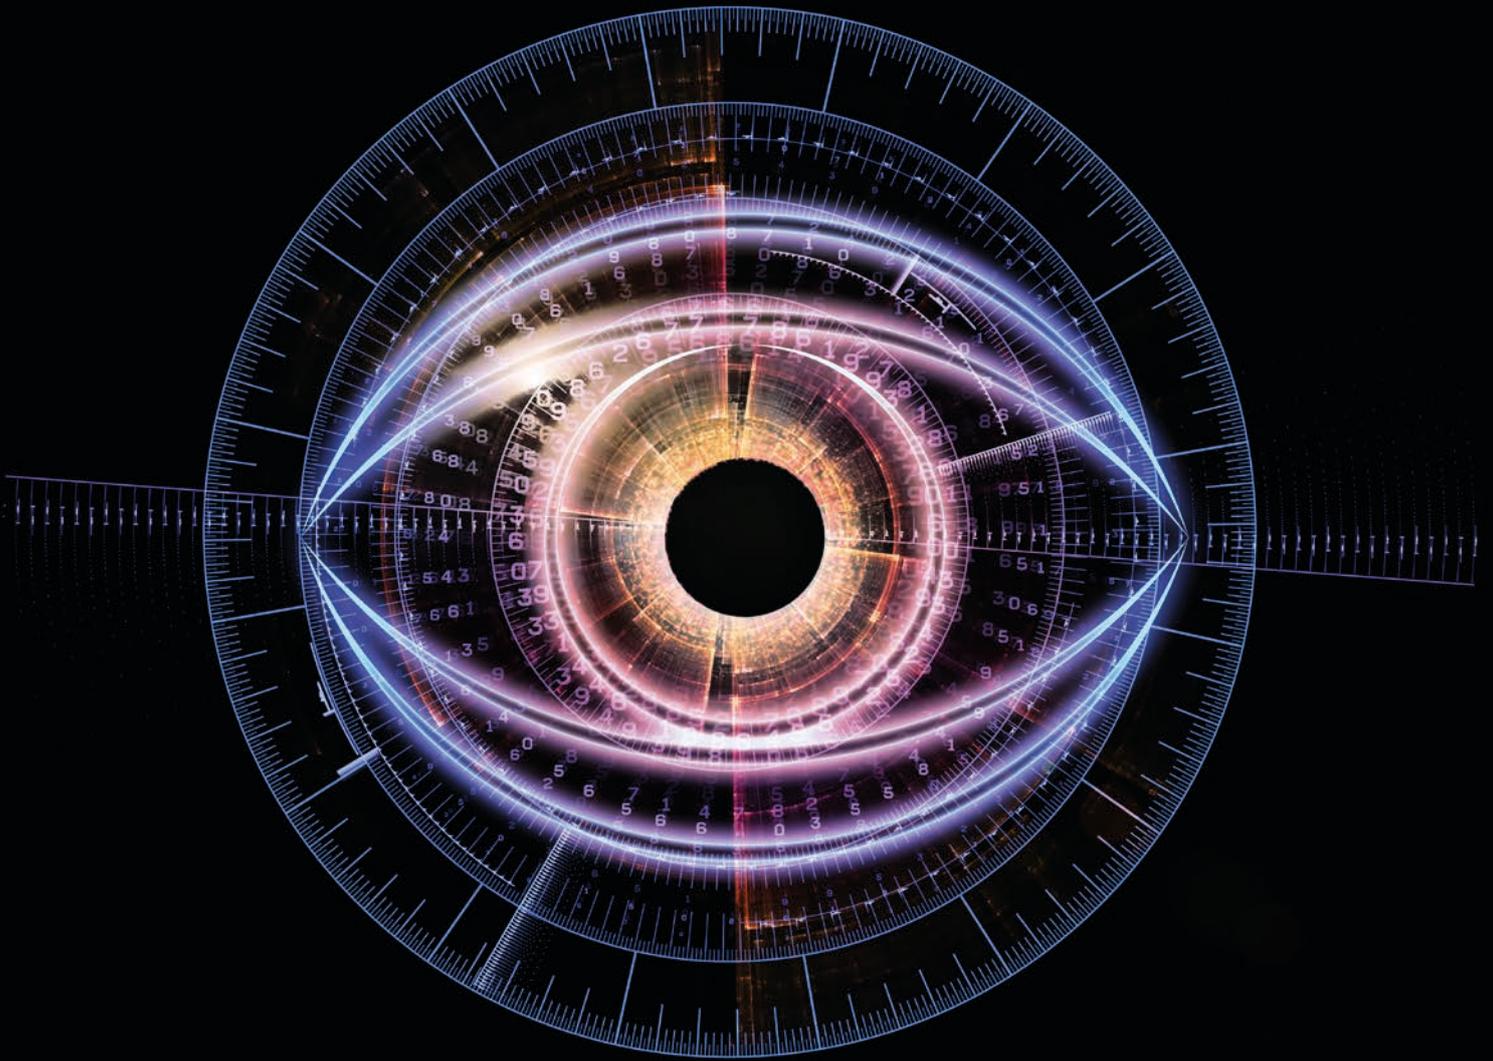


Deloitte.

