



Intelligent Document Processing

To build or to buy?

There are many solutions for intelligent document processing (IDP), including proprietary and vendor-led approaches. Learn how to identify the best IDP solution for your needs, and the considerations you should make whether you decide to buy or build your own.



Contents

03	An AI revolution in document and communications processing
04	To build or not to build
05	Build your own (BYO) IDP: Leave it to ChatGPT?
10	IDP as a service
14	An enterprise platform to scale AI and IDP

An AI revolution in document and communications processing

Almost every department, in nearly every function, requires some form of document or communication to complete its tasks. Processes such as recruitment and onboarding, accounts payable, sales, order management, and customer service are just a few examples. With larger-scale businesses, the requirement for document processing and information extraction increases exponentially. Yet this can be a significant cost center for most businesses.

Despite a wide array of solutions, a large part of document and communications processing is performed manually. This labor-intensive process is not only time-consuming and costly but also fraught with risks such as inaccurate data entry, overlooked inputs, and irregular working methods. These factors often lead to operational delays and damaged customer relationships across the entire business.

However, the fusion of automation and state-of-the-art AI in the latest intelligent document processing (IDP) experiences provides a solution. Automation's potential to perform repetitive, monotonous tasks, and the cognitive ability of AI enables IDP to accurately understand and process both documents and communications. It's no surprise, then, that [the market for intelligent document processing \(IDP\) is growing 28.9% each year and is expected to hit \\$17.8 million by 2032.](#)

Gartner defines IDP as:
“specialized data integration tools that enable automated extraction of data from multiple formats and various layouts of document content. IDP solutions ingest data for dependent applications and workflows and can be provided as a software product and/or as a service.”

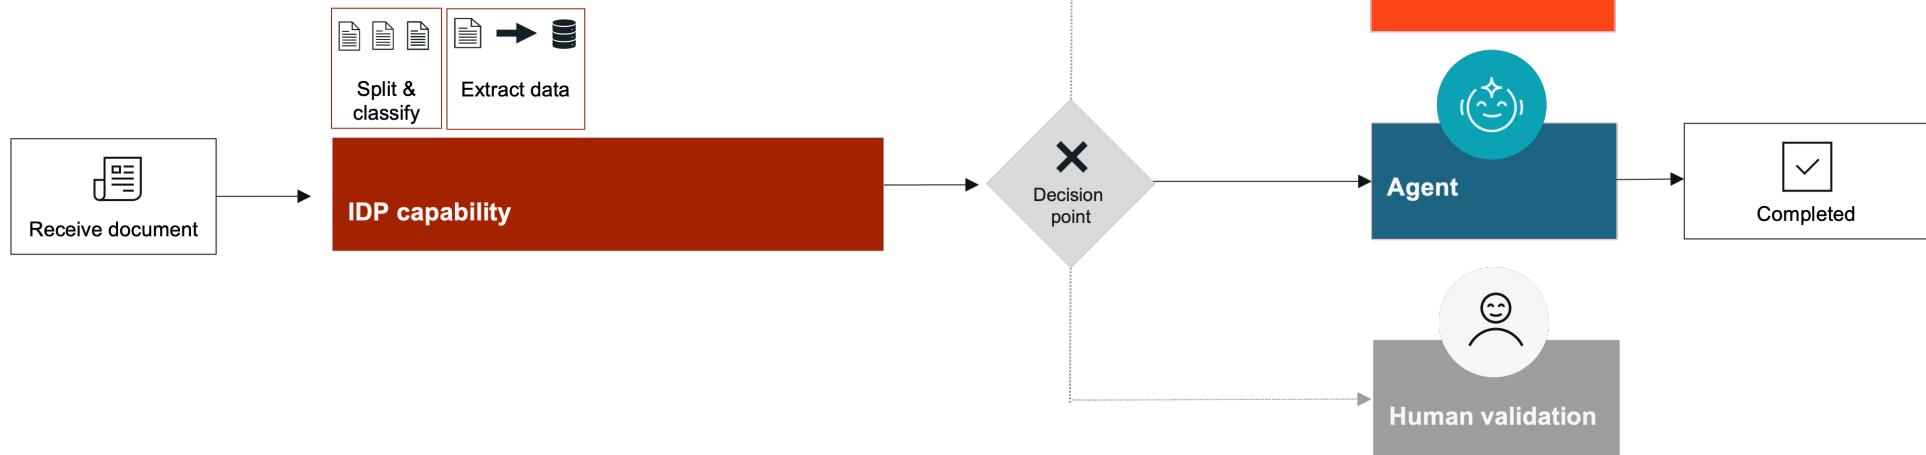
[Magic Quadrant for Intelligent Document Processing Solutions, 2025](#)

To build or not to build?

