



IDP

AI DEPLOYMENT

BFSI

# Documents Still Run Your Business.

## *They Shouldn't Slow It Down.*

How financial institutions are quietly unlocking 60–80% cost reductions in document-heavy workflows — and what our delivery teams have learned deploying IDP at scale.

Viracent Delivery Practice · 7 min read · Published 2025

## 78%

Average reduction in manual document processing time across BFSI deployments

## 6–9 wks

Typical time-to-value for a well-scoped IDP rollout on Microsoft Azure AI

## 3× ROI

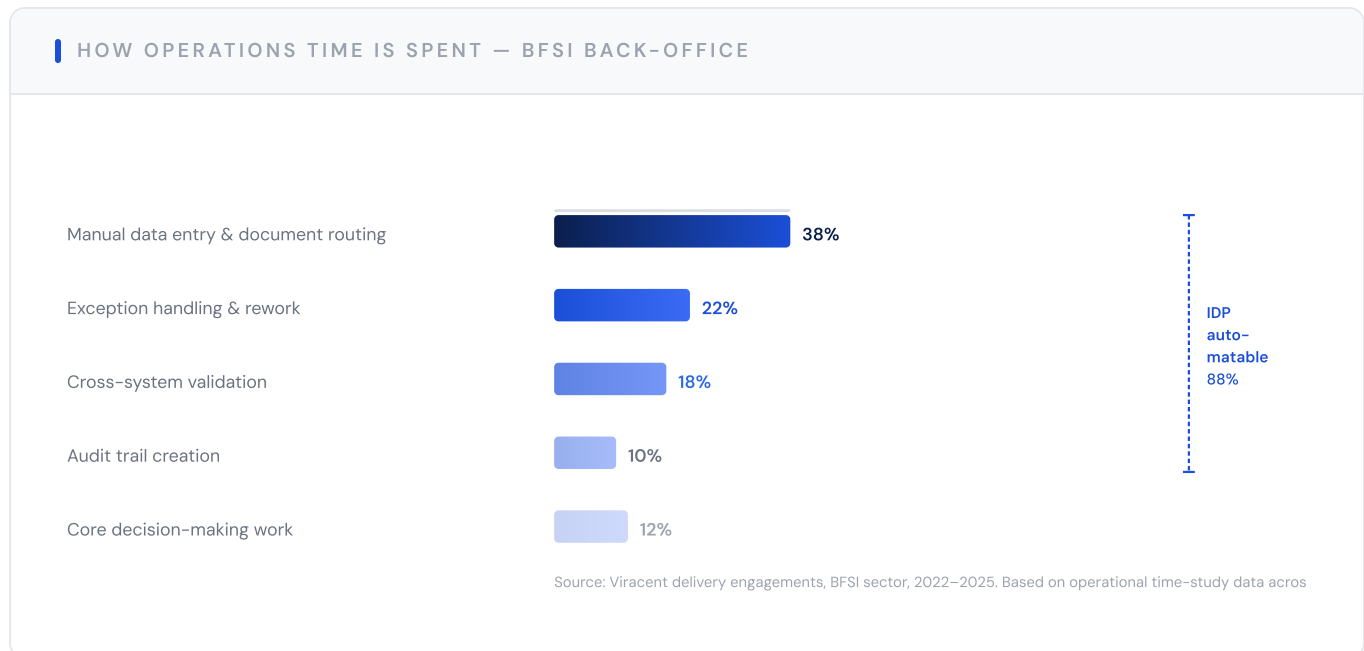
Median return seen within the first year when IDP replaces legacy OCR pipelines

# Where time actually disappears in financial operations

Walk the back-office of any mid-size bank, insurer, or NBFC today and you will find something incongruous: teams using state-of-the-art trading systems and real-time risk dashboards — yet routing loan applications, KYC packets, and claims bundles through manual data-entry queues that would look familiar to someone from 2005. Documents remain the single biggest friction point between a decision and its execution.

In our engagements with BFSI clients across India, the Middle East, and Southeast Asia, we have consistently found that between 30–40% of an operations team's productive hours are absorbed by activities a well-deployed IDP system can handle with greater speed and higher accuracy: extracting structured data from unstructured PDFs, cross-validating fields against system records, routing exception cases, and generating audit trails.

The problem is not a lack of awareness. It is a lack of clarity on where to start and a healthy scepticism about whether AI tools translate well from a demo environment to a live, messy document reality.