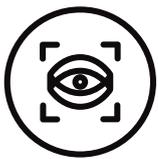


The time is now

Documentation Intelligence
in credit risk management

Contents

The benefits of Documentation Intelligence for credit risk management	01
Case study	04
Applications	06
Ease of implementation	07
Next steps for Documentation Intelligence	09
A tale of two banks	11
Speak to us	12



The benefits of Documentation Intelligence for credit risk management

The time has come for credit risk departments to do more for less

The rise of intelligent automation

Credit risk people the world over have grown used to hearing colleagues extol the benefits of end-to-end process automation. Particularly in banks' operations functions, cognitive automation and straight-through-processing (STP) have revolutionised the customer experience, fortified profitability and enabled sweeping programmes of cost reduction.

Over in the risk management world, the take-up of new technologies has been somewhat slower, to say the least.

Credit risk teams tend to work across a myriad of systems and processes with frequent human interaction and paper-based documentation. Just look at what a typical credit team does during the course of a day. There will be many high-skilled specialisms such as new business prospecting, financial spreading, credit analysis and capital planning. But many of these activities will likely involve a range of resolutely low-skill components such as data entry, copy-paste and office software formatting.

The smaller size of an average risk team (be it in the first or second line) also means that economies of scale exploited by huge operations functions can be harder to realise.

But things aren't what they used to be. Banks are actively considering the future of risk and how best to transform their risk operating models.

We have proven to clients that we can implement cost-effective automation using the tool we call Documentation Intelligence. It delivers the benefits of STP but without the need for huge up-front spending and without taking on nerve-shredding IT project risk.

Let's start with the three big reasons why STP fails in credit risk:

- systems don't talk to each other;
- paper trails still proliferate and processes depend on reams of unstructured data such as company financials, contracts, loan agreements, covenants or valuations; and
- you need a human to do something smart at regular intervals.

Cost-effective technologies can now address these three.

Cognitive automation is extending the reach of robotic process automation (RPA). The tactical use of RPA can get your systems talking to each other as you wait for the group-wide strategic solution to arrive. Because, let's be honest, that much-fabled strategic solution might never end up happening – which is why RPA isn't just tactical. It can be, and often is, strategic.

Optical character recognition (OCR) can digitise paper documents and make them readable by your systems, helping to automate at low cost large swathes of data entry work which is time-consuming, unrewarding and which humans often do poorly.

Finally, artificial intelligence (AI) and machine learning (ML) can replicate significant parts of what trained credit teams currently do, such as reading and interpreting financial statements.

We have integrated these discrete, modular technologies into a single tool:

Documentation Intelligence

Our cloud based solution takes standard RPA and OCR to the next level of sophistication by introducing machine learning powered extraction and analysis, combined with our business process expertise.

It helps you to transform complex credit risk processes that currently rely on data in unstructured documents and time consuming human analysis. This single tool extracts data and aggregates into the required format to increase speed and quality. As a result your credit risk staff can do what they do best; assess credit risk and serve customers.

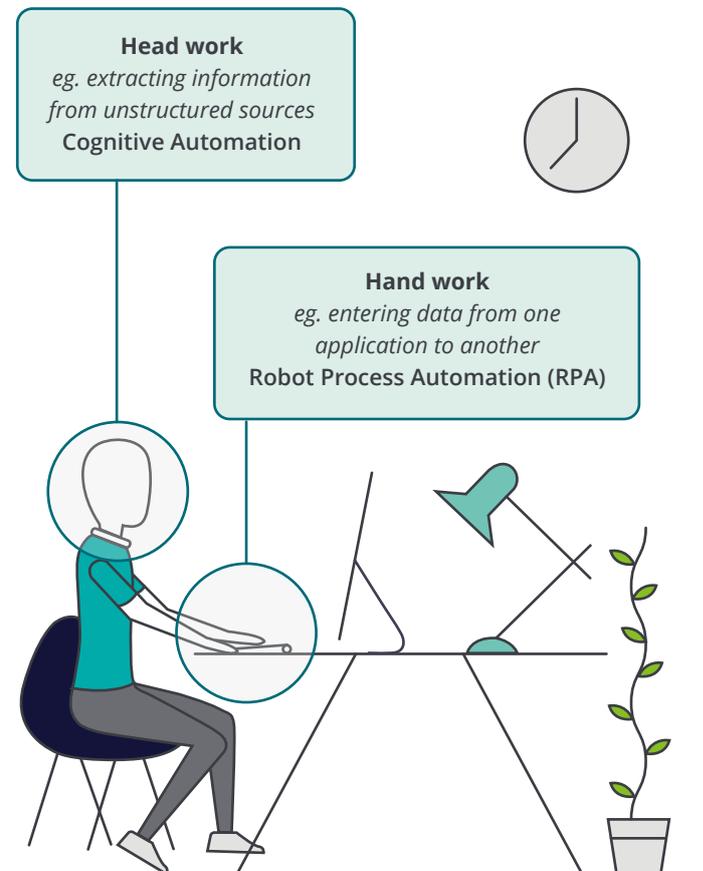
Clients who have tried Documentation Intelligence soon discover the wide range of advantages.

