pain points. Scenarios with large document volumes, high labor costs, high timeliness requirements, and high quality risks are most suitable for ADP. All four cases possess these characteristics, enabling quick results.

Second, setting end-to-end process automation as the objective. ADP is not merely a document processing tool but must be deeply integrated with business processes to form a complete automation closed loop. In the manufacturing case, the system not only recognizes orders but also automatically reviews, enters, and pushes for delivery. In the multinational enterprise case, the system not only recognizes invoices but also automatically performs three-way matching and approval decisions. This kind of end-to-end process automation maximizes business value generation.

Third, continuous optimization and iteration. Continuously optimize models and rules based on real business data, forming a data flywheel. In the manufacturing case, feedback from manual calibration feeds back to optimize confidence algorithms. In the insurance case, phase one focuses on 50 types of high-frequency documents, with subsequent gradual expansion. This continuous optimization mechanism ensures the system's 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 month 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**:

Foucs 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 a global leader in intelligent automation and agentic AI, dedicated to helping enterprises achieve intelligent transformation and close the work execution gap through digital worker solutions. Founded in 2015 and headquartered in Beijing, the company has branches in Shanghai, Shenzhen, Hong Kong, Singapore, and other locations, serving customers across finance, manufacturing, energy, retail, pharmaceutical, and numerous other industries.

As a pioneer in enterprise-grade AI applications, Laiye Technology has launched ADP (Agentic Document Processing), a next-generation document processing platform powered by large language models and agentic technology. ADP has successfully helped hundreds of enterprises achieve intelligent upgrades in document processing, significantly improving operational efficiency and reducing labor costs.

We look forward to partnering with you to usher in a new era of agentic document processing.

LAIYE



Website:<https://laiye.com>

Email: mkt@laiye.com

Contact us:

Scan the QR code to start the conversation now.