



DOCULABS

Intelligent Automation:

*Capturing
the Unicorn*



Overview

This eBook explains how to incorporate a new approach into your process management strategy. It outlines the scenarios where the technologies and products fit best, how to include them in your multi-year roadmap, how best to service and govern them, and how best to try before you buy with a focused Proof of Concept (PoC).

3 Takeaways:

What are differences among the various intelligent automation technologies?

How should I combine them?

How should I roll them out over the next few years and how should I manage and govern them?

Table of Contents

Defining Intelligent Automation	<u>3</u>
Understanding Your Capture Operation	<u>4</u>
What You Can Expect To Achieve	<u>6</u>
Overview of RPA	<u>8</u>
Defining Next Gen Process Automation	<u>9</u>
The Current Automation Platform Landscape	<u>10</u>
Intelligent Automation For LOB Processes	<u>11</u>
Intelligent Capture In A Perfect World	<u>12</u>
Machine Learning and Data Science: The Promise	<u>14</u>
The Results: An Example	<u>15</u>
Measuring Performance.....	<u>17</u>
Conclusion	<u>18</u>

Defining Intelligent Automation

Let's start by talking about the unicorn. The unicorn is the combination of AI, IC, RPA and workflow to create intelligent process automation that's worthy of the name. Until recently, intelligent process automation was a joke: most of the tools weren't very intelligent. They addressed tasks not processes, and they automated little because they had limited application and required human assistance. But now intelligent process automation worth its name is achievable with the combination of RPA, intelligent capture, AI and next generation process automation. RPA automates tasks, intelligent capture digitizes content to feed RPA, AI makes both RPA and intelligent capture smarter and next gen process automation orchestrates and manages all the pieces in the process. In this eBook, the focus is primarily on Intelligent Capture.

Intelligent Capture (IC) is used to capture the information on documents, categorize those documents and pages and extract relevant data for further processing using AI with technologies such as OCR, NLP, computer vision and machine/deep learning. The primary capabilities for IC are: classification, extraction, validation and conversion. Let's focus on the primary capabilities:

Multi-channel Capture	Multi-format Capture	E-forms	E-signatures
Classification	Extraction/Enrichment	Validation/Q&A	Conversion/Export
Content Analytics	Process Management/Automation	Task Automation (RPA)	Administration, Monitoring, Reporting

Figure 1. Intelligent Capture - Doculabs Material

- **Classification** is automatically identifying what pages, what documents and what document packages the system is getting.
- **Extraction** is recognizing the data that's on the pages and possibly enriching it – e.g., by judging sentiment or how angry the customer is, given the letter they just wrote you.
- **Validation** is checking for completeness and correctness of the information and document sets.
- **Conversion/export** is transforming the data into whatever standard format the downstream systems and people need – e.g., converting customer name into ALL CAPS with NO SPACES or PUNCTUATION. I'll mostly be focusing on the first three rather than conversion.

Intelligent Capture represents a technical revolution in document processing. Let's look at your likely capture operations and where Intelligent Capture can help.

Improved Document Automation Tools

Improved recognition and learning are separate capabilities, though they are typically combined in newer products. The results are much better than what you got from the older OCR tools. First, the new tools are better in how they recognize the text – in how they classify docs, paginate pages and extract documents.

In a sense, they have changed from how young children read to how adults read. Young children focus on letters, then words in a painful slow process. Adults use more context, including existing knowledge, the overall document and page layout and look, what the document's probably about, etc. The new tools may use computer vision, pattern analysis and other methods to take advantage of more context. Second, the new tools are better in how they improve that recognition – in how they continuously improve the accuracy of their opinion in light of evidence.

The learning capabilities range from not-so-intelligent to brilliant:

- 1) The human supervisor must do everything (e.g., drawing bigger zones in the template).
- 2) The human must provide them with positive and negative feedback (“these are correct answers” or “these are wrong answers”).
- 3) It's a mix of assisted and automated.
- 4) The system can brilliantly improve its recognition automatically 100% by itself.

Understanding Your Capture Operation

Let's look at where next generation Intelligent Capture can help your organization. The first step is to understand your capture operation - it may include:

- A centralized operation using one of the popular capture platform vendors.
- Some significant part that is outsourced to a for-profit capture vendor that does it for you.
- Some decentralized ingestion that takes in documents and data from multiple locations and multiple technology channels – paper, mfd, email, fax, smart phone and portal upload.

Your capture operation may look something like *Figure 2. Document Flow & Exception Handling*. Half of the capture process is in black arrows (the stuff going downstream). The other is depicted in red, which is the error and exception handling; it is equally messy.