Customers experience superior service

- from start to finish, credit decisions can take minutes, not months;
- your relationship managers have more time to spend on cultivating relationships in place of filling in forms; and
- there's less need to ask clients for paperwork that you can source and process on their behalf.

The bank can make better lending decisions

- cheap and consistent data can be used for models;
- credit professionals have more time to conduct 'what-if' analysis and benchmark peers;
- banks have the capacity to consider a greater range and volume of credit proposals, deploying their capital more effectively; and
- the underlying data is more accurate, with fewer lending decisions made on faulty numbers, erroneous decimal places or inappropriate currencies.



There are widespread efficiency and profitability gains

- better-informed lending decisions made in greater volumes at lower cost will feed through to higher profitability; in one proof-of-concept that Deloitte ran, we demonstrated that our client could save £2.5m-£7m over three years by automating just one of its credit processes and for part of the client's portfolio;
- more efficient processes allow staff to be deployed on activities that create genuine value; and
- potential to reduce use of off-shore contractors.

The bank can attract and retain higher calibre staff

- credit risk professionals get to spend more time performing highly-skilled, rewarding and productive activities; and
- a more profitable bank can afford more competitive salaries that in turn support recruitment and talent retention.

Digitisation enables a host of downstream benefits

- all the data in original loan documents becomes storable and searchable;
- your staff can now perform better trend analysis and benchmarking, using improved data quality to make better decisions, helping you moving from being reactive to proactive;
- other credit processes such as risk reporting, monitoring of covenants and compliance audit trails can be automated; and
- firms can commercialise the data by selling market insights.

Banks which adopt Documentation Intelligence also find that implementation is much easier than anticipated while the scope for deployment is far wider. Once a bank can seamlessly spread the complex accounts of multi-national, conglomerates, there are multiple uses of the results, ranging from model validation and stress testing to benchmarking and targeting new business.

The European CRO of a leading global bank, has tried it out for himself. "The end results are really impressive and I'm very pleased we have something that has such a demonstrable impact not just as a specific product but also in terms of further application. It has been presented to our management team with very positive results and feedback."

We explain why implementation can be so quick, cost-effective and unobtrusive on page 7. Next, though, we look at some of the use cases we are working with clients on (both in and outside of risk), including:

- Financial spreading.
- ISDA statements and collateral management.
- Surveyors' reports.
- Self employed income verification returns.
- Credit assurance.
- Asset quality review audits.
- Tax returns.
- Non-standard mortgage underwriting.
- Bank statement analysis for affordability assessment.

"The end results are really impressive and I'm very pleased we have something that has such a demonstrable impact not just as a specific product but also in terms of further application. It has been presented to our management team with very positive results and feedback."

European CRO, leading global bank



Case study

Financial spreading

A wholesale banking client wanted to improve the efficiency, reliability and productivity of its credit underwriting function. In order to work out what commercial terms it would offer to new and existing clients, the bank required relationship managers and their assistants to prepare extensive datasets from each client's financial accounts.

Unstructured and semi-structured financial data needed to be keyed into a number of the bank's systems, which were then used to generate a raft of metrics such as EBITDA, eligible collateral, loan to value and predicted profitability versus hurdle rates.

The process was cumbersome and prone to errors, inconsistent application and unexpected delays. Off-shoring had been trialled to bring down costs but the service level agreements (SLAs) sometimes led to turnaround times of up to ten days and the final quality was still some distance from great – hardly a way to impress prospective customers.

Even when it had been input, the data was hard to manipulate or share with other systems for other uses.

Using a combination of OCR, AI, RPA and machine learning, we were able to automate significant parts of this process. Our Documentation Intelligence system quickly achieved accuracy rates of over 90 per cent, meaning that automation was at least as accurate as financial spreading by humans and could be done at a fraction of the cost and time. We expect that our machine learning will soon outstrip human accuracy rates.

Some processes which had previously taken up to 10 days under the old SLA were performed by Documentation Intelligence in just three minutes.

The system was trained to be able to extract and standardise data from a wide range of company accounts and financial statements, displayed in highly disparate formats and using all manner of nomenclature. We used machine learning to teach the system about the underlying patterns of company accounts. At the easy end of the spectrum, the system had to learn that 'sales' 'turnover' and 'revenues' referred to the same underlying category. At the harder end of the scale, it had to reliably locate all of the accounting items needed for the bank's lending decision tree and then apply those rules to the data.

