


How much does document processing **really cost** your business?

In this guide we explore the total cost of ownership (TCO) of invoice processing, which includes much more than manual data entry costs.





In this guide, we focus on everyone’s favorite subject: the cost. Because one way or another, the more business a company does, the more invoices its Accounts Payable (AP) will eventually deal with. We unwrap the real Total Cost of Ownership of three basic approaches to data entry – [manual data entry](#), [template-based OCR solutions](#), and [AI-based cognitive data capture](#).

But it’s not just about a particular price tag. We are going to unravel an analytical framework that you can take and apply to your business to dramatically reduce the costs associated with [invoice data extraction](#).

Part 1: Invoice processing can be cheaper




None of us enjoys the routine aspect of Accounts Payable (AP) operations. Pushing invoices around, entering data to internal systems, gathering approvals, and getting payments done. It creates little value but is essential nevertheless. It is “just a cost center,” but the costs can be cut. Its operation is a fixed process, but the process can be optimized.

With modern technology like Artificial Intelligence (AI) and Robotic Process Automation (RPA), the [AP department can be completely transformed](#). New approaches can turn around all activities, process by process. This improves nebulous qualities like flexibility and transparency, but also very down-to-earth ones – a bill’s time-to-process, staffing issues, and cost.

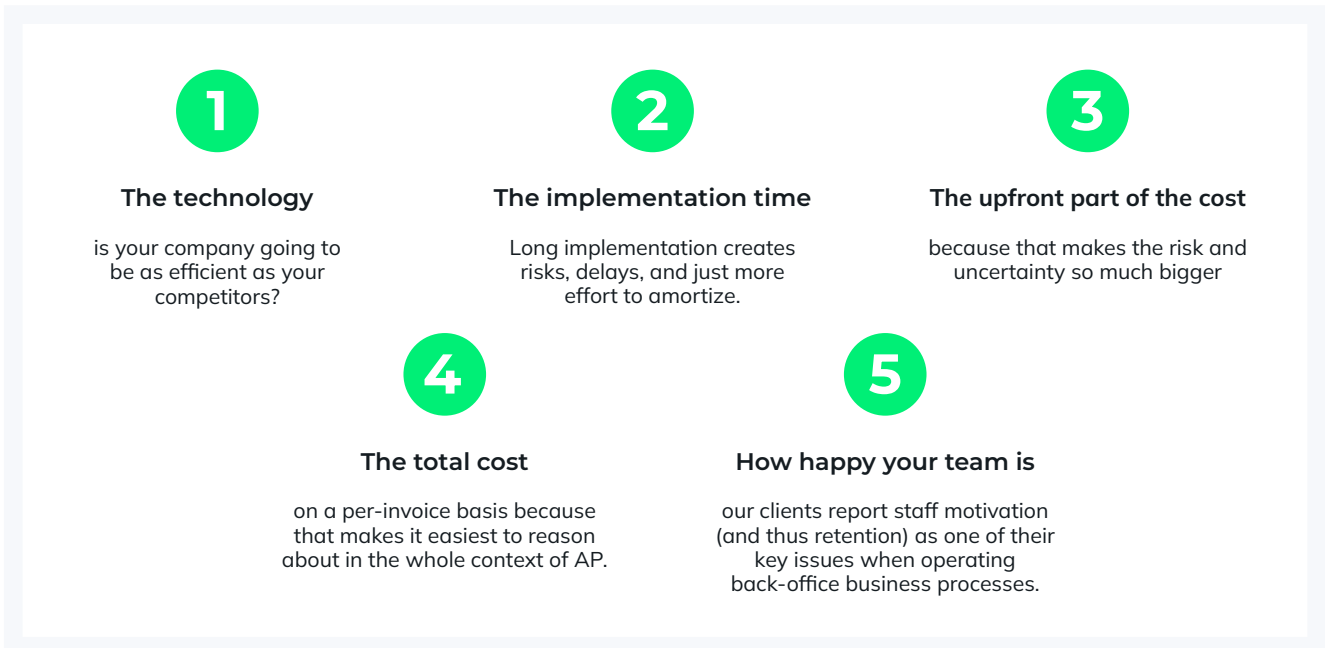
The actual TCO for OCR

We are going to start by comparing the three approaches head-to-head and delve into hard data. The numbers are based on a fairly typical model scenario. We will dive into all the details below – and also explain how all the numbers came to be. This should give you a blueprint to think about data entry TCO within your organization.