IDP typically blends numerous AI technologies, including large language models (LLMs) and image recognition, to help enterprises rapidly process documents and communications at scale.

As business leaders embrace AI agents—intelligent AI-based entities able to complete work for and on behalf of people—document-based processes are an ideal use case for agentic automation. Combining AI agents with robots and IDP capabilities, overseen by humans in the loop, makes for fast returns and has a major impact on efficiency.

However, this raises the question: to build or to buy? It's important to assess all options and consider which approach to IDP can deliver the highest return on investment, the best performance, and the fastest time to value.



Build your own (BYO) IDP: Leave it to ChatGPT?

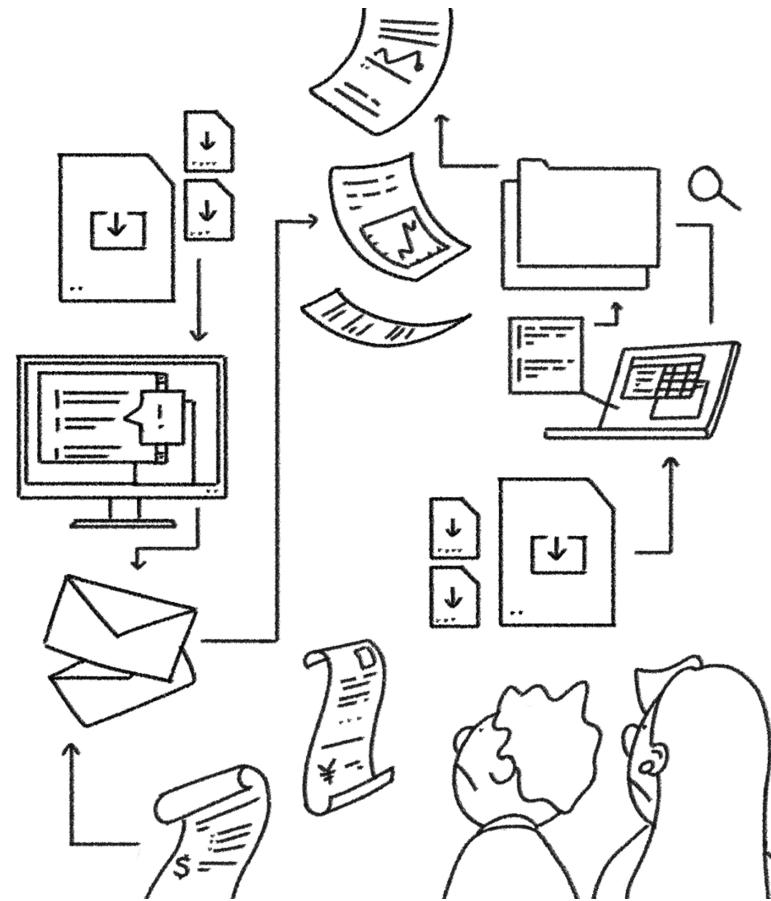
As LLM capabilities continue to expand—evolving from purely text-based to multimodal vision language models that can also understand images—business leaders are now asking: **‘Why can’t we just process all our documents using ChatGPT or a similar LLM?’**

Indeed, with the growing availability of LLMs and supporting APIs, enterprises have never had more tools to help them build custom IDP systems. However, lots of tools in your toolbox doesn’t make the actual task any easier.

In a BYO IDP system, each component (from language understanding to data extraction and automation) will need to be built from scratch or, more likely, sourced from multiple third-party providers. Using an external LLM for language and visual understanding capabilities is only one part of a much larger, more complex system.

A BYO IDP system provides a business with end-to-end ownership and greater customization potential compared to vendor solutions. They have the flexibility to adapt their system to changing business needs, without having to work with another organization.