By speeding up the process, the bank's credit sanctioners and relationship managers were able to improve the customer experience, consider each new loan in more depth and ultimately make better-informed lending decisions. The bank had not been targeting cost reduction in this phase of the project and chose to deploy the freed up time of its staff on searching out further opportunities to support its strategic growth ambitions.

While efficiency and customer experience were the main project drivers, Documentation Intelligence had a number of other benefits:

- **Prospecting**

Previously, there was a clear limit on the number of potential customers the bank could explore. Speculative customer research was too costly and time-consuming. By automating the financial spreading, the bank was able to industrialise the process of looking for wholesale customers.

- **Monitoring covenants**

Once the information on each customer had been digitised, it was straightforward to extend the process to the monitoring of the covenants agreed in the loan contract. Alerts could be automated and fed into other risk systems such as limit monitoring and risk appetite.

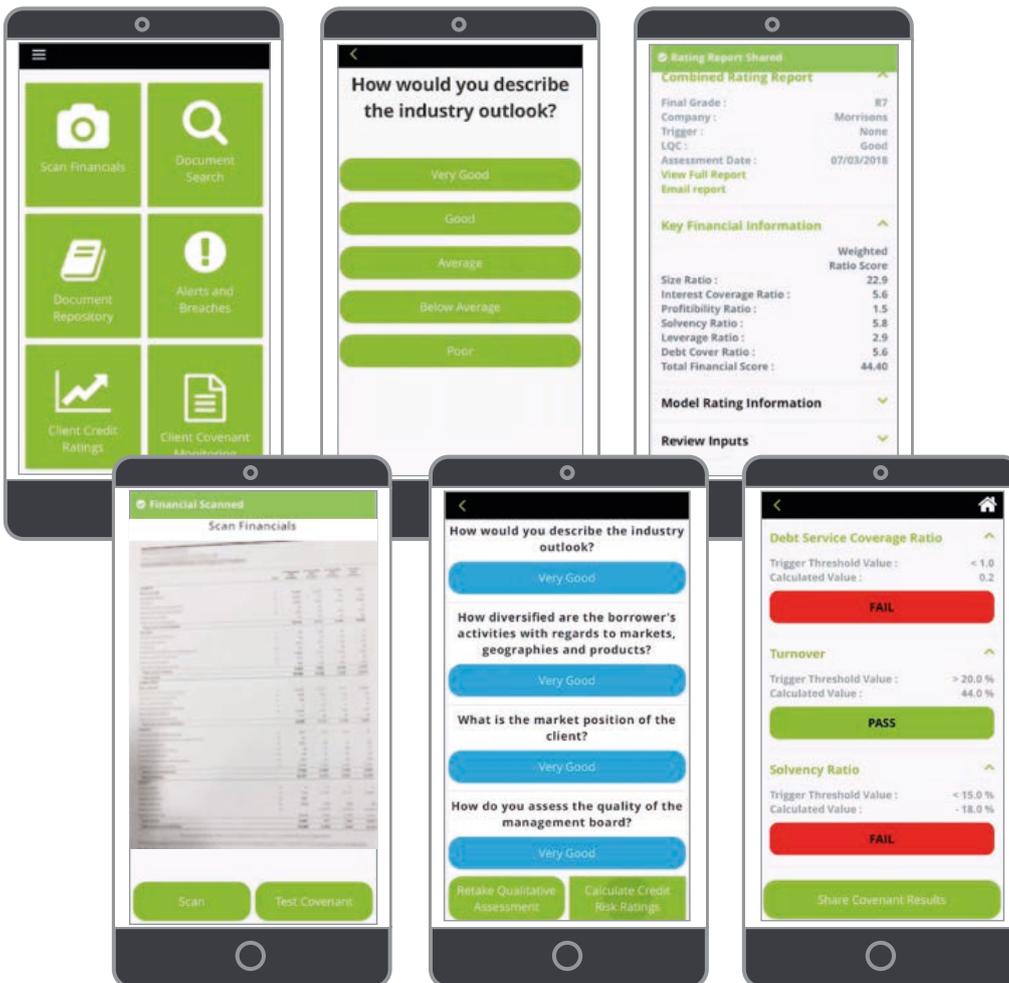
- **Covenant compliance documentation**

If the system can track, in real time, whether covenants are being complied with, it is simple to bolt on other, automated processes. Teams can automate how they notify, report and generate certificates confirming ongoing compliance with covenants. There is substantial scope to replace most of the mundane, manual process usually given to the credit underwriting function.

- **Dynamic annual reviews**

Automation of loan file reviews makes it much easier for a firm to stick to its policies. Overdue reviews can become a thing of the past. It's easy to set up on-demand or ad hoc reviews.

Documentation Intelligence of credit risk spreading integrates desktop, cloud and mobile computing. Its user-interface features all the intuitive, time-saving features that we've come to expect from smart phone apps. Navigation between pages is logical and our staff can search, scan, swipe, upload, download and share at the touch of a screen.