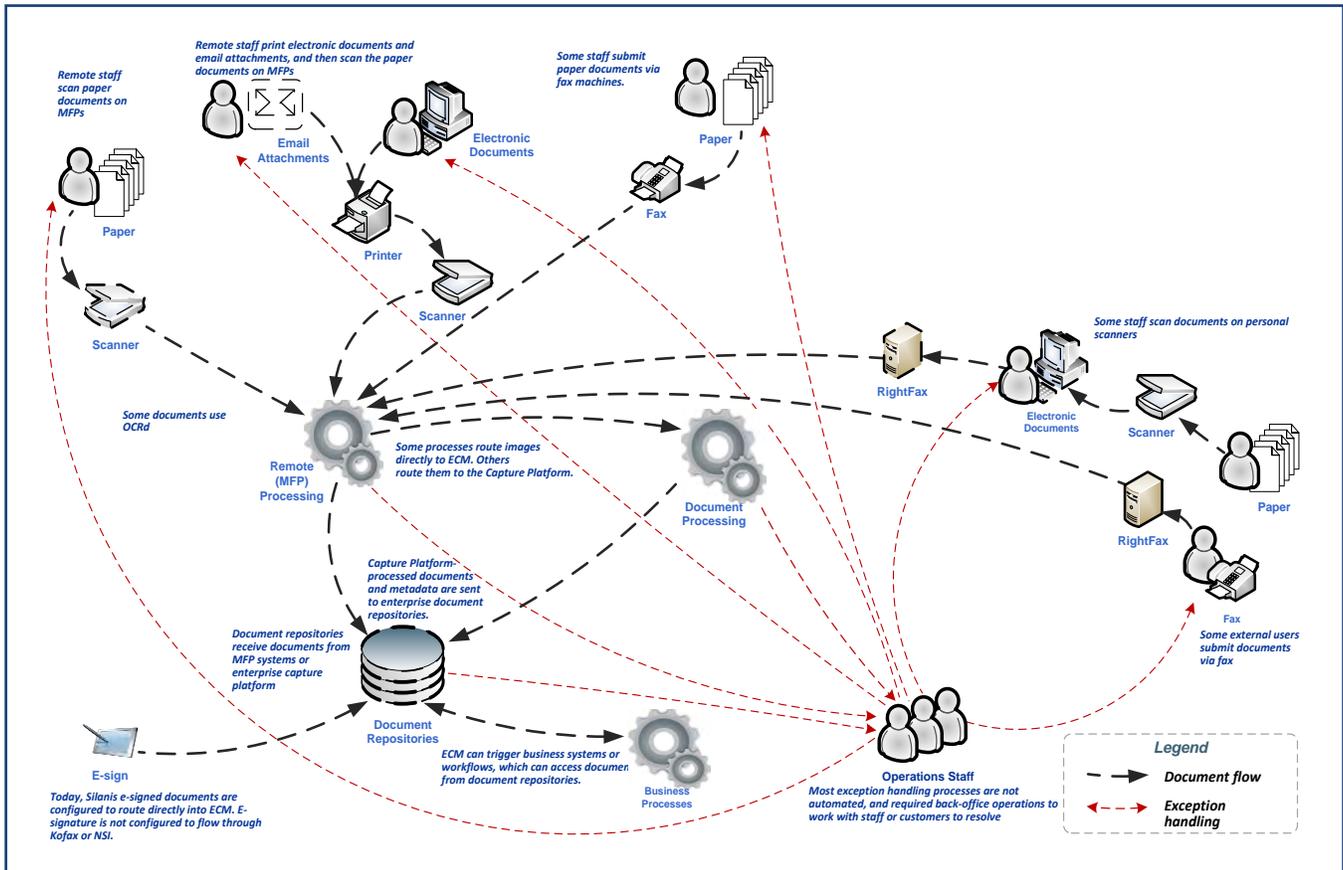


Figure 2. Document Flow & Exception Handling - Doculabs Material

It suffers from the same disorganization, but it is also too late, too damaging and too expensive. Often mistakes don't get addressed until they are far downstream, meaning lost time, unhappy citizens and inefficiency. It's possible to address these issues with document automation.

Many operations have literally been doing this since the late 1990s. For many of them, it's a mess. There are opportunities to greatly improve efficiencies in all these scenarios with the next generation intelligent capture tools.

Look For Labor Reduction

There is a real business case for using the next generation Intelligent Capture tools to reduce the labor in the capture workflow.

The following table illustrates some of the potential that’s just sitting there – here are some typical costs from inhouse and outsource capture operations:

Service	Unit	Ex. Peer	Industry	For-Profit	In-House
Sort & Prep	Per Hour	\$35.50	\$35.28	\$34.15	\$36.40
Scanning	Per Image	\$0.0546	\$0.0398	\$0.0298	\$0.0497
Data Entry (Onshore)					
Tier 1 (<10 Keystrokes)	Per Index	\$0.0650	\$0.0891	\$0.0665	\$0.1117
Tier 2 (11-30 Keystrokes)	Per Index	\$0.0950	\$0.2264	\$0.1630	\$0.2898
Tier 2 (ave. 50 Keystrokes)	Per Index	\$0.3750	\$0.4589	\$0.3460	\$0.5717
Retrieval/Lookup/Research					
Tier 1 Research	Per Hour	\$47.75	\$49.70	\$49.00	\$50.40
Tier 2 Research	Per Hour	\$63.75	\$58.50	\$59.40	\$57.60
Automated Recognition	Per Image	\$0.0225	\$0.0206	\$0.0132	\$0.0280

Figure 3. Typical Costs of Capture Operations - Doculabs Material

Doculabs collected these costs from 25 years of capture benchmarking research. Your costs are likely similar to what is displayed in the *Peer* column.

