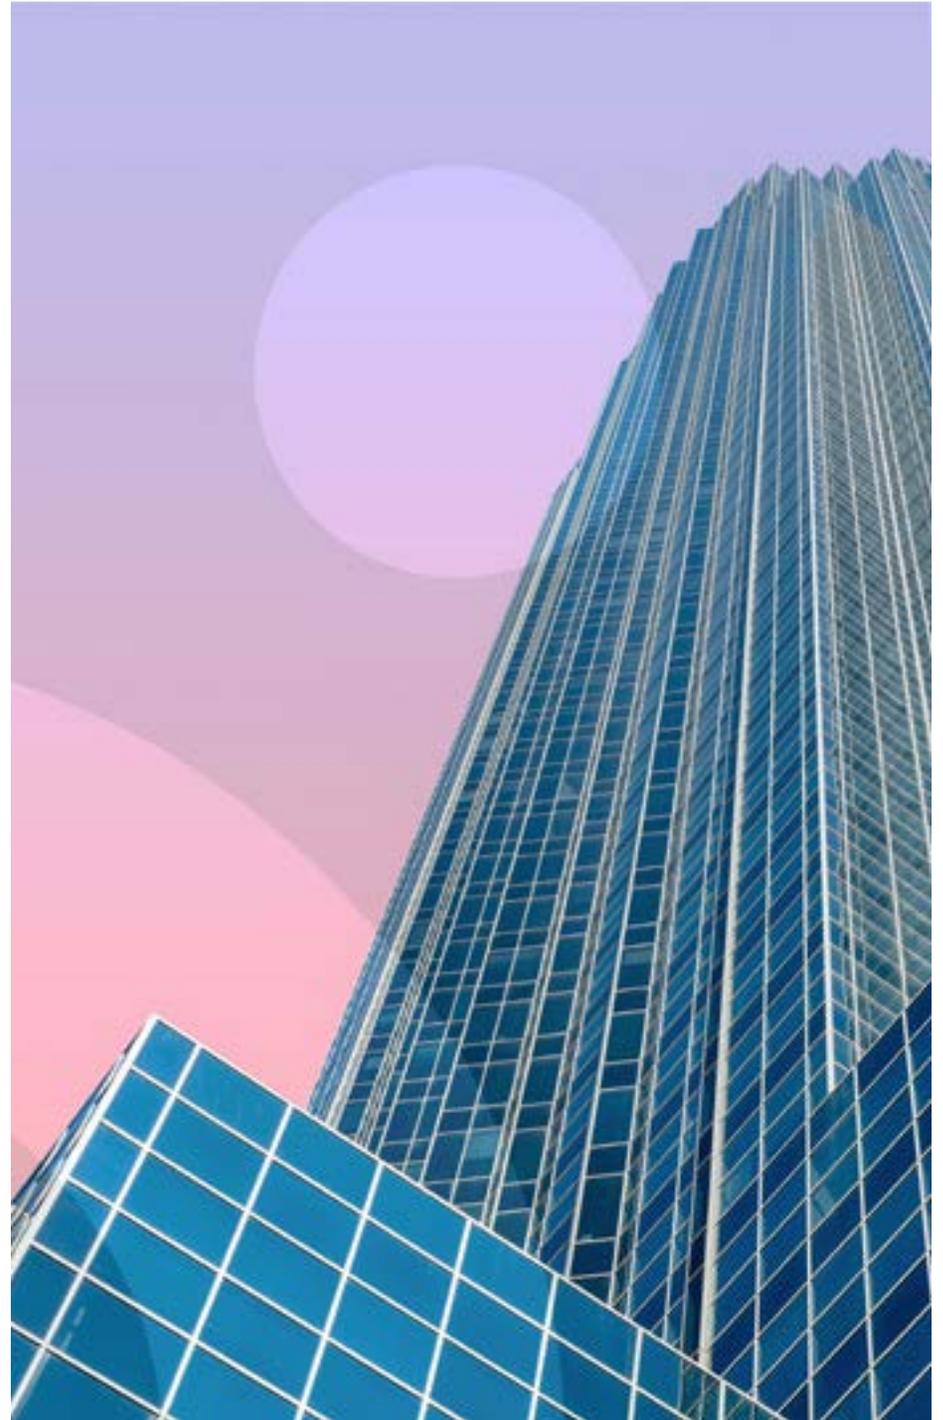




The Ultimate Guide to Intelligent Document Processing

Where to start and how to succeed

Document processing is changing fast, with agentic automation delivering value across industries and domains. Learn how to automate document processing end-to-end—with high accuracy, at speed and scale.