	MANUAL	TEMPLATE-BASED OCR	AI-POWERED IDP
Technology stack	Paper	On-Premise	Cloud & Browser
Implementation time	0	Months	Hours to weeks
Upfront & impl. cost	0	>\$100,000	\$0 to \$10,000
Total per-invoice cost	\$2.03	\$1.03	\$0.45
Workforce morale	Unhappy	Content	Loving



As you can see, the TCO in our eyes is not just the cost itself. We think about:



You simply need to have the staff do this work for the organization to function, and how much you spend comes second.

But now, let's delve into the meat of the matter: the technical analysis of the total cost.

How did we get our data?

At Rossum, we strongly believe in data, searching for the actual truth and transparency. Third-party data in this domain isn't always accurate, and published estimates routinely vary by an order of magnitude. Our estimates come out at the very low end of what's commonly published. Our data is more valuable to our purposes, as it's lifted directly from our experiences with clients and analyzing their results.

When a new client comes to Rossum with an interest to improve their [data capture process](#), our automation experts take a deep look at their metrics and experience and measure the efficiency of their invoice data capture.

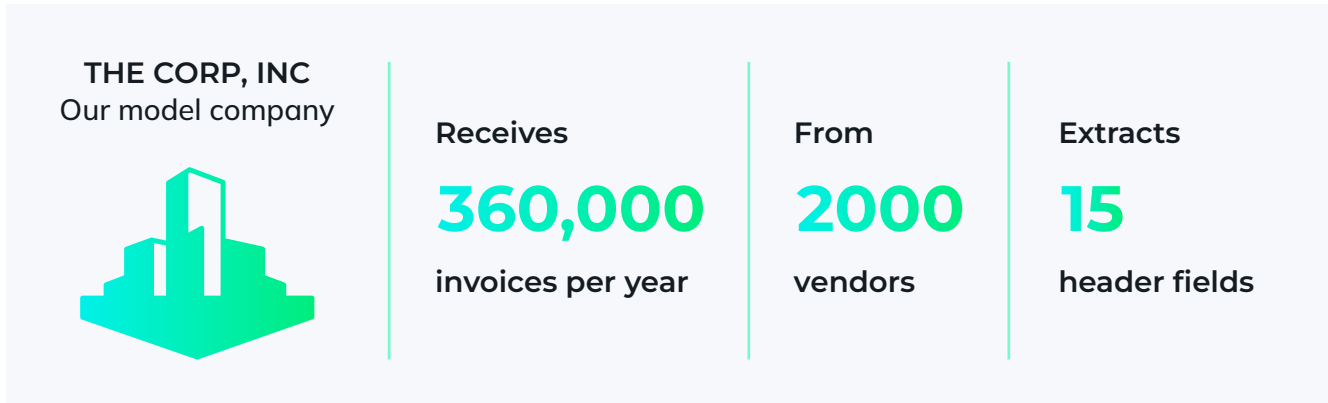
This helps the client to receive the best options for them, but it also gives us insight into the market of all the data entry processes in the world. And the insight on cognitive data capture comes from our own long-term production clients.

Meet our model company

For the purposes of the TCO calculation, we picked some particulars and calculated costs for a model example of an imaginary company. To give basic context to the numbers above, take this hypothetical into consideration:

The Corp, Inc. is based in Central Europe, operates in a retail segment, and is processing 360,000 invoices per year or 30,000 invoices monthly...