However, in most cases, these benefits are outweighed by the key challenges of the BYO approach.

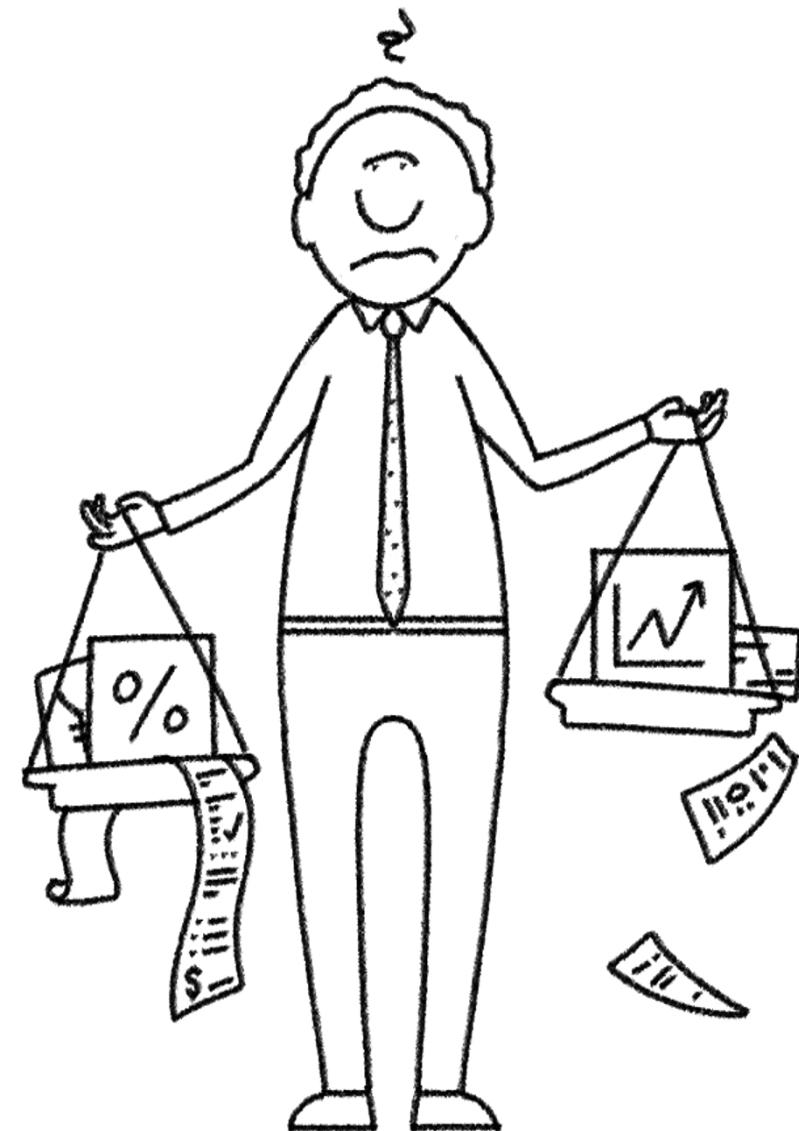


Cost

It's a common misconception that BYO is cheaper than paying for IDP as a service. In most cases, this isn't true in the short or long term.

Developing and then maintaining your own IDP system demands significant time and expensive specialist talent. You need software developers to create the platform and user interface (UI), data scientists for data preparation, pre- and post-processing, and many other experts for tasks like monitoring performance (which you'll need to create your own reporting dashboards for), as well as auditing and logging. Consider that even third-party LLMs require AI specialists and engineers to fine-tune the chosen model to exact business requirements.

Maintaining your own platform also necessitates continuous updates and resources. Any BYO system that requires data annotation will need training materials for users, and these will have to be updated to reflect your user interface.



Risk

Depending on rare and expensive technical talent to keep a system operational is risky. Due to cost and talent shortages, these teams tend to be small. They may face limitations in the number of use cases and business units they can realistically support.

Talent attrition can also make a system non-performant or non-viable in the long term. There's the ever-present risk that project funding could be pulled as well.

These challenges are compounded when AI skills and model fine-tuning are required. [According to Forrester research](#), 30% of AI decision makers see a lack of technical skills among the greatest barriers to adopting GenAI at their organization—the most frequent response.