Contents

Section 1	03	Document processing has been around for decades. It's now changing fast
	04	An intelligent answer to a paperwork problem

Section 2	05	Get your documents processed intelligently
	06	Delivering value across industries
	07	Enterprise-wide improvement with intelligent document processing
	08	Complex content: the evolution of IDP
	09	Introducing UiPath IXP

Section 3	10	Documents & communications are key
	12	Unlocking success with end-to-end automation: Where to start
	14	Analyzing the business impact and outcomes Measuring success

A new era for document processing - driven by AI

Enterprises run on documents and communications. Almost every process, in almost every department, needs a document or message to get the job done. Hiring and onboarding, accounts payable, sales and order management, customer service—the list goes on. The larger the organization, the more documents it needs to process and the more information it needs to extract. And, as every enterprise knows, managing it all can be a major financial and resource burden.

Despite the plethora of solutions available, both document and communications processing remains, for many, an intensely manual task. It's time-consuming and full of risks—of inaccurate data, missed inputs, and inconsistent ways of working. All of which creates frustrations in teams and operational bottlenecks across the enterprise. But things are changing.

Agentic automation—combining powerful AI agents, automation robots, generative and specialized models—offers exciting new opportunities to redefine document and communications processing. Not only to remove manual headaches, but to radically streamline daily operations within finance, customer services, human resources, sales, and more—through intelligent document processing (IDP). And it's nothing short of transformational.

This paper provides a guide for enterprises looking to embrace the true potential of IDP and agentic automation—discussing where to start and how to succeed.

Everest Group describes IDP solutions being used in: “a wide variety of use cases across different business functions and verticals. IDP adoption not only helps enterprises achieve cost savings but also improves their workforce productivity and employee and customer experience”.

[Intelligent Document Processing \(IDP\) and Unstructured Document Processing Products PEAK Matrix® Assessment 2025](#)

An intelligent answer to a paperwork problem

IDP has come a long way in a short time—combining traditional document processing with OCR, automation, specialized and generative AI (GenAI) to read, interpret, and act on data automatically. Now, intelligent, autonomous AI agents are leveraging IDP to rapidly and accurately automate document-based processes.

Enterprises are embracing the latest, AI-powered IDP for good reason. Inefficient document processing may negatively affect businesses:

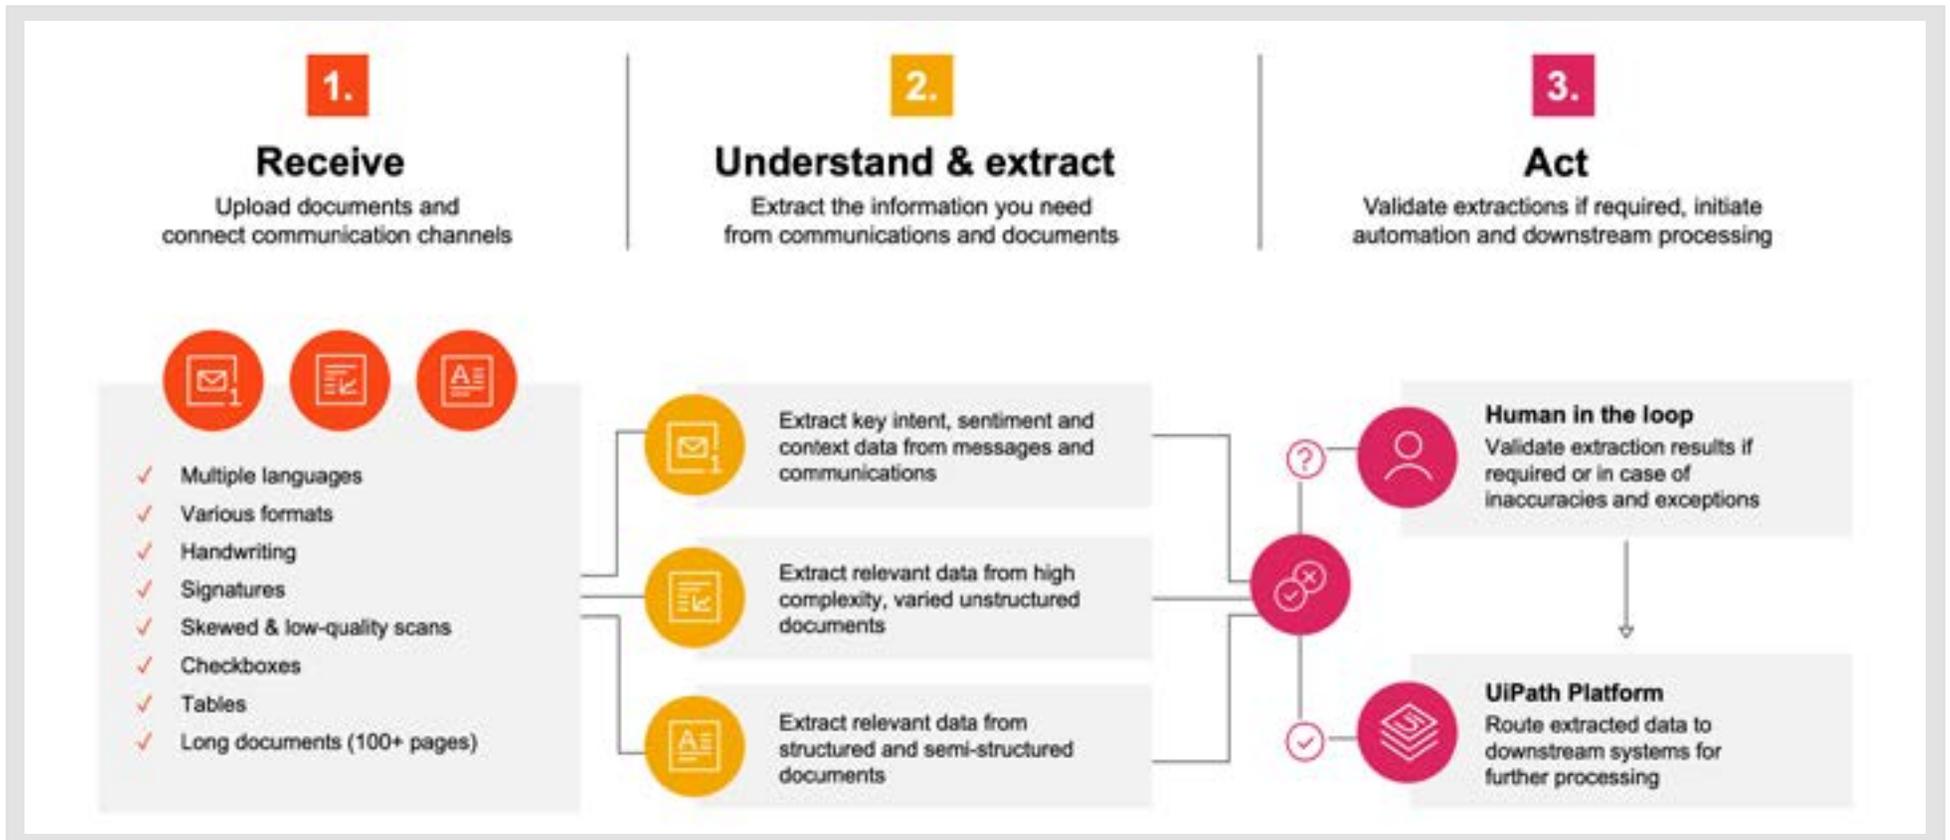
- **Limited business growth and scalability**
Labor-intensive document and message processing limits ability to scale efficiently and capture market opportunity.
- **Poor customer experience and third-party vendor management**
Complex, unstructured data mandates human decisioning, slow onboarding, and servicing.
- **Increased financial and compliance risks**
Higher chance of data input errors, missed information, and incorrect procedures.