Some processes which had previously taken up to 10 days under the old SLA were performed by Documentation Intelligence in just three minutes.

How we applied machine learning to financial spreading

In order to automate the spreading process, we first deployed OCR to digitise a wide diversity of financial accounts. None of the OCR engines on the market were able to digest every type of document we threw at it, so we used three of the best and developed an algorithm that would let us predict which OCR engine to use on each style of document.

Linking up the outputs of the OCR engines to a comprehensive database of financial taxonomies, and using an initial suite of data classification algorithms, we produced a preliminary learning model. The model was able to classify by document type, section, page and line item. The accuracy of such models is rarely high enough to use, especially for sophisticated business cases such as the analysis of financial information which can be tabulated in a multitude of ways – all of which are perfectly compliant with accounting standards and principles.

We then ‘trained’ the first generation model with positive reinforcement when it was right, user-specified corrections when it was wrong and additional algorithms developed to overcome the specific challenges of company accounts. We ramped up the accuracy with each generation – incorporating rules on page position, font size and document hierarchy along the way – to achieve rates of accuracy that not only matched but exceeded those delivered by humans alone.



Applications

ISDA statements and collateral management

Most large banks have ISDA Master Agreements in place to govern the netting and reporting of derivative exposures they hold. Working out what collateral you hold is a logistical challenge for many financial institutions, let alone efficiently calculating where, how and when it could best be posted.

Documentation Intelligence can be trained to read and understand ISDAs. They are not only able to predict when it might be necessary to post or call upon additional collateral – but can also incorporate machine-learned algorithms to optimise the use of available collateral.

It's then only one small (robotic) step to add on the automated process to carry out the collateral management itself.

Non-standard mortgage underwriting

With the rise of the gig economy, banks should expect more and more of their customers to have non-standard proofs of income. Running the affordability checks on residential mortgages can be onerous. The credit sanctioners have to construct a miniaturised P&L for each potential customer.

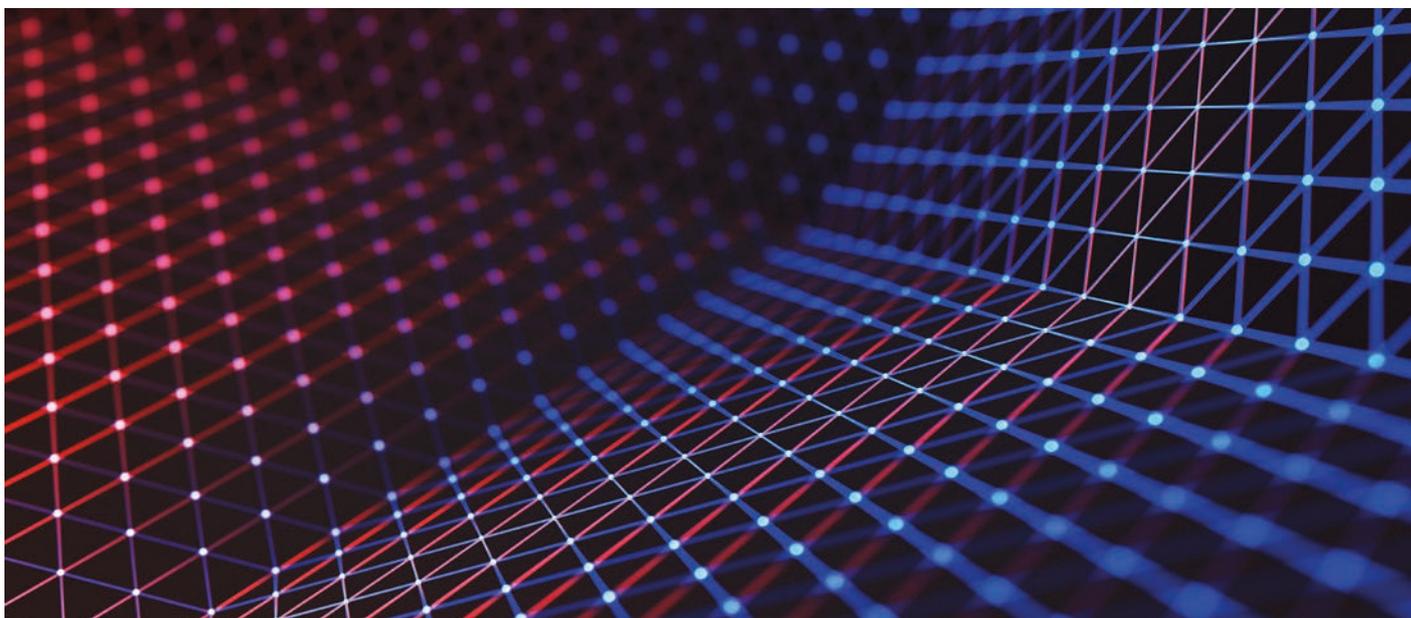
Documentation Intelligence can be configured to allow banks to read bank statements, credit card bills and tax returns to extract the data needed to analyse and compute each customer's risk profile: assets, revenue streams, liabilities and volatility of earnings.