Their invoices come in as digital PDFs and scanned images (we are excluding EDI, supplier portals, or P2P platforms) from 2,000 different suppliers, with a natural supplier rotation of around 20% per year. Invoices are, on average, 7.8 pages in length, but line item data is not necessary. This makes it just 75 data fields to be captured. The fields are mostly numerical (ranging from amounts to dates and VAT identifiers), averaging 6 digits per field.



Invoice processing may involve up to three other departments that need to be contacted to validate/approve information (procurement in particular). For the Head of AP and managers, we estimate yearly wages of \$48,000; for the operators transcribing the data, we will go with \$24,000. In this comparison, we are going to use fully loaded costs, meaning all direct and indirect costs of human operators included – from the hiring and training costs to computer and office overhead to managerial costs stemming from the company structure. Fully loaded costs are therefore \$96,000 and \$48,000, respectively.

One-time costs of new technology can be quite significant – to be fair, we dissolve them over a period of three years. That would be the typical amortization horizon with the rapid pace of new technology in our model company.

Part 2: Manual typing is expensive



Meet the main metrics: Keystrokes and FTE Load

How much an invoice is going to cost is now really a function of two things, how much you **pay for the process itself**, and how much **effort your workforce spends on operating it**. Let's think about the latter for a minute.

What does "data entry" actually mean in practice? We consider the time spent per document, per time period, per employee as a valuable way to measure data entry. On principle, it may sound naive at first, but we have found that "how much you type" is a great proxy for the effort you spend on the data entry task.

Keystrokes

78
kpm



250
kpm

Let's measure the **number of keystrokes**: the number of times a human operator needs to input information, using a keyboard, mouse, or other means of physical input. For example, retyping 723456789 takes 9 keystrokes, meanwhile validating an already extracted number 723456789 can take just a single mouse click.

Next, let's think about the keystroke speed of the operator. In practice, the typing speed we have measured for a Rossum client when entering field values from scratch was 78 kpm (keystrokes per minute).

Do you feel that's too slow? Let's assume a base keystroke speed of 250 kpm for continuous typing. But the snag is that invoice data entry is very much not continuous typing. Instead, every individual invoice field comes with a fixed time overhead – to move to it in the user interface, locate it on the page, take in its value first, etc. The 250-to-78 kpm difference corresponds to 3.8 seconds per field of fixed overhead. This is realistic: when going "maximum speed": it would be difficult to jump field to field with less than 2 seconds per field. When bogged down in such routine and repetitive tasks all day, every day, we can assume twice the time is reasonable, before the error rate spikes.

We are very lucky that our model company has such a simple data entry case – we did not even consider searching internal or external databases, communicating with other departments, and resolving discrepancies.

The Full-Time Equivalent (FTE) Load is the effort of one employee as if working full time on data entry tasks. For example, if 9 employees each work on data entry one-third of the time, you would have a data entry FTE Load of 3.