***"The first thing we do on any IDP engagement is audit the actual documents — not the clean samples the client prepared for us. That***

*audit almost always changes the scope in useful ways."*

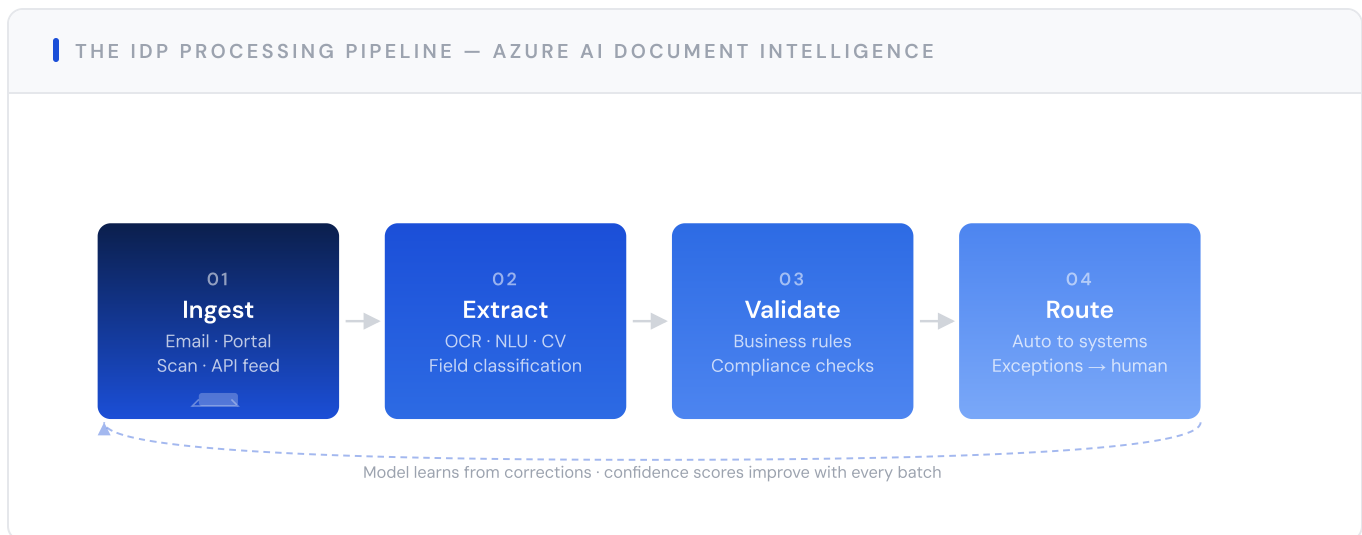
— MANISH, VIRACENT DELIVERY LEAD, DOCUMENT INTELLIGENCE PRACTICE

WHAT IDP ACTUALLY IS

## Beyond OCR — understanding the full capability stack

Intelligent Document Processing is not a single product. It is a layered capability that combines optical character recognition, computer vision, natural language understanding, and workflow automation into a system that can read, interpret, and act on documents in a way that approximates — and in high-volume scenarios, exceeds — what a trained human operator does.

In a mature IDP deployment on Microsoft Azure AI Document Intelligence (our preferred platform as a certified Microsoft Partner), the processing chain typically spans four stages:



The critical distinction from legacy OCR is the validate and route layer. Traditional OCR extracts text and hands the problem back to a human. IDP applies business logic to the output — and learns from corrections over time, improving confidence scores with every batch processed.



### **Start with one high-volume, well-defined document type**

Clients who scope their first IDP deployment around a single document class — say, loan application forms or insurance claim submissions — see measurable ROI within the first quarter. Trying to solve all document types simultaneously creates a training data problem and delays value realisation by 3–6 months.



### **Data quality is the deployment's real constraint — not the AI**

In over 70% of our BFSI engagements, the first two weeks are spent cleaning and labelling training data, not configuring the model. Clients often underestimate how variable their own documents are across branches, time periods, and intake channels. A document audit before kickoff is non-negotiable.



### **The human-in-the-loop design is not a compromise — it is a feature**

IDP does not eliminate human oversight; it focuses it. Our best-performing deployments maintain a human review queue for low-confidence extractions, typically representing 5–15% of total volume. This keeps accuracy high, satisfies regulatory audit requirements, and builds the team's trust in the system faster than a fully automated approach.




### **Integration is where projects quietly stall**

Extracted data has to land somewhere — a core banking system, a CRM, an ERP. In legacy BFSI environments, API availability is often limited or inconsistently documented. We now build an integration discovery sprint into every project plan, and where native APIs are unavailable, we deploy lightweight middleware adapters rather than waiting for IT to open access.