The rows are the various tasks that are performed in a capture workflow. Note *Automated Recognition* at the bottom – that’s cheap because it has very low labor costs. You can get significant improvement by moving the data entry labor to the *Automated Recognition* row. But that’s not all. You can also significantly reduce the sort and prep labor and the exception handling labor (what I call here the *Research* activities). General QA, error detection and error correction can also be reduced or avoided.

Many organizations only look at how they can use IC for reducing data entry, rather than the other labor activities in the capture workflow. They are leaving savings on the table.

What You Can Expect To Achieve

Now let’s address what kinds of problems these tools solve that the older generations of capture technology could not.

The table below shows easy, moderate and difficult problems in the primary document automation processes (classification, extraction and validation):

Complexity	Classification Capabilities	Extraction Capabilities	Validation/Q&A Capabilities
Easy	Page Classification	<ul style="list-style-type: none"> English machine text- structured layout Mark and signature detection Barcode recognition 	<ul style="list-style-type: none"> Document type, page number and order, signature presence, date validation, field completeness Data validation with source system, name/entry validation (against e.g. registry)
Moderate	Document classification <i>(single page type may belong to multiple documents)</i>	<ul style="list-style-type: none"> English machine text - unstructured Multilingual machine text - structured Multilingual machine text - unstructured 	<ul style="list-style-type: none"> Document field value consistency
Difficult	Package classification	<ul style="list-style-type: none"> English handwriting - structured English handwriting - unstructured Multilingual handwriting - structured Multilingual handwriting - unstructured Metadata enrichment (e.g. sentiment analysis) 	<ul style="list-style-type: none"> Signature correctness Package completeness Package field value consistency

Figure 4. Document Automation: Levels of Complexity - Doculabs Material

As a first glance, you can probably solve the easy problems with the solutions you have today, assuming you've kept up with newer versions of your capture software. Your biggest opportunity in the next couple years will likely be to focus on the *Moderate* problems and some of the *Difficult* problems. After that, you can consider tackling the more complex, *Difficult* tasks once you have that set in happy production.

For example, if you can do a good job at *Document Classification*, you don't have to do much document preparation. These new tools can automate - or near automate - this task better than previous tools. Once you tackle that, in a few years, you can take on package classification. This is where you identify a whole set of documents and prepare extraction and validation for what to expect.

In terms of extraction, you can really crank up accuracy on machine text well beyond 99%. You can also address handprint and handwriting, first in structured English and then later in unstructured other languages (e.g., Spanish).

For validation, you can check for completeness and correctness in entire documents – what's called NIGO (Not in Good Order) processing.

Later on you can do NIGO processing for entire packages of documents such as in a loan application. Or, check not just signature presence, but correctness.

So what are some typical use cases that would be worth pursuing with next generation Intelligent Capture that formerly were intractable?

- 1) **Change of address forms** with English, French and Spanish machine print and handprint, multiple checkboxes. The multiple languages, handprint and checkboxes were often non-starters, but are straightforward now.
- 2) **Survey forms** with 5 free-form comment boxes. You might want to separate and paginate the surveys and extract the comments. (We can also use the tools to do the simpler extraction tasks like barcode or machine text reading or use our platform capabilities for those tasks.)
- 3) **Invoices** in your AP process which come from a wide variety of smaller suppliers. We want to separate the invoices and pull the relevant data from them. This is a case where next generation Intelligent Capture brings automation from 85% to 99% - an incremental but significant improvement.

Overview of RPA

A natural place to begin digital transformation using intelligent automation is with robotic process automation (RPA). RPA tools automate steps of a process by mimicking the manual steps a human worker would take when using existing application software.

Multi-channel Capture	Multi-format Capture	E-forms	E-signatures
Classification	Extraction/	Validation/ Q&A	Conversion/ Export
Content Analytics	Process Management/ Automation	Task Automation (RPA)	Administration, Monitoring, Reporting

Figure 5. RPA or Task Automation - Doculabs Material

It's been highly successful in back office clerk activities in financial services and insurance, call center and typical swivel chair activities. Examples include document and data download, transaction processing; high-volume data entry, repeatable, computer-centric processes; as well as double and concurrent data entry into old and new systems during migrations.

RPA is quickly maturing and best used for repetitive and rule-based tasks. It's a significantly more sophisticated evolution from macros and scripts, and is often deployed tactically – as a standalone, 'duct tape' repair. Or, else it's deployed more strategically, with BPM or case management tools that manage entire processes.