According to a McKinsey and Company report, the combination of AI and other technologies can automate work that currently absorbs 60-70% of employees' time.

Process your documents intelligently

IDP solutions are capable of automatically processing documents in a multitude of structured, semi-structured, and unstructured formats (from forms to emails). Using a range of AI-powered technologies, documents are digitized and classified, and then data is extracted and interpreted with little or no manual intervention. Afterwards, that data may be used by an AI agent to complete a task, or get sent to an employee for exception handling and correction.

Put it another way: AI agents can use IDP tools to read documents of any difficulty, pull desired information, and pass it to where it needs to go. All with a high level of accuracy and reliability. Average handling time (AHT) of the document is dramatically reduced. Errors are removed, and humans can instead focus on higher-value tasks.



Delivering value across industries

Using IDP, it's possible to massively reduce manual intervention and radically streamline a near-infinite range of document-intensive processes.

Extend this out across all lines of business, and the potential impacts are huge. Customer requests can be screened by AI agents to quickly match the request type and action items to enhance customer experience. Loan and mortgage applications can be scrutinized with the same results. Time to compliance falls in finance and healthcare—improving customer and employee experience. Pretty much anywhere there's a document and/or communications-based process, IDP can enhance and streamline for better results and higher business impact.

Reducing the overall cost of processing huge volumes of documents and messages is just one driver among many. 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 (according to IDC's State of Content Services Survey, June 2023).

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

- [Thermo Fisher Scientific](#) cut invoice processing time by 70% for 824,000 documents processed annually
- [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
- [Canon USA](#) finance department saved 6,000 hours on processing thousands of invoices annually
- [Hiscox](#) broker services team cut process lead time by 300% for requests to enter workflow

Enterprise-wide improvement with intelligent document processing

Enterprise

- > Hiring & onboarding
- > Finance processes (e.g. accounts payable & receivable)
- > Sales & order management
- > Customer service & support

Banking and Financial Services

- > KYC – Know Your Customer
- > Mortgage application
- > Client & customer onboarding
- > Trade settlement & amendment

Insurance

- > KYC – Know Your Customer
- > Underwriting processes
- > Claim processing
- > Customer queries and communications

Healthcare

- > Patient registration
- > Consent and assent processes
- > Health insurance claims
- > Health history

Public Sector

- > Immigration applications
- > Benefit applications
- > Unemployment verification
- > Citizen questions & feedback

Manufacturing

- > Sales & order management
- > Accounts payable & accounts receivable
- > Purchase order processing
- > Customer and vendor communications