# How to frame IDP as a strategic investment, not an IT project

The most effective business cases we have seen for IDP in financial services are built around three numbers: the cost of current manual processing per document, the error rate and its downstream cost (rework, compliance penalties, customer churn), and the volume trajectory over the next 24 months.

For a mid-size insurer processing 15,000 claims documents per month with a 4% manual error rate, the math typically produces a compelling case for IDP investment within the first conversation. But the less-obvious value driver is speed to decision. When KYC verification drops from three days to four hours, or when a loan pre-assessment that required two analysts can be completed overnight, the competitive impact is direct and measurable.



## TYPICAL VALUE DRIVERS BY BFSI SEGMENT

- **Retail Banking:** Loan origination, KYC packet verification, account opening forms — volume is high, variability is moderate, ROI is fast
- **Insurance:** Claims intake, policy renewal documents, medical records processing — complexity is higher, but accuracy gains are substantial and audit-trail requirements make IDP a natural fit

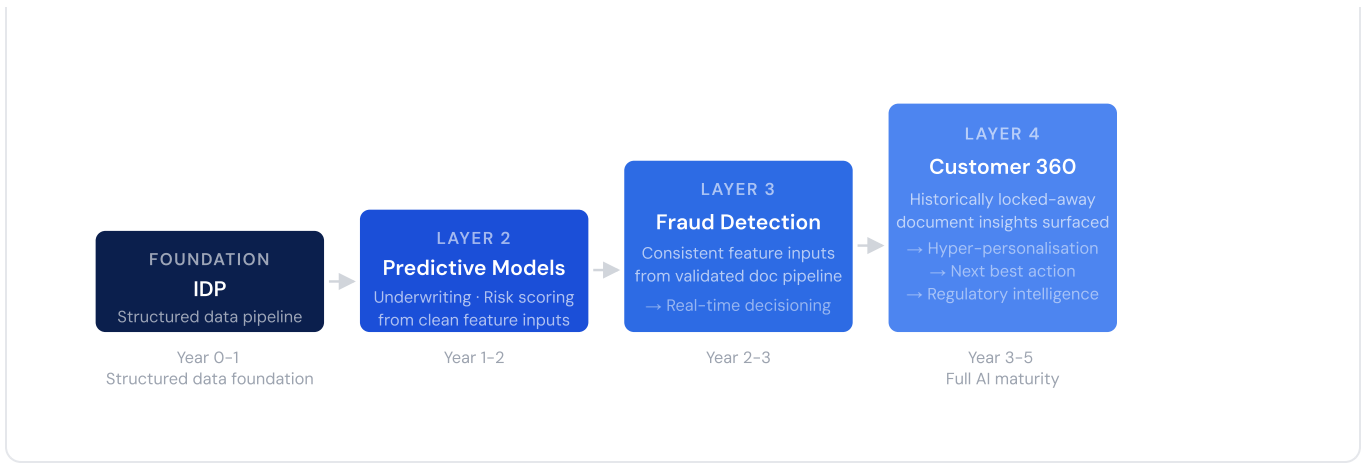
- **NBFCs and Fintechs:** Due diligence packs, invoice financing, trade finance documents — often under-resourced operationally; IDP provides enterprise-grade processing capacity without enterprise-grade headcount
- **Wealth and Asset Management:** Onboarding packs, mandate documents, regulatory filings — low volume but high compliance sensitivity; IDP's structured audit trail reduces regulatory risk materially

## IDP as the first step in a broader AI journey

We encourage clients to treat an IDP deployment not as an end-state but as a foundation. Once you have a reliable, structured data pipeline from your document estate, the next layer of intelligence becomes significantly easier to build: predictive underwriting models that feed on clean application data, fraud detection systems with consistent feature inputs, customer 360 profiles that include historically locked-away document insights.

The organisations that gain the most from AI over a five-year horizon are those that start with the unglamorous work — getting their unstructured data into a usable state. IDP is the most practical, highest-ROI entry point to that journey available to financial institutions today.

What separates a successful IDP deployment from a failed pilot is almost never the technology. It is governance: clear ownership, a realistic phased plan, an honest document audit, and a delivery partner willing to say no to scope creep in the first sprint.



## Thinking about an IDP deployment?

[Speak to our team →](#)

Our delivery team offers a complimentary document-readiness audit for qualified BFSI organisations — typically a two-hour working session that produces a clear picture of deployment scope, effort, and expected returns.

© 2025 Viracent Private Limited

Field Insights are produced by the Viracent delivery practice. Views represent practitioner perspectives, not investment or legal advice.

[viracent.com](https://viracent.com)  
[ContactUs@Viracent.com](mailto:ContactUs@Viracent.com)