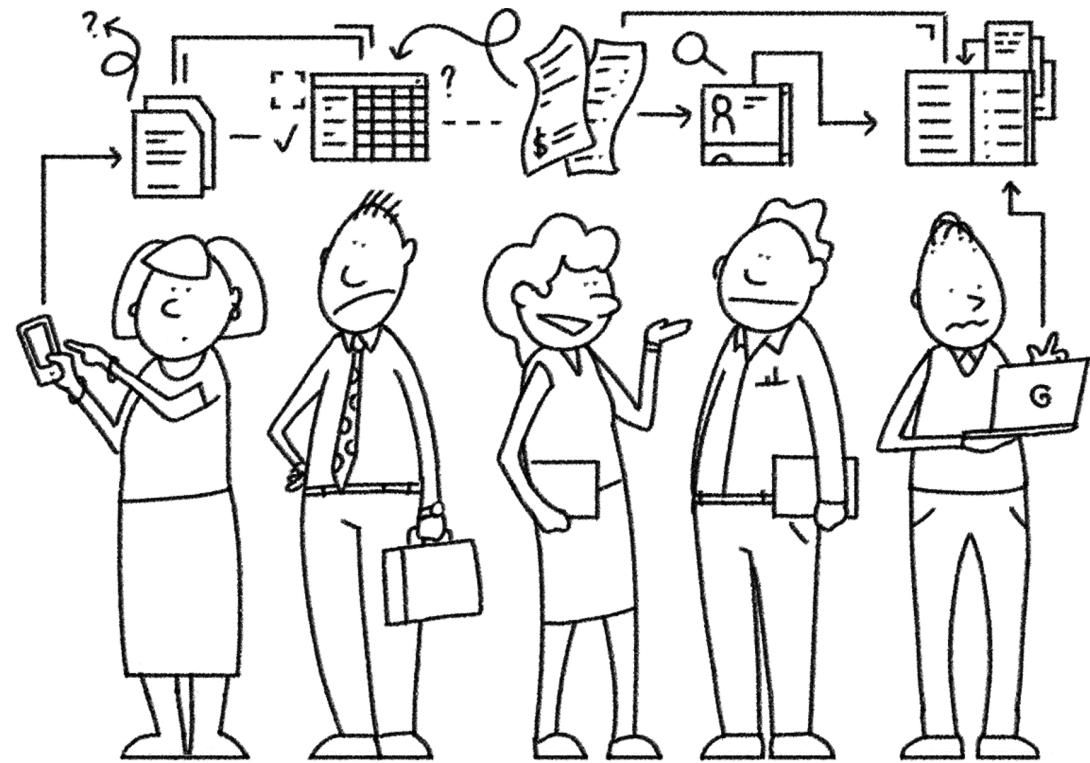


Complexity

When you build an IDP system, you are solely responsible for complex AI model and platform governance. Indeed, systems built for complex use cases may require hundreds of AI models that have to be managed.

For instance, a large bank might need several hundred models fine-tuned for various use cases to achieve the necessary level of accuracy. Even if an AI system can process documents and communications ‘out-of-the-box,’ a significant amount of prompt engineering or context gathering is still needed for acceptable performance. Scaling this to hundreds of use cases would be incredibly difficult, as you would have to annotate data, benchmark, deploy, and maintain hundreds of prompts.

There are many hidden costs involved in an IDP system you’ve built yourself. Every component of the system is an important decision, and each technology requires specialist skills and increases technical debt (along with increased risk). Inevitably, BYO is a heavy lift and means slower time to value. Higher lifetime cost is likely due to talent, governance, and maintenance requirements. It’s no surprise that 69% of technology decision makers consider document extraction and routing use cases highly difficult to deploy, according to Forrester analysis for UiPath.



Technical limitations

LLMs are a transformative technology with immense business value. However, they have specific limitations that curtail their utility in enterprise-level, intelligent document processing (IDP) and related automations.

Dead-ends

Some problems simply can't be solved through prompting alone. Whether this involves an LLM's core extraction capability or its performance on a specific field, table, or document, no amount of prompting can generally solve certain problems at scale. With an exclusively prompt-based solution, this can lead to frustrating experiences we describe as a 'dead-ends.'