Surveyors reports

Documentation Intelligence can help in any situation where companies rely on a mixture of systems, people and processes. As such, it has enormous potential across the whole of banking – particularly in those areas most affected by regulation and the need to evidence compliance. The processing of surveyors reports is a prime candidate for Documentation Intelligence because they are complex documents that come in a range of formats requiring reasonable amounts of experience to interpret.

An updated survey or valuation can be swiftly 'read' and used to update loan files, covenant monitoring, if all is well, or flagged to the Credit or Legal teams if not.

Documentation Intelligence can help in any situation where companies rely on a mixture of systems, people and processes.



Ease of implementation

The ease, speed and simplicity of installation generally come as a pleasant surprise. Documentation Intelligence interacts with your front end much like a human would.

In many cases, providing a user name and password for the system to use is all that's needed. Machine learning does require machines to learn, and that process can take time. But even so, the timescales range from weeks for a proof of concept to a few months for full implementation. Many of our clients are already on their RPA journey and we can align to firms' existing governance.

Plug and play

We offer a service and retain the delivery risk, aligning our incentives with yours. Development and delivery are based on a simple philosophy of cloud-based, open source technology stacks. All in all, we design Documentation Intelligence as an unobtrusive, user-friendly solution that your IT functions can smoothly implement to the existing estate (see following page).

Unlike some of the large tech vendors, our multi-skilled teams include colleagues steeped in credit risk experience and banking regulation.

Jan-Hinnerk Fahrenkamp is one of our directors who has implemented Documentation Intelligence for credit risk. "I grew up managing credit risk at banks," he says, "I always understood where my team made the biggest contribution to profits: generating

leads and getting deals through Credit Committee. Our goal is to help banks automate – intelligently of course – everything else. Everybody wins."

As new technologies come online offering more accurate OCR or next generation RPA, we upgrade the modules behind the scenes to serve you seamless, continual improvement.

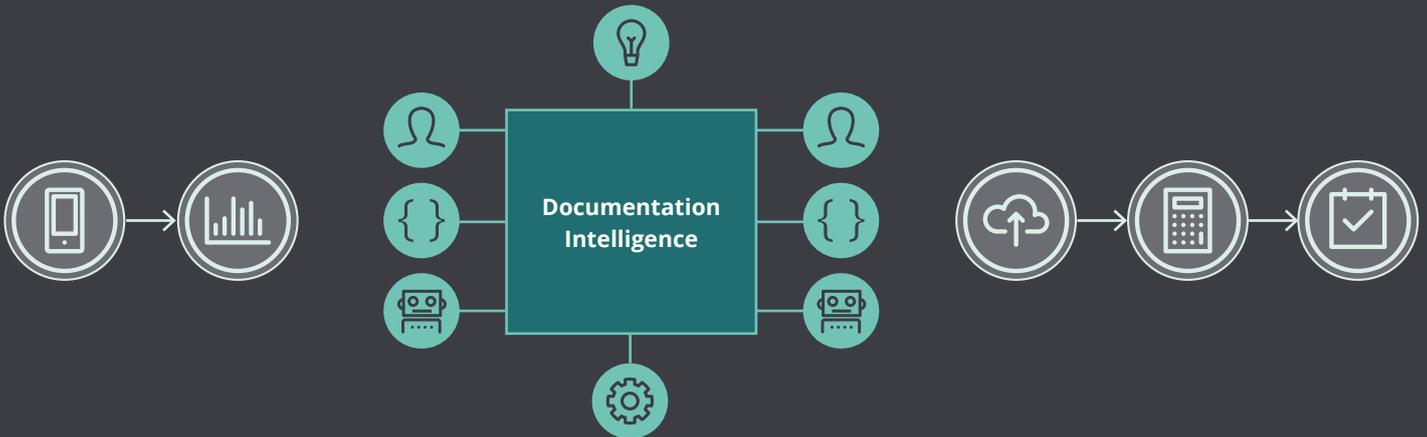
Clients purchase a licence to run the software, with little maintenance required. It's important to remember that as a bank's internal systems and risk policies change, so must the automation software. Furthermore, we shouldn't forget that Documentation Intelligence relies on human maintenance to process exceptions and keep the algorithms well trained. Clients generally find that it is more cost effective to buy the maintenance and exceptions processing as a service rather than expending in-house resource on occasional reconfiguration.

Moreover, as a managed service provided to clients across the industry, we can share the costs of regulation-driven updates across a portfolio of clients, reducing costs and easing the compliance effort.