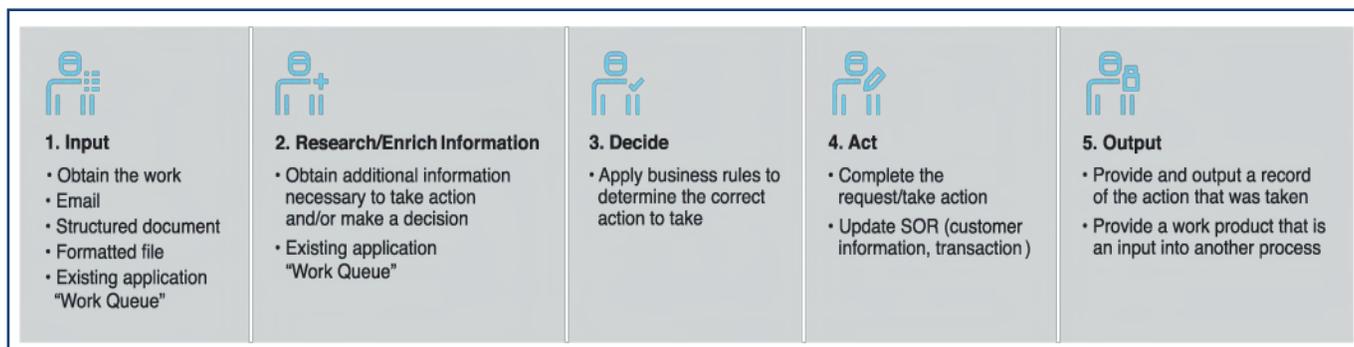


Figure 6. RPA typically includes answers to 5 steps: input, enriching, decision, action and output - Doculabs Material

By now it should be clear where RPAs are a traditionally good and bad fit, and where they are helped by AI and IC. They tend to work well with processes that are rule-based, simple to moderately complex, stable, mature and documented. If AI makes them smarter, they can do better with less structure and more complexity. They also historically required digital processes with structured data, and many processes otherwise suitable for automation do not have neatly structured data inputs, making them unfit for RPA solutions. But if you attach IC to the front end, you can address processes that include paper and dumb images.

Defining Next Gen Process Automation

Next Gen workflow emerged from past and current generation process management. It's useful to take a quick tour to understand what's old and what's new with it. Workflow has evolved from a few categories:

- **Simple Work Routing**
Basic request-submission-review-and-approval routing and tracking. Often addressed through lightweight ECM solutions and document collaboration tools.
- **Content-Centric Workflow**
Designed for more structured and repeatable processes that involve human interaction and document routing. Often addressed by the workflow capabilities of an ECM solution. Broader BPM-centric solutions also play in this space.
- **Business Process Management**
Designed for structured processes that may involve both human-to-system and system-to-system process orchestration. BPM differs from workflow in providing more advanced data aggregation, transformation, integration from multiple sources and business activity management (BAM) including process monitoring, analytics and simulation.

- Case Management

These solutions are designed for processes that are highly variable, dynamic and ambiguous. They rely on knowledge workers to determine how to structure the work throughout a complex “case” based on decisions needed, other teams to involve, information needed, etc. Most of the leading BPM vendors offer case management solutions or at least capabilities, as do most of the leading ECM vendors who provide content-centric workflow.

Digital Automation Platforms

A Digital Automation Platform (DAP) is a supplier and analyst-driven term for the newest process automation platforms. They include the capabilities of the preceding types, are low code and light weight and can be configured by *citizen developers*. DAPs improve and include characteristics of traditional workflow, BPM and case management. Benefits include:

- Maximize capability for building apps
- Maximize flexibility for updating apps
- Integrate to many information / content sources
- Built-in compliance and analytics

The Current Automation Platform Landscape



Figure 6. Decision Making, Self Service & Automation

New platforms have several advantages.

Where they come into play with the unicorn is in orchestrating, coordinating and managing various IC and RPA pieces, providing a much more efficient, stable and automated process:

- Aggressive Best-in-Class peers focus on using “light” platforms to improve decision making, self- service and automation – and to replace other legacy platforms.
- Leading solutions are comparable on core capabilities, but differ more significantly in emerging capabilities.
- Core capabilities include user experience, business rules modeling, case management, document support, process and flow design, analytics, integrations, enterprise scalability, etc.
- Emerging capabilities include AI and RPA integration, low-code, mobile engagement, customer journey metrics, etc.

Intelligent Automation For LOB Processes

The three major types of low-hanging fruit for these combined solutions are:

- 1) Automating everything better: from 80% to 99%.
- 2) Automating things you couldn't before: handprint and handwriting, package level validation.
- 3) Attaching to the front-end of RPAs. RPAs need semistructured text or data to work – paper and dumb images are indigestible to them. Stick IC to the front of an RPA and you can automate more. *Note* that you can insert IC and RPAs at any stage of a process, not just the front end.