For example, extracting data from tables in documents and communications can see significant problems. Third-party LLMs frequently make mistakes like skipping rows, mixing up columns and rows, or extracting data where it's missing.

Automation readiness

With simple documents and just one or a small number of samples, it's very easy for an LLM to get an accurate extraction—at first glance. However, seeing fields extracted in a chat experience is not sufficient for automation. You need to drive consistency in the structured schema that's extracted, and you need data typing to ensure the data is in the correct format for your automation.

Attribution

Attribution, or references/citations, is key when automating a process. Attribution is how a model shows a user where a given field has been extracted from a document. Because human review remains necessary to automate important documents at scale, it's critical to provide clear attribution for model outputs. Without it, a user would need to manually review the entire document, which defeats the purpose of automation.

Third-party LLMs do not provide reliable attribution and have been shown to erroneously invent citations in many cases. A lack of reliable attribution limits the utility of LLMs in document processing.

Confidences

Confidences—the degree to which a person can trust a specific output from a model—are critical for ensuring human-in-the-loop review is both effective and efficient in document processing. Users need a simple metric to help them decide whether a model output should be reviewed by a human. Third-party LLMs don't provide confidence metrics, leaving the only option to manually review every document, grinding the process to a halt and defeating the purpose of AI-enabled automation.

IDP as a service

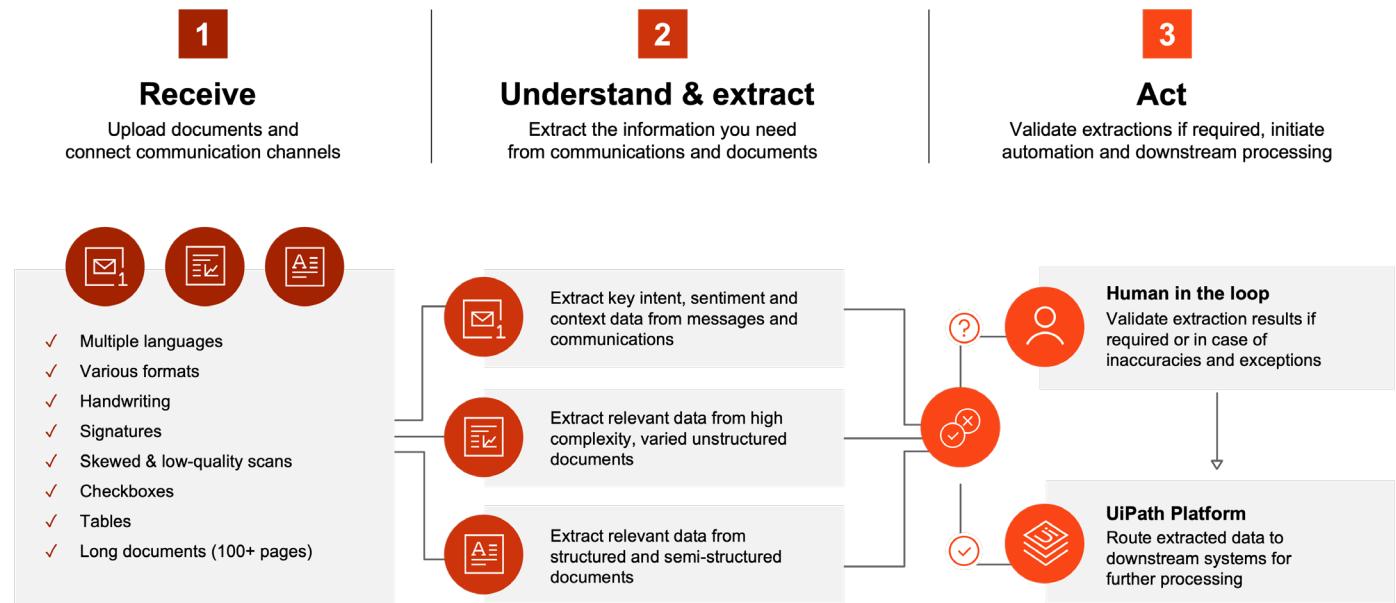
The alternative to building your own system is purchasing IDP as a service from a third-party vendor.

There are two main approaches to this:

1. Purchasing IDP as a **point solution** and integrating it with the rest of your enterprise tech stack.
2. Purchasing IDP as part of a larger **solution or platform**. This provides access to additional capabilities, such as automation, as needed.