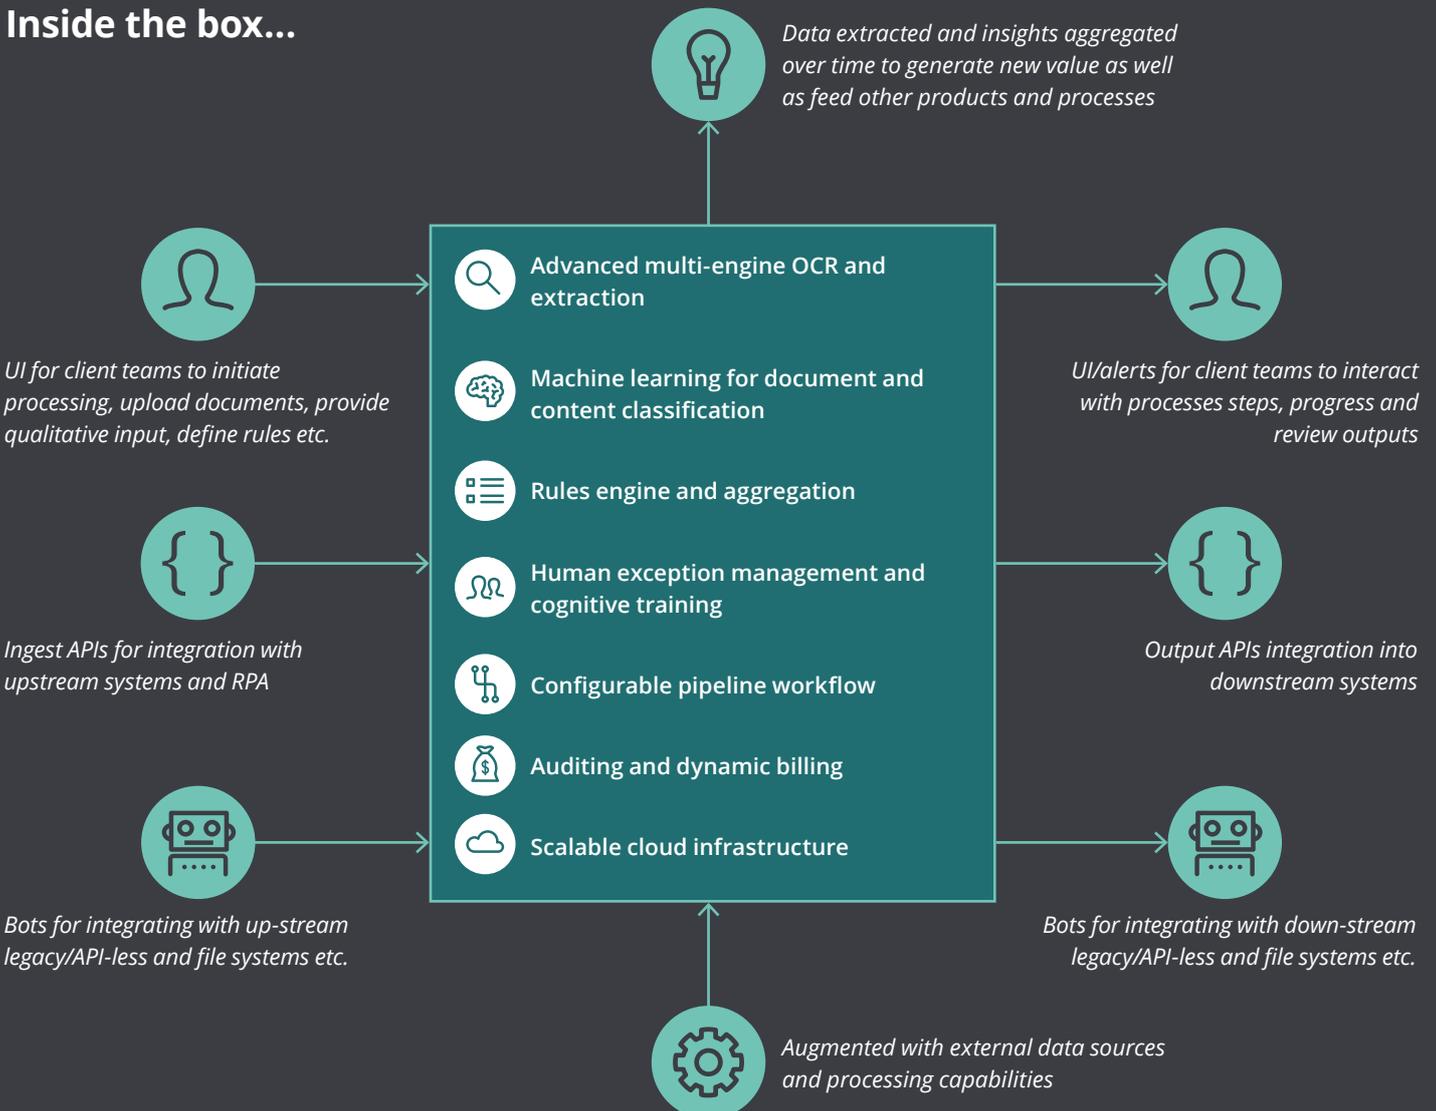
Documentation Intelligence

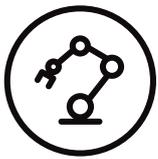
The driving force behind digital transformation is the digitisation of processes across every industry and within every business...