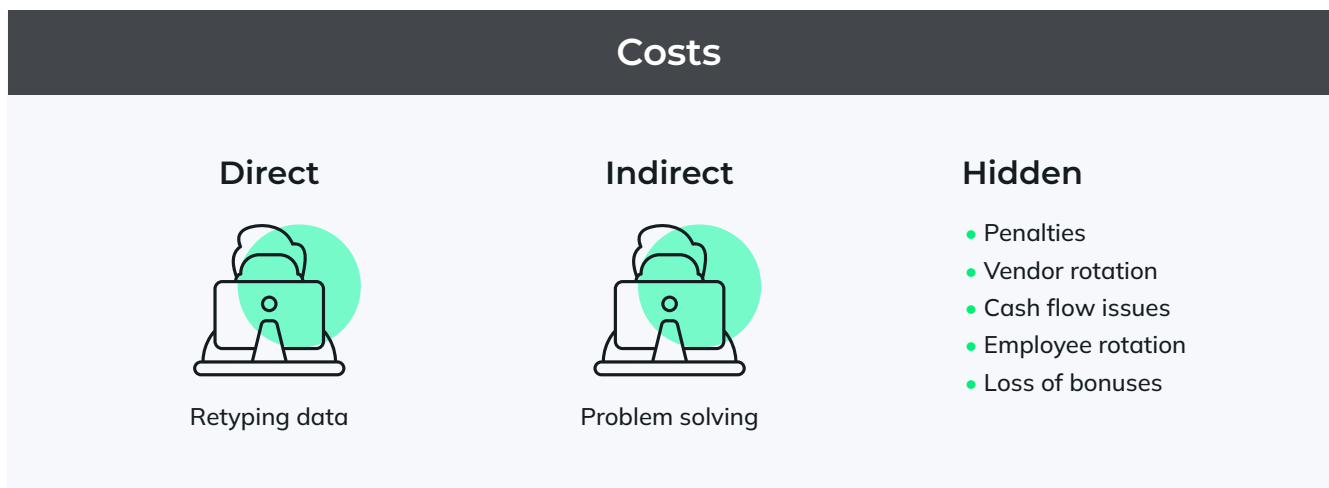
The key contributing factor to FTE Load in a company is the number of keystrokes necessary to process every invoice field on each document, combined with the keystroke speed of the operator. We will assume the fixed **78 kpm** speed in our analysis, even though the speed of the system and the user interface quality is doubtlessly going to be an advantage for modern tools. We are further assuming **75% time efficiency** for a back-office worker in our model company.

Not all of the nominal time spent will ever be productive time, and out of 8 working hours, at least 2 hours will be spent on inevitable overhead – from coffee and stretching to workplace discussions and administrative agenda.

The FTE Load then obviously directly influences the direct cost of the process – how much you spend on your staff doing data entry is simply the FTE Load times the employee wage.

The cost structure of AP data entry

When it comes to assessing the costs of manual retyping of the data, it is easy to fall into the trap of taking only direct costs into account. By multiplying the time employees spend manually entering the data by their wage, we can conclude the cost of data entry fairly straightforwardly. But the structure of costs is more complex. Let's take a look at how it works at our model company.



The actual costs

In the Total Cost of Ownership, we need to consider not just direct costs, but also indirect and hidden costs. Let's take a look at what the costs are for our model company in the case of manual invoice data extraction so we can measure the efficiency of invoice data capture.

- **Direct costs** are the fully-loaded costs of the employees of the AP department, in terms of the FTE Load, which ultimately stems from keystroke count.
- **Indirect costs** stem from additional effort associated with the process when it goes awry. They are associated with problem-solving such as identifying and correcting data entry errors or other associated issues, such as duplicate payments. This is significant – 72.5% of invoices require some type of re-working (as estimated by the IOFM, Institute of Finance & Management) when handled manually.
- **Hidden** (intangible) costs talk about everything else besides staff effort that is lost. They include penalties for late payments, loss of bonus for early payments, cash flow issues, vendor issue escalations, vendor rotation due to poor communication, employee rotation due to an inefficient process, etc.

Manual invoice data capture

When it comes to assessing the costs of manual retyping of the data, it is easy to fall into the trap of taking only direct costs into account. But the structure of costs is more complex. Let's take a look at how it works at our model company:

Direct effort spent: 111 seconds per invoice

The manual process is certainly heavy on human effort, totaling **105 keystrokes per invoice**. That's simply 75 fields times 7 keystrokes per field (6 characters on average, plus one keystroke to confirm and move to the next field). At 78 kpm, an invoice is completed in **81 seconds** on average.

However, a manual process means that every invoice must be handled manually as a whole rather than presented automatically by a system – either in its physical form or opened and arranged on the screen if the document is digitized or scanned. This easily takes an additional 30 seconds per invoice in practice, making the total a neat 777 seconds. That makes the actual speed 32 invoices per hour. At 75% time efficiency, **FTE Load is 3,840 invoices a month**.