Purchasing IDP as a service gives businesses less control over platform development compared to a custom-made system. However, platform providers will work with their customers to ensure the system evolves to meet their needs.

There are also many other advantages to consider.



Time to value

Implementing an existing IDP platform is usually faster than developing a new one. Established platforms are tried and tested over many years of use in large enterprises across diverse use cases. Training and enablement resources are already created and, in many cases, professional services support is provided to help users hit the ground running and start generating value from their implementations.

BYO systems built around third-party LLMs require costly and time-consuming finetuning and prompt engineering before they are ready for business use. By contrast, the AI underlying IDP as a service is typically designed around rapid, low-code training experiences.

UiPath Intelligent Xtraction and Processing (UiPath IXP) leverages the latest LLM innovations while giving users the tools to evaluate, fine-tune, and improve the accuracy of their outputs. This enables ordinary business users and AI models to actively collaborate and accelerate time-to-production and value.

70% reduction in time

Thermo Fisher Scientific cuts invoice processing time by 70% for 824,000 documents processed annually



The solution has delivered a 70% reduction in the time it takes to process invoices, with about 53% of all invoices being handled without any human involvement. This helped our P2P finance team dramatically reduce the workload of eight full-time employees who were managing about 824,000 invoices annually.

Luis Cajiao, Senior Manager, Smart Automation Global Business Services, Thermo Fisher Scientific

Reduced risk

Opting for IDP as a service substantially lowers various forms of risk. You are no longer reliant on expensive internal talent to keep your IDP system operational. The vendor owns system maintenance, platform and model governance. IDP solution providers are also expected to deliver the highest levels of enterprise-grade data security.

Cost of ownership should also be considered. Hasty development, imperfect code, or the use of soon-to-be-outdated technologies in a BYO IDP system all raise the risk of technical debt. Over time, this necessitates costly corrections and system upgrades.

Purchasing IDP as a service greatly reduces risk and technical debt for the buyer. Vendors prioritize technological upgrades to stay competitive, adopting the latest capabilities, continually iterating and improving their service. They deal with the necessary rework, testing, updating, and replacing obsolete components.

→ Simon Knowles, Head of Technology, vabble

“Don’t attempt to build the models yourself. Attention-grabbing headlines make it look simpler than it is. Find a vendor that was doing AI before the LLM hype. Competence in adjacent AI processes is needed such as unsupervised learning, data preparation, and basic ML calculations.”

LLMs Promise Document Automation Glory: But Should We Believe It?, Forrester, May 10, 2024

Scalability

A custom system requires custom-made connectors and APIs to integrate with the relevant enterprise systems. Depending on the size and complexity of the business, this may require hundreds or even thousands of combined development hours.

Established IDP platforms will have pre-made connectors for the most popular enterprise systems, enabling fast integrations and time to value. With a cloud-based IDP system, you can also expect access to the latest AI upgrades and capabilities without eating the cost to build or integrate those capabilities yourself.

For all these reasons, IDP as a service should be the preferred option for enterprises seeking fast time to value, accuracy, and reliability in their IDP system.

Custom-builds invite considerable technical debt, heightened risk and liability in the long term. IDP as a service places the burden of responsibility on a specialized platform honed by years of competition and iteration.

50% business growth

Dexcom automates the prescription intake process, reducing the AHT by 80% and saving 200,000 hours a year.



There was no additional headcount added. When we were growing at 50% per year, from 300 to 600 prescriptions a week. We didn't have to hire additional help because we had Document Understanding able to manage the intake.

Stephen Sikes • Director of Operational Excellence at Dexcom

An enterprise platform to scale AI and IDP

UiPath IXP is the ideal choice for enterprises looking to implement best-in-class AI and IDP, safely and successfully. IDP is embedded in our end-to-end automation framework, enabling you to automate document and communications processing. You can also leverage prebuilt capabilities, customize them, and bring in AI agents and third-party components to automate complex enterprise processes end to end.



Unlock enterprise data safely, at speed and scale



Unlock enterprise data

Identify and extract the information you need from documents and communications.

Automate more, and extend end-to-end automation into areas previously untouched.