Documentation Intelligence is our new plug-in service that enables true end-to-end digitisation



Inside the box...





Next steps for Documentation Intelligence

As you consider how best make Documentation Intelligence work for your business, we've found that six considerations can help shape your strategy, speed up development and quantify success.

Gain consensus around the real business drivers

For some if not many firms, Documentation Intelligence is all about efficiency – time for the risk department to engage with the cost:income ratio. For some, it may be a productivity and capacity play. While for others, it may have more to do with sharpening the employee proposition (you've already made smart hires; you want to improve their productivity and keep them feeling busy, happy and loved).

Aim for high accuracy on unseen cases

Unless accuracy rates are high, your staff will spend more time correcting the system than they save on the initial data entry. Accuracy rates of around 90 per cent are highly desirable. Achieving those kind of rates takes time but is well worth it in the long run. Be wary of any solution that seems too good to be true. The true test of any system of document digitisation is its ability to apply what it has learnt to cases it hasn't seen. Documentation Intelligence produces confidence intervals around each automated data point, so that you have a clear appreciation of the thresholds of accuracy.

Hoard financial data

Loan applications, customer accounts, tax invoices, financial reports, sector analysis and Companies House filings all contain masses of information. You may not know how to use it right now, but automate the extraction and processing and your teams will soon find a use. For example, some of our clients have found that back-testing their low-default models became much easier when they could run their lending rules on rare historic cases of default and see if their own models would have captured the risk.

Keep your IT colleagues on board

We strive to make implementation as easy as possible, but the risk department is seldom the only stakeholder to consider. Keep the IT department up to date with your plans and learn from their experience.

Plan now for the next generation

The biggest change is without doubt the initial implementation of Documentation Intelligence. That said, you do need to consider in your current plans how easy it will be to update your systems, refresh the technology, keep the machine learning up to date and modify the RPA in line with general business changes.

Governance

Just because it's new, doesn't mean you can dispense with old-fashioned governance. The snag is that standard ways of controlling risks need to be updated. We have rolled out risk and control frameworks to manage Documentation Intelligence based on the following key criteria:

- *Decision accuracy*: Check that out of the box models are performing as expected and that hand-crafted rules provide sufficient coverage.
- *Live system performance*: Actual results may not mirror accuracies recorded in pre-production; and machine learning models may underperform when presented with new data.
- *Data integrity/bias*: Self-learning algorithms are open to data errors due to the small number of (potentially biased) contributors and human intervention.
- *Decision opacity*: The inability to rationalise the process behind a credit decision can expose firms to a number of regulatory and compliance risks.
- *Data security*: Access rights need managing and data encryption and data recovery plans need to be developed.

The true test of any system of document digitisation is its ability to apply what it has learnt to cases it hasn't seen. Documentation Intelligence produces confidence intervals around each automated data point, so that you have a clear appreciation of the thresholds of accuracy.



A tale of two banks

Success in business is about what you prioritise.

The following thought experiment traces two banks through very different decision trees. Spoiler alert: we want you to be 'Bank B'.





Speak to us

Damian Hales

Partner, Risk Advisory
dahales@deloitte.co.uk

Jan-Hinnerk Fahrenkamp

Director, Risk Advisory
jfahrenkamp@deloitte.co.uk

Jack Webb

Senior Manager, Risk Advisory
jawebb@deloitte.co.uk

Joel Miller

Senior Manager, Risk Advisory
jgmiller@deloitte.co.uk

Catalina Hallett

Manager, Risk Advisory
challett@deloitte.co.uk

Report Author

Dan Oakey

Associate Director, Risk Advisory
doakey@deloitte.co.uk

Paul Garel-Jones

Partner, Risk Advisory
pgareljones@deloitte.co.uk

Andrew Berry

Director, Risk Advisory
andrberry@deloitte.co.uk

Graeme Carmichael

Senior Manager, Risk Advisory
gcarmichael@deloitte.co.uk

Charoula Stamou

Manager, Risk Advisory
cstamou@deloitte.co.uk



This publication has been written in general terms and we recommend that you obtain professional advice before acting or refraining from action on any of the contents of this publication. Deloitte LLP accepts no liability for any loss occasioned to any person acting or refraining from action as a result of any material in this publication.

Deloitte LLP is a limited liability partnership registered in England and Wales with registered number OC303675 and its registered office at 1 New Street Square, London EC4A 3HQ, United Kingdom.

Deloitte LLP is the United Kingdom affiliate of Deloitte NSE LLP, a member firm of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"). DTTL and each of its member firms are legally separate and independent entities. DTTL and Deloitte NSE LLP do not provide services to clients. Please see www.deloitte.com/about to learn more about our global network of member firms.

© 2019 Deloitte LLP. All rights reserved.

Designed and produced by 368 at Deloitte. J16512