3,840 invoices a month per employee may sound like a staggering amount in a manual process, but achievable in extremely efficient operations. It just goes to show how great a case for manual entry we have built in our model company when other analysts report that the average can be as low as 7,000 invoices per month per FTE.

Sometimes, an error is corrected quickly – as soon as it is communicated, the document is looked up, then re-synchronized across systems, and so on. But occasionally, the additional time to fix an erroneous invoice easily reaches 30 minutes (communication with other departments, vendors, banks...). In extreme cases, it can take hours or even days, when more employees and external contractors are concerned.

Out of the 30,000 invoices a month received by our model company, The Corp, Inc., that means 3,750 invoices reworked over 5.55 minutes each. That's **an extra 2.9 FTEs dedicated to reworking!**

TCO of manual invoice data extraction: Average cost per invoice is \$2.03


We are close to determining the TCO of manual invoice data extraction. In terms of direct FTE Load, processing 30,000 invoices requires 7.8 FTEs. Indirect load represents an additional 2.9 FTEs. That's 70.7 FTE Loads in total, with a likely team structure setup being 8 junior FTEs and 2.7 manager FTEs (including the Head of AP, involved particularly in corrections). The costs for managers are 2.7 x \$96,000, for typists 8 x \$48,000, totaling **\$643,200** per year for both direct and indirect costs combined.

The hidden cost is the trickiest to estimate, its structure varies market by market and also in terms of invoice sizes. Since it is influenced by the process' speed and accuracy, we think about it in relation to the indirect cost, and lean towards the hidden cost which is roughly half of the indirect cost portion (27.2%) on average. That would make it 73.6% of the total effort cost, or an additional **\$87,475**.

Total cost: \$730,675 per year for 360,000 invoices. The resulting cost per invoice is \$2.03. It is important to add that the efficiency of manual setup decreases in time because of the rising costs of human labor and the complexities of the invoices and validation processes.

As we explained, our fictional company extracts only a limited amount of data, furthermore only numeric! If a company retypes more data fields, including tables and descriptions of goods, for example, the price will rise accordingly, and very dramatically (the number of data fields may move by an order of magnitude).

This number can reach, according to Sterling Commerce, a staggering \$72.30 for processing, depending on the complexity of the process. Some resources put the price per invoice in manual processing as low as \$7 but they may follow a different TCO methodology.



We have calculated the TCO of manual invoice data extraction for our model company which is setting the bar for invoice data capture. We have discovered that this manual method is not only quite expensive but also presents a big strain on the morale of the typists – nobody likes to spend time correcting errors stemming from one erroneously typed number.

Part 3: The real cost of invoice automation

Template-based, on-prem OCR solution

During the 1980s, the rise of computers in the enterprise sector marked a big change in data capture. Optical Character Recognition technology, converting printed text into data, enabled extracting invoice data using templates automatically. This template is either a **zonal OCR** with fixed page locations for individual data fields, such as total amount or supplier address, or a **rule-based OCR** configured with a sequence of if-then rules telling the software where to look for specific information. This software is installed on company premises, set up, and continuously updated by supplier technicians and used by the internal AP department.

Common misconception

The most usual mistake when it comes to [template-based OCR systems](#)? That the cost for processing an invoice with a template system would equal the price of the document as determined by the license.

The actual cost structure

On-premise software carries several costs that are often neglected, from the one-time setup to regular maintenance. The initial implementation and setup of the system are typically quite time-consuming and

expensive — and are performed by supplier experts. Trained staff to perform the actual invoice data extraction. One of the related costs is the training, which can take just half a day, but retraining is needed relatively often due to personnel rotation. We need to take into account all those costs so we can calculate the TCO of template-based invoice data capture.