The latest in GenAI

GenAI embedded and leveraged throughout IXP to provide faster time to value and enhancing user experience.



AI Guardrails

Production-grade model controls, built-in safeguards for data protection, human in the loop, compliance, and governance for responsible AI use.



UiPath Platform

IXP is part of the UiPath platform that connects hundreds of data sources and enables action at scale.

The UiPath Platform provides an industry-leading solutioned approach to IDP. We provide extensive platform capabilities and controls over LLMs, enabling them to be rapidly scaled across the enterprise:

Enterprise-ready

UiPath first-party models are governed under the strongest enterprise controls. The UiPath Platform delivers robust rules-based access control (RBAC), model versioning, and extensive performance guardrails for our first-party Document Understanding and Communications Mining models. Human in the loop is also built into our platform experience, ensuring generated outputs are properly validated.

Furthermore, UiPath manages all third-party models through the UiPath AI Trust Layer, delivering governance, trust, and security for GenAI. This means zero data retention and external model training with your business data.

A flexible, open model strategy

AI models are advancing quickly, and the best model now may not be the best model for your use case in the future. We are constantly investing to create the best specialized LLMs for core business processes and are constantly improving them. The UiPath model family, including our proprietary models for IDP are testament to that.

Yet our AI strategy remains open. We provide the tooling to integrate your favored third-party or proprietary LLMs, and manage them under the UiPath AI Trust Layer. You have the flexibility to use a combination of the best models in class for every task across your use cases.

Optimization for improved accuracy and reliability

UiPath IDP is optimized for advanced data extraction with techniques like RAG and curated system prompting. Context grounding means IDP models that are safer, more performant, and accurate. Our Validation Station interface shows evidence of where the extracted information was found in the document, meaning generated outputs can be easily verified with evidence.

Project-based experience

Our innovations are guided by real, on-the-ground IDP deployments in major enterprises. Our users simply define their document type and can leverage multiple models.

They can also evaluate performance, monitor, and manage model versions—all key capabilities for deploying AI, maintaining, and then scaling it across the business.

Cutting-edge AI research for IDP innovation

UiPath combines the latest AI models and techniques to provide the best performance for our customers and their unique use cases.

UiPath IXP will evolve to intelligently orchestrate and combine multiple models—selecting the most effective approach for each document processing task—while incorporating the latest advances in pre- and post-processing.

Proven success

UiPath has also recently been recognized as a Leader for Intelligent Document Processing (IDP) in the Gartner® Magic Quadrant™ for Intelligent Document Processing Solutions.

[Read the report.](#)

“UiPath is investing in agentic automation via its IXP offering, deeper AI/LLM integration including its native SLMs DocPath and CommPath, vertical use cases, and expanding Autopilot and agentic IDP capabilities for more autonomous and intelligent document processing.”

Top business benefits that most companies experience from adopting IDP are optimizing the supply chain, increasing customer satisfaction, increasing employee satisfaction, eliminating data silos for all stakeholders, and creating greater access to documents across the business.

This is borne out in the real world through the experience of UiPath IDP customers:

- [Green Point](#) saw a 98% reduction in document processing time, cutting 10,000 hours of manual work per year and achieving 100% compliance with regulations
- [Teads](#) achieved a 50% reduction in support and operation workloads, with millions of dollars gained annually through increased focus on strategic tasks
- [Banca Progetto](#) achieved 90% automated document verification and reduced handling time for automated tasks by 68-70%
- [Expion Health](#) increased the volume of insurance claims handled on a daily basis by 600%
- [Covestro](#) streamlined repetitive HR processes to save 85% of time spent on manual document processing
- [Omega Healthcare](#) increased productivity by 100% in accounts receivable and achieved 50% faster turnaround time
- [Hiscox](#) broker services team cut process lead time by 300% for requests to enter workflow

Ready to supercharge?

See how the IDP capabilities of the UiPath Platform deliver fast time to value and help scale state-of-the-art AI throughout the enterprise.

[Learn more about UiPath IXP](#)



“In less than nine months after deploying the UiPath solution, we processed about 40,000 invoices, or about 4,500 monthly. We initially had a goal of processing 75% without human intervention, but achieved about 90% straight-through processing during that time period.”

Thomas Earvolino
Director of Financial Systems, Canon USA