What's an example of using RPAs within capture processes?

They can help Quality Assurance processes or exception handling, or where research is needed. In most QA processes, the bad document is sent to a human researcher who tries to resolve the problem by hooking into multiple different systems. It's very labor intensive.

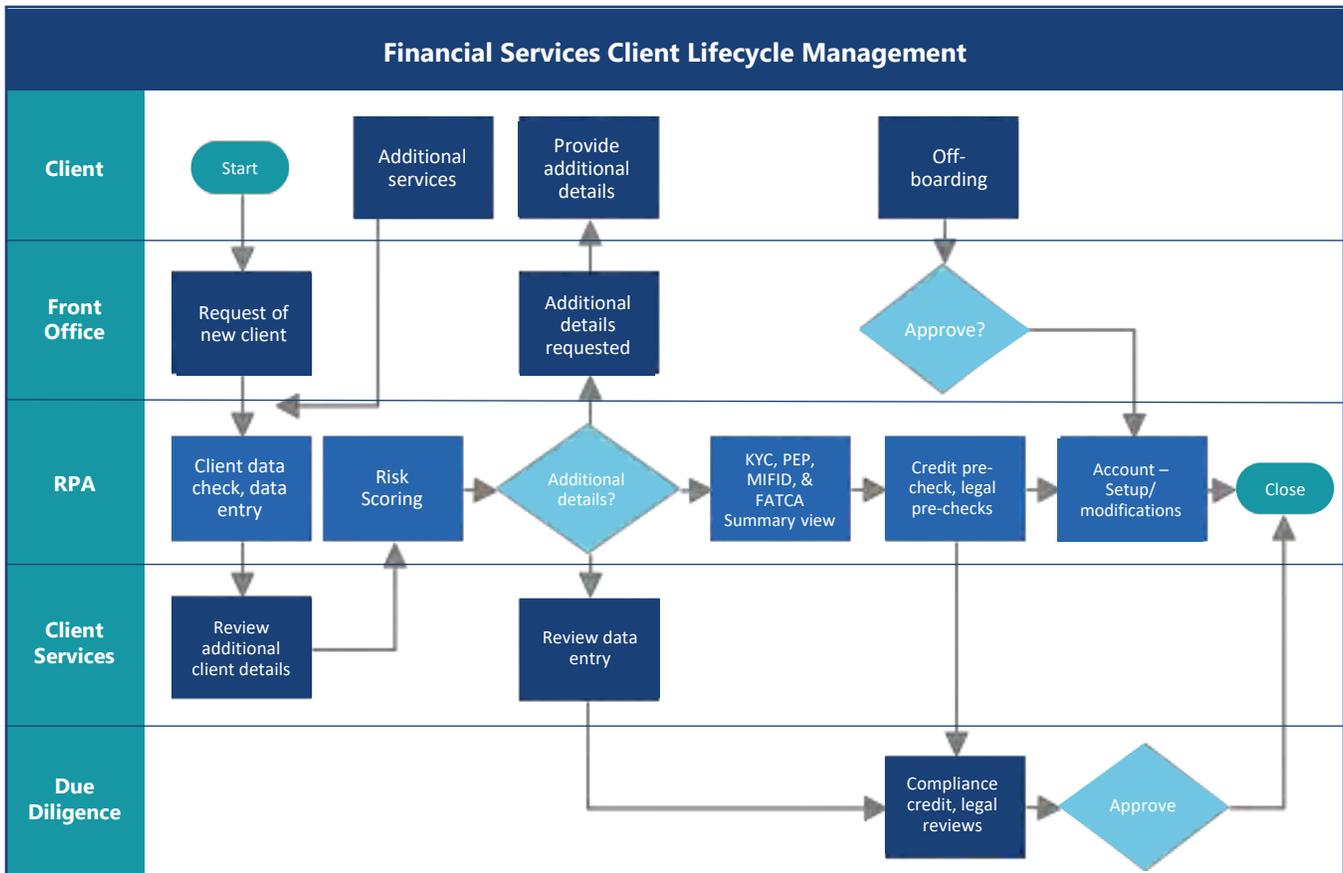


Figure 7. Financial Services Client Lifecycle Management - Doculabs Material

Do a Proof of Concept

So what should be your first step in your roadmap? We strongly recommend doing a Proof of Concept (PoC). Doculabs and Parascript created an [eBook](#) that walks you through the PoC process. In a nutshell, the POC process has 4 big steps, each of which consists of several smaller steps:

1. **General Preparation**, where you plan what you want from the POC and make some of the big decisions like who you're going to test and why;
2. **Detailed Preparation**, where you get everything ready for all testing;
3. **Actual Testing** of the software tools against your own docs and environment; and
4. **Evaluation** where you evaluate your results and decide what Next Steps you should take.