The system configuration is specific for every vendor and needs regular updates, due to natural vendor rotation and changes in invoice layouts. Our model company, The Corp, Inc. has 2,000 unique suppliers and the invoice traffic follows the typical structure of 70% of volume generated by 30% of suppliers. This translates to 600 layouts configured to cover 70% of incoming invoice volume.

Since setting up a template is expensive, it is never done for all suppliers. Even though 600 configured layouts might be on the high end for a template-based system, we will roll with this figure to stay conservative. From our clients, we hear template coverage rate reports of anywhere between 20% and 80% of volume.

Because template-based systems can process only invoices from vendors that are input into the system, the remaining 30% needs to be processed manually, the same way we described for manual data entry (except the 30-second invoice open time - we will correct this figure to 5 seconds in an automated workflow).

The invoices that fit the template also need manual verification and possible correction. This aspect of cost is very often forgotten, even though it will account for quite a significant piece of the cost, as we will see later.

Effort spent: 3.1 direct FTE Load, 1.1 indirect FTE Load

The 30% of invoices need to be processed manually, taking 87 + 5 seconds per document as analyzed above. For the 30,000 invoices total of our model company, the manual workload is **2 direct FTE Load**. Rework brings an additional **0.9 indirect FTE Load**.

For the 70% of invoices where a template kicked in, validation needs to be performed for each invoice. In the Rossum validation interface, purely the validation part in the model company would take 70.5 seconds; we will generously allow the same speed for validation in legacy template systems.

In approximately 70% of fields (optimistically), validation would spot a mistake and manual correction would be required – taking into account the KPMs as well as observed ergonomics of these systems, this would take another 5 seconds per field on average, or 7.5 seconds per document of correction time on average.

Multiple field errors often go together in a single document (e.g. the scan was faulty, automatic validation failed, etc.). That helps the speed of validation as the process is smoother for the user when it is not interrupted by correction efforts.

Cost per invoice: \$1.03

Before we consider the effort costs, we will take a look at the costs associated with the OCR system itself. We need to stress at this point that those costs can vary by hundreds of percent depending on the vendor. We are estimating them from our experience, our client's costs, and industry knowledge. For the purposes

of this comparison, we have used the average price of the most used template-based OCR invoice extractors.

One-time investments are the **original implementation and hardware/software purchase** (\$700,000 and \$75,000, respectively) – remember we are amortizing this cost over three years, translating to a cost of **\$38,333 per year**. The implementation effort is 50 to 200 man-days and apart from the process integration, system configuration, and setup, template configuration is also necessary.

Every year, a **license** needs to be purchased with fees for additional pages (\$20,000), and updates to vendor templates and system setup that can be expressed as a percentage of the original setup (\$ 25,000), totaling **\$45,000/year**.

The effort of operating the process translates to 4.2 FTE Loads. Assuming 1 manager FTE and 3.2 junior FTEs, this represents the direct and indirect cost of **\$249,600 per year**. Following the same principle as for manual process analysis, the hidden cost will represent a further 75% premium (**\$37,400**).

Hidden costs will have a different structure as the nature of errors will vary compared to a manual process. More care has been given to the process integration, catching some errors and speeding up the process. However, new hidden costs will be associated with managing the external consultants (or internal IT) that maintain the OCR software.

The total cost of this setup is \$370,333 per year, putting the price per invoice at \$1.03. Other sources, such as analyst group PayStream Advisors, estimate the cost less conservatively – as much as \$4 per invoice. It is important to add that the efficiency of this setup decreases in time, raising the average price.

Cognitive, cloud-based solution

Software as a Service (SaaS) is taking over the world, presenting a lightweight alternative to on-premise systems in all kinds of industries.

The other driver of change is machine learning. A cognitive [cloud-based OCR solution powered by AI](#) has the potential for an even higher rate of automation than a template-based system but does not need the costly and time-consuming implementation and updates. In fact, a prototype process can be set up in days.

The secret? Artificial Intelligence can be trained, like humans, to read the data in a non-linear, non-standard process, therefore not needing any pre-built templates.