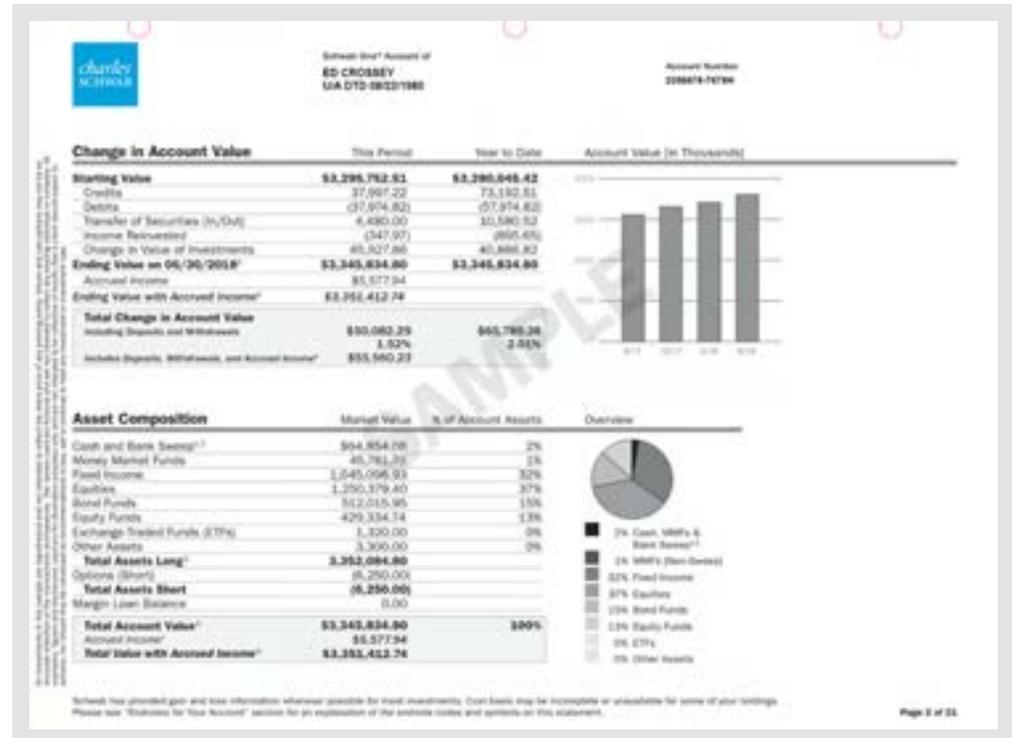
Complex content: the evolution of IDP

Traditionally, IDP solutions have struggled to understand and process complex, unstructured documents. This is because IDP has long depended on rules-based methods or specialized AI models pre-trained to understand and process documents in a consistent, structured schema.

Think of an invoice—many clearly structured ‘fields’ where important information like customer names, dates, and invoice numbers can be found.

Compare this to a complex contract, or a brokerage statement where key information is buried within masses of text, often without a standardized format. These complex documents are typically very long and their structures varied. They can have complex elements like nested tables, charts, images, and very long page counts, often in the hundreds.

However, in the last couple of years, an answer has been found. IDP solutions have begun to combine the use of specialized AI with state-of-the-art GenAI models. These are usually large language models (LLMs) similar to those powering UiPath Autopilot and Anthropic’s Claude. These LLMs are capable, not just of absorbing large amounts of multi-model data, but also of reasoning and interpreting it to understand its meaning.



By combining structured, specialized AI with flexible GenAI, the latest IDP solutions enable businesses to process the bulk of their enterprise data, whether it's structured, semi-structured, or unstructured.

When combined with AI agents, robots, and people, entire processes involving complex, unstructured documents and communications are automated end-to-end.

Introducing UiPath IXP

UiPath has provided leading IDP for years. But our approach has developed to now incorporate advanced GenAI.

We call our new IDP capability **UiPath IXP (Intelligent Xtraction and Processing)**, where the 'X' not only represents the idea of 'extraction' but also a growing number of diverse content types, including complex and unstructured documents. UiPath IXP is built for fast time to value and a seamless user experience to take the pain out of complex document processing of all types.

UiPath combines the use of specialized AI models for structured documents with state-of-the-art GenAI for complex unstructured content. Our inference-first experience means you can provide the model instructions (just like a prompt) and it will extract the information you need and put it into the format you specify.

UiPath IXP also offers precise controls to ensure the accuracy of our IDP outputs. Users have the ability to write prompt instructions at the level of individual fields and are given tools to evaluate model accuracy. They can perform closed-loop learning to improve performance. Our models also give confidence scores for every prediction so you can manually review when the model isn't entirely sure. Finally, our models also provide attribution, giving sources and references for all their predictions.



Multi-modal data classification and extraction



New prompt-driven capability for unstructured, and high-complexity documents



Configurable and flexible user experience that leverages GenAI and specialized models to reduce time to production

The screenshot displays the UiPath IXP interface. At the top left is the UiPath logo. The main heading reads "Choose the right tool to process your data". Below this, there are three tool cards:

- Communications data**: Use Communications Mining for short-form communications data, like emails and workflow tickets.
- Unstructured and complex documents**: Use Generative Extraction for documents with free form text, multiple tables, graphics and inferred values, like reports and contracts. A "Preview" link is visible next to this card.
- Structured and semi-structured documents**: Use Document Understanding for documents like forms, invoices, and purchase orders.

At the bottom of the interface, there is a search bar labeled "Search for dataset", a "Columns (4)" dropdown menu, and two buttons: "Administration" and "Create dataset".