How to Govern Your Process Automation Initiatives

Another question is what kind of governance structure you should have as you expand your RPA and capture capabilities. We recommend that you start with an extension of capture or workflow, or with the RPA group that's trying to extend and grow to the next level. Mature organizations have a strong centralized matrix organization with oversight committee, cross-functional involvement in platform direction/roadmap, and alignment with enterprise initiatives and business unit priorities.

Intelligent Capture in a Perfect World

When it comes to intelligent capture, there's the perfect world scenario, and there's the reality. So in a perfect world, you get a hundred percent accuracy out of a hundred percent automation of your tasks without any need to have any type of human intervention. If we were to look at this in terms of its Intelligent Capture parts: the first one is the number of document tasks that you typically must deal with as part of your workflow. Those tasks can be things like reorienting images that are scanned in and doing document identification, separating one document from another locating information and keying it in all those representative visual tasks.

In a perfect world, you would have a hundred percent of these processes automated. Let's suppose that you've got every single one of your document processing tasks automated.

The second part of it is how many of those tasks that are automated are accurate. So that's the next step in this process is that you want to have every single one of those task results be accurate.

The third part is that of those tasks that you can automate and that are accurate, you want the system to be able to identify which can go straight through without any type of human intervention. In most cases, systems — even though they present accurate information — can't tell the user what data is accurate from what data is inaccurate. So, unfortunately in most implementations, it's often necessary to perform a hundred percent verification of that data.

Intelligent Capture in Reality

We've talked about the perfect world. Let's talk about the reality. So let's take that first step where we try to automate all our tasks. No system is perfect — the amount of automation you can get out of the system depends on how much time you're willing to put into configuring the system and testing it. The reality is that most systems — due to many factors — only automate a percentage of the tasks, leaving the remaining percentage to be manually verified and executed.

We start with the number of tasks that can be automated, and then we have the number of automated tasks that are accurate. Of correct automated tasks, only a percentage of them can flow straight through. The reality is that most systems require *manual verification* because of the level of effort necessary to optimize the identification, document separation, data location and extraction so to ensure correctness, few tasks achieve straight through processing.

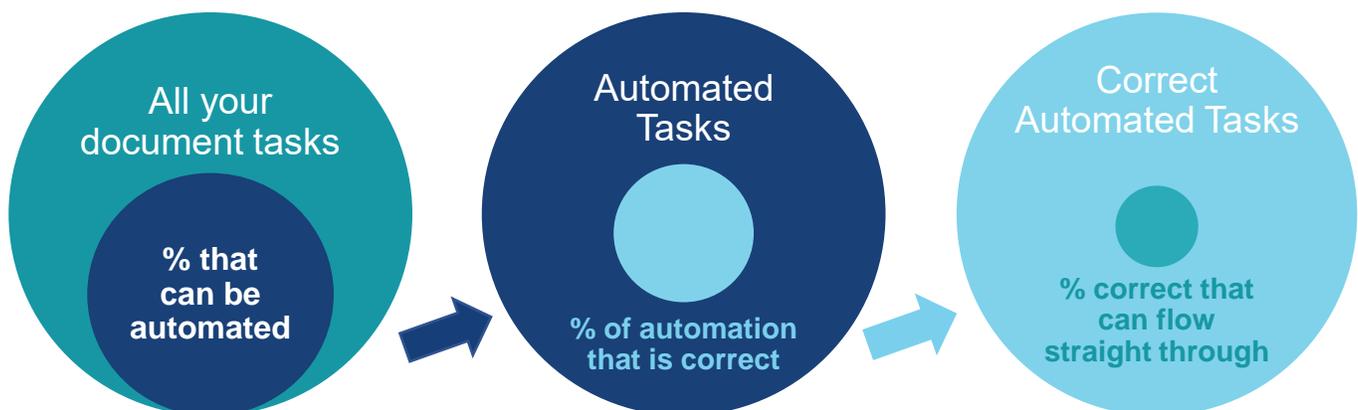


Figure 8. Current Capture Reality - Parascript Material

If the system can't tell between accurate and inaccurate data, then you're going to have only a small percentage of that data flow straight through – leaving the remainder of task to be completely manually reviewed.

There are many reasons why the actual system performance fails to meet expectations. While we always want a 100% perfection, the reality is that within a configuration, or setting up an intelligent capture system, there are many steps that must be completed very well.

The first step is data preparation, which consists of understanding the scope of your documents – not only the number of documents that you'd like to automate, but also the characteristics of those documents:

- Is there a high degree of variance in terms of image quality or data layout?
- Anything that affects the amount of comprehensiveness in terms of when you create and configure a system?
- Is it going to know how to deal with the documents?

Once you have the wide array of samples that you want to build a system around, you must take the time to configure it. Configuration can often require technical capabilities, and that often requires a significant amount of investment. The next step is testing and tuning, which is an iterative process of testing your output and optimizing it. Once you get it into production, the unfortunate reality is that things change. Documents change. Layouts change. You onboard a new client or customer who has a new type of document, for example, and this means you have to go through all these steps again.

Machine Learning and Data Science: The Promise

Machine learning is not a technology or a solution unto itself - although sometimes it's presented that way. However, machine learning offers real benefits when you combine it with data science. You reduce the amount of configuration, iteration and investment required for traditional intelligent capture solutions. Ultimately, data preparation, configuration, extraction, testing and tuning – and then operational maintenance – can all be automated and optimized as a compute-time operation without requiring direct human involvement. That's the great thing about applying machine learning. Leveraging machine learning and data science means that a significant amount of the data *can* be automated at high precision. This is possible when you have machine learning software crunching through enormous volumes of data. You end up with a greater amount of automated tasks performed at a greater level of accuracy.

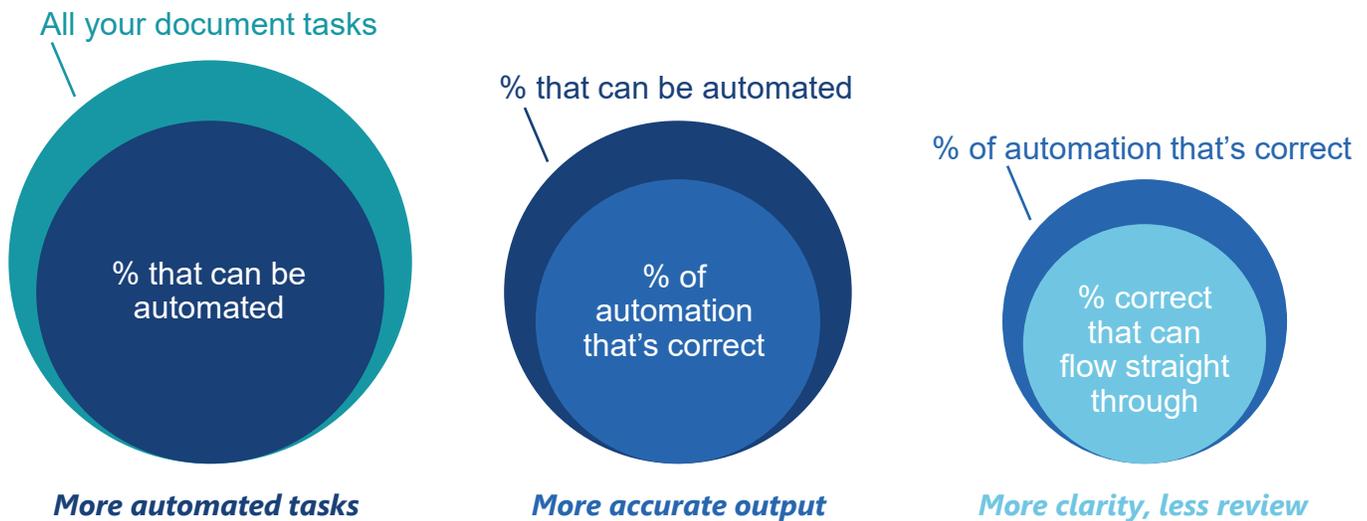


Figure 9. ML & Data Science. Leveraging machine learning and data science means that a significant amount of the data can be automated at high precision - Parascript Material

The Result: An Example

Let's look at a scenario using a real use case (refer to *Figure 10. Example Scenario* on page 16). We have an organization that must process manually 30,000 pages per day with 10 data fields on average per page. The human accuracy for manually processing these pages that they've analyzed is about 95%.

This means that there is 5% error, and the automation target (as always) should be at that accuracy level or ideally better. So looking at it from three perspectives we have on the left side, a hundred percent manual operation.

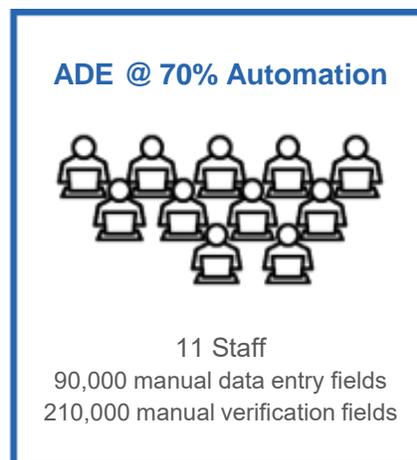
If you calculate it, and you can derive these in terms of how much time is needed to process each individual page and to perform manual data entry on a single field, it will take roughly the equivalent of 21 staff on a daily basis to do these tasks manually.

When we apply the traditional capture, we're talking about a lot of investment in terms of analyzing documents and configuring the system and then testing and tuning. Ultimately, what we get on average is from 60% to 80% task automation. As you can see here, we effectively reduce staff or the equivalent staff required by half, but all the data must still be verified. So we're not completely optimized yet in terms of the amount of tasks that are automated or the amount of accurate tasks that are automated. And still, we also don't know which data is good versus which data needs correction. Staff must still review data and do minor corrections.

When we apply a machine learning to the tasks of both analyzing and curating data and automating the system configuration, we see a significant uplift in terms of automation and efficiency with increased accuracy. As you can see here, we can move the needle effectively to getting about 90% of those tasks automated at 98% to 99% accuracy with the equivalent of two staff to take care of just that amount of information that needs to be reviewed. You can go from a 100% manual process, applying machine learning all along the way to automation that truly achieves high straight through processing.

Example Scenario

- 30,000 Pages/Day With 10 Fields Per/Page
- Human Accuracy: 95%
- Automation Target: At Or Better Than Humans (95% Accuracy Or Better)



Everyone Using Capture Is Here

Data Science + Machine Learning Can Get You Here

ADE=Automated Data Entry

Figure 10. Example Scenario - Parascript Material

Measuring Performance

We've talked about the power of machine learning, where it fits and where the promise of it lies within Intelligent Capture. It can take a process performing on average with about 50% efficiency and move it to 90% efficiency or better.

However, there's still the factor of measuring and knowing how your system is performing. There are really two important aspects of this.

Unattended Automation Rate

First, there is measuring overall system accuracy. When you're measuring system accuracy, you're measuring *all the output*. So let's say you've got 10 fields on each page. What you're doing is you're measuring the number of fields that are being output *correctly* or *accurately*. This allows you to understand how much data entry you can remove, but it doesn't let you understand how much of that data can flow *straight through the system*. This leads us to the second measurement, which is measuring the *unattended automation rate*.

Confidence Scores

This allows you to look at not only the total output of the system and how much of that output is accurate, but you also measure something called the *confidence score*. Now with intelligent capture, you often get a confidence score associated with each automated path. So you've got a confidence score for each field that has automated data entry, or you have a score associated with each document class assignment. Every automated task has a corresponding *confidence score*.

Thresholds

To assess the unattended automation rate, you're looking at accuracy and setting a reliable confidence threshold. This threshold is a reliable breakpoint in those confidence scores that allows you to accept all those answers above the threshold as accurate. You can measure this to a fine degree. Really, the ultimate goal here is to move from an average off intelligent capture implementation to something that's modern, almost completely unattended automation.

Conclusion

In conclusion, the mechanics of integrating intelligent capture tools like Parascript with RPA or other automation technologies depends on many factors including the workflow, the scope of the documents, etc. As discussed, in an automation workflow, RPA is typically involved with initiating the process as well as collecting data from different systems. It might even screen scrape some data off websites. And then, initiate a request for the import of document. Examples of this are a lending scenario or claims adjudication at the point of initiating the request of documents. This request might be a notification via a mobile app to take a photo of the document or through a web browser. This is where intelligent capture tools are important so those documents can be uploaded and presented through an API.

All RPA solutions provide some level of API to hand-off documents to other types of systems using Intelligent Capture software that can be configured to meet the requirements of systems. RPA is defined by the process. If you can define the process, you can define the scope of documents; and if you can define the scope of documents, you can define what data needs to be located, extracted and verified from those documents. Intelligent Capture by virtue of being in the middle of that RPA process knows what it's likely to deal with and accomplishes these classic automation tasks such as classification of documents or separating those documents into individual discrete pages. IC can also do other types of analysis, like locating specific information and reporting back summaries. All this information is presented back to the RPA system through the API (or file-based mechanisms) so that as RPA has its necessary data, it continues along with its workflows. Typically it is a simple handoff. The RPA system handles all the other types of data synthesis and incorporates that within the workflow. It's relatively seamless and allows for much higher straight through processing.

About Doculabs

Doculabs offers Information Management Strategy Consulting for ECM, CCM and IRG. Managing unstructured content has emerged as both a risk and an opportunity in meeting operational challenges and profitability goals. Doculabs consultants understand and successfully navigate the deep complexities of unstructured content for its clients across the world. Doculabs consultants deliver trusted answers to your strategic and technology enterprise content management system questions. From analysis and development of strategy, to design, content cleanup and migration, our information management strategy consultants support our customers. Visit [Doculabs](#).

About Parascript

Parascript software, driven by data science and powered by machine learning, configures and optimizes itself to automate simple and complex document-oriented tasks such as document classification, document separation and data entry for payments, lending and AP/AR processes. Every year, over 100 billion documents involved in banking, government, and insurance are processed by Parascript software. Parascript offers its technology both as software products and as software-enabled services to our partners. Our BPO, service provider, OEM and value-added reseller network partners leverage, integrate and distribute Parascript software in the U.S. and across the world. Visit [Parascript](#).

For More Information:

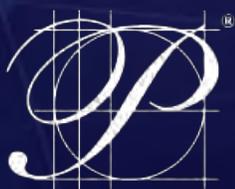
Doculabs

www.doculabs.com
312.433.7793
info@doculabs.com

Parascript

info@parascript.com
parascript.com
(888) 225-0169

DOCULABS



PARASCRIPT®

©2020 Doculabs & Parascript, LLC. All Rights Reserved.