The user training for a template system can take up to half a day, but the cloud-based solution features an intuitive interface, reducing the time needed for setup. The whole training for brand new users can be done in fifteen minutes, improving employee satisfaction.

We often encounter the situation when a new, AI-powered AP department needs only a few personnel: a head of AP, a manager, and a typist. The time previously put into invoice extraction is usually invested towards long-term projects and improvements.

The actual cost structure

Just like any other OCR system, a cognitive data capture solution will come with some implementation and license costs. This is alongside the operation effort costs and hidden costs that we revisit below.

Implementing the AI solution is a simple matter. As with template-based OCR, the essential components are the process design and integration with other company systems. But the OCR configuration is a matter of hours and its integration into a simple workflow with robots can be accomplished with a few man-days of effort in most ERP systems.

The key difference when comparing to template-based OCR is that the configuration does not involve layout by layout setup. The solution comes with already high accuracy all over the map and further adapts to the processed invoices based on a user. And it even improves by itself, just by the virtue of the OCR user base widening over time – that means the more diverse layouts seen, the better out-of-the-box performance.

Not only does this mean massive savings on implementation, but there is no longer a sharp 70-30 distinction between manual and automatic processing of documents.

All documents are processed automatically in a cognitive data capture solution, the only variation is in the degree of accuracy.

Full [automation of document processing](#) may be attainable on a case-by-case basis. But just to be on the safe side, let's consider that all invoices still need to be manually validated. Sometimes, manual corrections are going to be required. Over time, we do see the error rate drop well below the 70% level assumed for template-based systems – this is due to all the automatic learning facilities.

For the case of our model company, we would aim at 95% field accuracy in Rossum, with most of the corrections involving just a few keystrokes to correct a typo (we observe 7.43 correction keystrokes per document, and 95% accuracy means an average of 0.75 fields wrong per document).

Effort spent: 1.3 direct FTE Load, 0.3 indirect FTE Load

We can finally calculate the TCO of cognitive invoice data extraction. First, all invoices are processed automatically, validated, then corrected. The validation takes 70.5 seconds – remember, the estimate for template-based OCR was actually based on Rossum's measurements. We will assume the correction time per field to be 4 seconds with very few keystrokes, but the fixed overhead of 3.8 seconds remaining.

With 0.75 corrected fields per invoice and 5 seconds fixed per-invoice overhead (just like with template OCR), that makes **the total direct effort 18.5 seconds per invoice**. For the volume of our model company, this means **1.3 direct FTE Load**. We again consider reworking that process to take the same effort as in the manual data entry example, but apply only to a portion of the 70% of manually corrected documents again, causing just a **0.3 indirect FTE Load**.

Cost per invoice: \$0.45

The license's cost varies according to several factors. Most importantly, the volume processed and fields captured. For our model company, it is going to be on the cheap side due to the small set of fields, all of the standard AP fields, and no special enterprise requirements. We are going to use around **\$40,000 per year**, even though the actual price may well be south of this number.

To make room for a complex solution analysis at the beginning and for iterative process improvements after the initial deployment, we are budgeting **\$10,000 for the implementation** – but you can certainly save on that significantly too. Annually, let's budget 25% again (\$2,500) for system maintenance, e.g. reconfiguration due to changing business requirements. The user training costs are negligible compared to the template solution (the training takes 75 minutes).

All in all, this means the cost of **\$45,833 annually** for the solution itself as we are amortizing the implementation over three years again.

The effort of operating the process translates to 7.6 FTE Loads. Assuming 0.6 manager FTE and 7 junior FTEs, this represents the direct and indirect cost of **\$105,600 per year**. The hidden cost will again represent half of the indirect cost, thus now only a 9% premium (**\$9,500**).

Finally, putting this all together, we now have the TCO of cognitive invoice data extraction: a **\$160,833 total cost per year or \$0.45 per invoice**. And that's while still being rather pessimistic about the benefits of cognitive data capture – and optimistic about the other methods.