Combined in a single AI-powered IDP solution, UiPath IXP makes it simple to automate document-intensive processes, reduce manual paperwork, and increase operational efficiency right across the organization.

The Document Understanding capability addresses complex document processing and removes manual data extraction headaches by empowering agents or robots to read, understand, and act on documents using AI understanding. Automations can now work at speed and with a high accuracy level that cuts the risk of errors—automatically determining the location of document data, even when the type or format of the document changes.

Depending on the document type and business process requirements, Document Understanding offers different document processing methods – from rules-based to AI models. There are out-of-the-box and custom specialized AI models tailored for specific business use cases, including insurance claims reconciliation, know your customer (KYC), accounts payable and accounts receivable automation, and more. GenAI strengthens the solution as it can be leveraged in a number of ways, in particular, to classify documents, extract data from unstructured documents, and validate the output. Additionally, GenAI and LLMs within Communications Mining support data extraction from unstructured communications like emails.

Communications Mining extends automation into entirely new business areas. In most organizations today, communications—email, chat, calls, tickets, and more—are routed to the correct department. Employees read it, copy the details, and manually key them into relevant systems. Not any more. By connecting UiPath to all these messages, inefficient manual processes are eliminated and customer requests automated to cut costs and scale faster.

Both Document Understanding and Communications Mining leverage GenAI to accelerate time to value in IDP implementations as well as to address more automation challenges—like finding answers in complex unstructured documents or generating an email response to a customer request.



What Everest Group says about UiPath:

“UiPath maintains its position as a Leader in the Intelligent Document Processing (IDP) Products PEAK Matrix® Assessment 2025 by achieving strong growth driven by innovation. Some of UiPath’s latest advancements include an expanding library of pre-built solutions for use cases across industries and process areas, Autopilot for UiPath Studio, a new prompt-driven capability for unstructured, and high-complexity documents, and ongoing investments in agentic automation.”

[Intelligent Document Processing \(IDP\) and Unstructured Document Processing Products PEAK Matrix® Assessment 2025](#)

Unlocking success with agentic automation: where to start

IDP gets even more superpowers when it's part of end-to-end agentic automation. The UiPath Platform is a foundation for **enterprise-scale IDP—from process discovery to analytics and optimization, offering multiple deployment options from cloud to self-hosted.**

Customers choose UiPath Automation Cloud for its time to value and the operational benefits of the platform as a service approach. You can also leverage the integrated platform experience with on-premise deployments, or on public or private clouds.

Whether SaaS or self-hosted, you get all the enterprise-grade compliance and security capabilities—and more, including complete audit trails, pluggable integration with Azure Key vault, and IT-specific governance through UiPath Orchestrator with Azure AD integration.



A good place to start an IDP journey is to look at what the business is trying to achieve—and the impact existing manual processes are having on its ability to achieve those outcomes. Are labor-intensive processes limiting the ability to scale, efficiently, or capture new market opportunities? Is the volume of documents and extended processing times affecting customers or suppliers? Are mundane, manual tasks having an impact on employee experience, leading to retention and recruitment issues? Are data input errors

and missed information creating high rework levels and process inefficiencies? With the business impact understood, it's important to consider the kinds of documents and communications causing the bottlenecks, and their processes, layouts, and formats. These will be different for different departments—emails and tickets in customer service, invoices and purchase orders in finance, contract agreements and order amendments in sales, etc. Communications are part of

many document-intensive processes, so it makes absolute sense to use IDP for both documents and messages. Think about any process that includes a message (email or support ticket) with a document attached to it. There's plenty of untapped automation opportunities across customer service, IT support, human resources, and finance, to name just a few.

It's also useful to identify the measurements used to quantify the performance of the current document

understanding process—to identify areas for improvement and help you design target metrics.



Here are some questions to help you get started:

- What document types are being processed?
- How are documents being processed today? What works well and what can be improved?
- What system(s) are the documents received in?
- What system(s) does the document need to be sent to?
- What system(s) is the extracted data used in?
- What does the end-to-end process look like? Is it documented?
- What is the current average handling time to process a document?

Analyzing the business impact and outcomes

IDP solutions reduce the processing time and improve data accuracy to avoid duplication and rework. IDP delivers on all counts, and the ultimate outcomes include money saved, money earned, and costs prevented.

These are the building blocks to more meaningful business outcomes:



Accelerated productivity
Automate highly-manual document processing tasks to save time and costs



Better customer experience
Decrease the response time to deliver better customer experience



Higher accuracy
Reduce human errors to mitigate financial and compliance risks



Happier employees
Help employees escape from the mundane chores and focus on higher-value tasks



Money saved

Time is money. If teams can reduce manual hours on document reviews, data entry, exception management, and decision making, the cost of the process will naturally fall.

IDP applications analyze and extract data from multiple documents simultaneously. And, by extracting and processing documents using predictable methods, error and rework rates plummet. Average handle time (AHT) improves and the door opens to significant savings.

With fewer employees required to handle document processing tasks, organizations can reduce labor costs associated with hiring, training, and managing additional staff. A human in the loop approach—bringing machine and human intelligence together—ensures efficient human-agent and robot collaboration. This delivers a higher level of confidence in the overall process, and its impact is especially tangible in more complex scenarios.

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

Money earned

Cost reduction may be king, but IDP has a significant impact on increasing revenue, too. Faster processing onboards new customers more efficiently, increases conversion opportunities in contact center environments, and enhances customer satisfaction and referral business. And with AI-powered automation increasing team capacity, enterprises can accommodate higher volumes of customer requests, sales orders, and transactions—serving a larger base and capturing more market share.

600% increase in volume

Expion Health accelerated the volume of insurance claims processed daily by 600% which resulted in improved client experience and employee engagement.



We have gone from processing about 75 claims a day in the past to now frequently processing as many as 500 claims or more in a day. That's almost a 600% increase in the volume of claims we can handle.

D.S. Suresh Kumar, Chief Transformation Officer, Expion Health

Costs avoided

Manual processes often lack consistency and standardization, especially when multiple individuals or teams are involved in handling documents and communications. Manual data entry is especially prone to human errors, such as typos, omissions, or misinterpretation of data. According to [Gartner](#), poor data quality costs organizations an average of \$12.9 million. IDP applications ensure consistent and standardized data extraction and processing across documents and formats. This not only reduces the inconsistencies, misinterpretations, and reworks that create customer, compliance, and cost issues, but also enable faster business growth at lower cost.

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

Measuring success

To properly assess the ROI of your IDP deployment as well as to evaluate the business outcomes mentioned above, it's important to use analytics tools like UiPath Insights but also to understand which data is critical in your case. While there are so many metrics to track, the most important measure of success in IDP is average handling time (AHT). It measures the time it takes to process a document from start to finish, including any time spent on handling exceptions and corrections by humans. It provides a comprehensive view of the real efficiency and performance of the document processing solution.

Another metric that is usually tracked in any automation processes is the straight-through processing (STP) rate, though the STP rate alone doesn't show the real ROI when it comes to IDP. STP is calculated by dividing the number of documents processed automatically by the total number of documents processed. A higher STP rate indicates that the system can process a higher percentage of documents automatically, with no human intervention. A lower STP rate only indicates that there are more instances when human validation is required, yet it doesn't really measure reduction in manual effort with automation compared to manual document processing.

Measuring the accuracy of your solution is a critical factor in any IDP operation. But it's easy to get bogged down in definitions. There's AI model accuracy, OCR accuracy, extraction and classification accuracy, field level accuracy. Depending on your document and your process, it's important to understand each of these. However, the key accuracy metric that matters more than all these is outcome accuracy as it directly impacts the AHT. Outcome accuracy is the sum of all these parts, and also includes the added impact of business rules and human validation.



Spending less time on a process is a good thing—and translates directly into cost savings. AHT looks at how many hours your people are spending on the manual process against the time it takes with IDP in operation. This allows you to calculate the cost savings as well—the math is pretty simple.

Ready to supercharge?

Find out how UiPath can intelligently automate document and communications processing for your business.

[Learn more about UiPath IXP](#)

[Contact us](#)