And what will the future bring for users of cognitive data capture systems? While the cost for fixed processes is going to increase over time, the total cost of a cloud-based system will continue to decrease. New technology improvements become available to all users, and the automatic AI learning means a continuous further reduction in effort required. This means that in the long term, automation of the AP process is only going to increase further for AI-based OCR users.

AP Department Structure*		
MANUAL	TEMPLATE-BASED OCR	AI-POWERED OCR
11	5	2
<ul style="list-style-type: none"> 1 Head of AP 8 Typists 1.7 Managers 	<ul style="list-style-type: none"> 1 Head of AP 3.2 Typists 	<ul style="list-style-type: none"> 0.6 Head of AP 1 Typist

**We are rounding the number of employees up. The cost for different methods of data capture stem mainly from the wages.*

What will tomorrow bring?

Invoices need to be processed in every organization, the faster, cheaper, and more accurately, the better. We have calculated the costs of invoice data capture for three frequently-used solutions. If your process is primarily manual (your AP retypes the data themselves), you will pay \$2.03 per invoice, with template-based, on-prem \$7.03 per invoice, and using a raising alternative, AI-powered cloud solution, your costs are going to be \$0.45 per invoice. We can definitely agree with content service provider Hyland that AI is revolutionizing data capture.

Today, moving from a fully manual process to a cognitive data capture service for our model company means a staggering 4.5x reduction of costs for an AP process.

Those numbers illustrate the current status of the fundamental differences between those processes. Other powerful tools to increase AP process efficiency certainly exist – a comprehensive Purchase Order system, robots applying automated verification steps to catch manual entry errors, or even a procure-to-pay EDI platform for frequent suppliers. But of course, all these tools compound with a cognitive data capture service just as nicely.

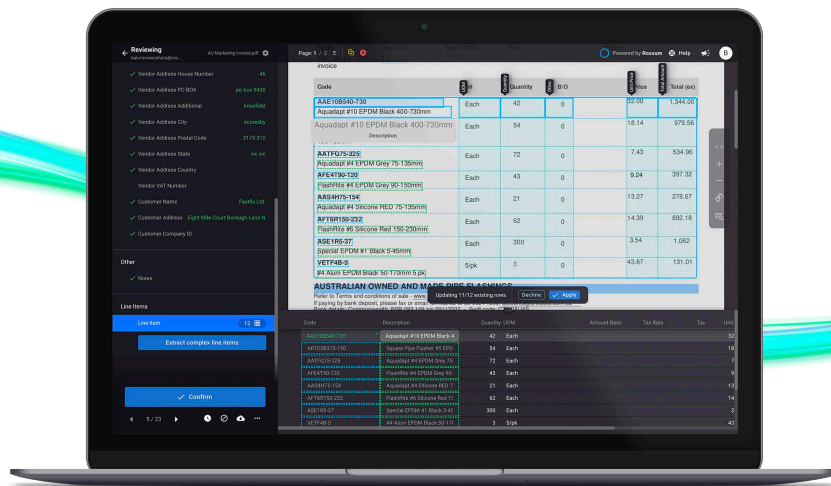
Outsourcing the AP data entry process completely may sound like another nice alternative, and it can indeed do wonders to your bottom line. However, it is a project that is fraught with risk – you are giving up control, and e.g. quality issues may take a long time to iron out during implementation.

What will tomorrow bring? The gap between the manual, OCR, and the cognitive system will only grow further with the increase in the volume of the invoices, or any extra complexity in the process (such as line items, vendor address capture, or purchase orders). Even today, staff willing to retype data are increasingly difficult to find. An automated solution won't just help your bottom line; your employees will be free to accomplish more fulfilling, less menial tasks.

The Rossum Platform

A single platform designed to minimize the time your organization spends processing documents.

[Book a demo](#)



Follow us:

