

IBM Automation Document Processing

Founded 1911 | HQ Armonk, NY | 380,000 employees | \$77.1B revenue (2019)

IBM Automation Document Processing is well thought out, from packaging it as a cloud service module (Pak) to the functionality's breadth and specificity. It would have been easy to rest on their Datacap success and add incremental improvements. Instead, IBM has reimaged, reframed, and rebuilt their capture experience.



The Company

IBM is one of the world's largest IT vendors, with \$71.1 billion in revenue (2019) and 380,000 employees. Headquartered in Armonk, NY, the company has operations in 177 countries. IBM has a long history in document capture – one that was bolstered by the 2010 acquisition of Datacap. In fact, for the past decade Datacap has been the cornerstone of the firm's capture business. However, 2020 has seen a major shift in direction with the release of IBM Automation Document Processing, the focus of this report.



The Technology

IBM Automation Document Processing (ADP) is one of the firm's newest capabilities within the Cloud Pak for Business Automation, released into general availability in December 2020. Cloud Paks are containerized software solutions, so in simple terms, they are pre-built, integrated solutions that run on top of the Red Hat open hybrid cloud platform. Or, to put it yet another way, they are essentially out-of-the-box applications that can run on any cloud. This, in and of itself, is a significant shift for IBM in the document capture market.

Traditional document capture is typically (but not always) run on-premises and requires a lot of configuration and management; it also typically requires running multiple modules to achieve a single goal. ADP leverages artificial intelligence (AI) including IBM Watson, removing the need for various modules, set-up, and configuration. ADP provides the ability to

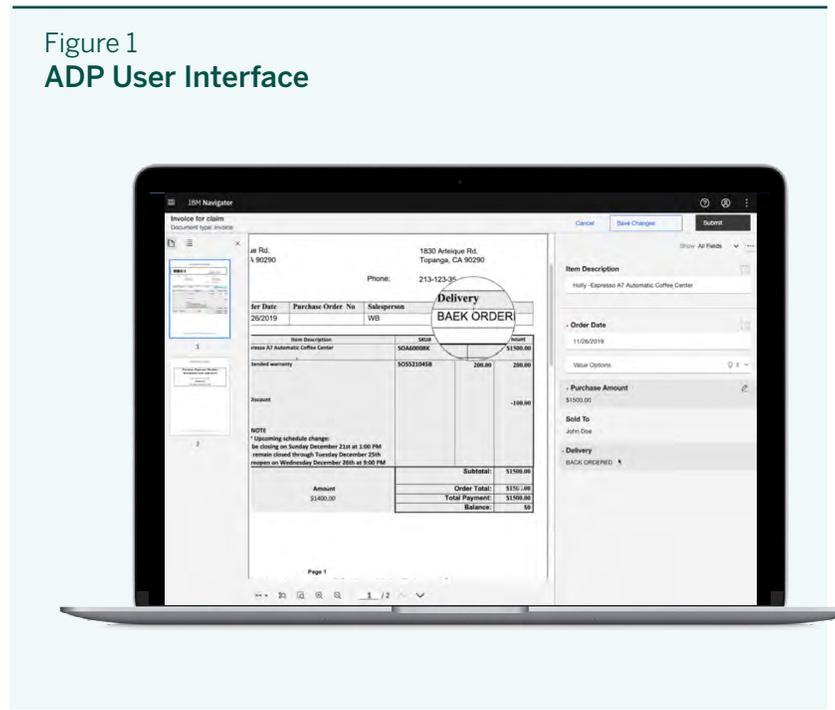
extract data from unstructured documents, add to, refine, and correct the extracted data, and then trigger integrated tasks and processes with that data within business applications and workflows.

ADP is a pretty sophisticated and potent product at its heart, but its complexity is largely hidden from the end user. For example, the deep learning capabilities are pre-configured to meet general document classification and data extraction needs and can be accessed, trained, and managed in a no-code user interface. Furthermore, though ADP does leverage complex neural networks (deep learning) techniques, it also includes HITL (Human in the Loop) functions to ensure that you remain in charge of accuracy, including the ability to set and refine confidence levels in the capture process. This no-code, human-centric approach is further exposed by providing out-of-the-box capture templates and pre-trained algorithms. Though more are expected, at launch the product comes with out-of-the-box templates (algorithms) for:

- Invoices
- Bills of lading
- Utility bills

To be clear, though ADP goes a long way to providing a full-featured document processing and capture product, it will require you to carefully design and implement it to meet your specific needs, as will any other similar system. For example, you must set up and define your particular operating parameters and requirements before going into production. Even so, ADP is one of the easier-to-use systems we have seen, primarily as so much has been pre-configured in advance, and its reliance on AI negates the need for many traditional pre-production activities such as setting and adjusting zonal templates of scripting overlapping locators and field identifiers.

Figure 1
ADP User Interface



We mentioned that ADP leverages deep learning, a complex and powerful AI technique based on neural networks. Though it is not the only firm that uses deep learning for general image and document capture, few use it for complex enterprise capture needs. Deep learning is a major step up in the AI world from machine learning (the most commonly used system) and brings some challenges and benefits. The biggest challenge is that deep learning is a black box environment by its nature, meaning that it is inherently unexplainable; you don't know how it came to its decision. Hence the importance of the HITL features in ADP, where you can at least semi-supervise its activities and outcomes. HITL is particularly important as, by default, the AI will automatically correct and enrich data, improving its structure, for example by correcting text characters that were recognized incorrectly through the OCR process.

The use of deep learning adds a lot of rich functionality that would traditionally have been only available as add-on modules, if at all. For example, ADP provides not only the regular text field extraction you would expect, but also advanced features such as tables and variable form extraction. It can also construct logical “objects” that can be grouped together. And, custom extractors can be built or applied from IBM Watson Studio. Finally, it’s worth noting that deep learning, if used correctly, can dramatically reduce the number of samples and the time required to teach the system what to capture.

All in all, then, what ADP provides is an AI-led, pre-configured, enterprise-grade, cloud software solution. But, arguably just as (if not more) importantly, the capture activities in ADP provide a foundation for an end-to-end document processing. Ultimately, processing documents through their lifecycle is the goal of any document capture operation, but in reality, most document capture products stop at the initial capture phase. ADP, though, has been built to feed into broader lifecycle document activities, particularly when used in conjunction with IBM FileNet Content Manager and Enterprise Records Manager through the available integrations. Indeed, these content repositories greatly benefit from the increased visibility of the information contained within the documents processed and subsequently stored on them, making it possible to automate enterprise search, auditability, security, and governance. Similarly, it is possible to integrate ADP with existing workflows and process management systems and to feed robotic process automation (RPA) systems automatically. There are integrations to IBM content management products for ADP and there is a REST API that allows calling services directly and making calls to external applications.

Our Opinion

On the one hand, it is not surprising that IBM, a company already synonymous with AI, should exploit its expertise by applying it to document processing. What is surprising is the sheer ambition of the ADP launch. Amazon, Google, and Microsoft have been making strides in leveraging deep learning and computer vision for document capture. Still, IBM has a deep and long history in the space and an understanding of enterprise customer needs that few can match. ADP is well thought out, from packaging it as a cloud service module (Pak) to the functionality’s breadth and specificity. It would have been easy to rest on their Datacap success and add incremental improvements. Instead, IBM has reimagined, reframed, and rebuilt their capture experience.

Advice to Buyers

Today, so close to the launch of ADP, this product will be of most interest to existing IBM customers and partners. It’s not an upgrade on their past document capture products; it is something new and will require careful consideration and planning to use effectively. However, we like it a lot and recommend that you look closely at it if your organization has significant enterprise document processing needs. Used effectively, it could simplify and streamline your end-to-end document lifecycle management activities.

SOAR Analysis

Strengths

- Deep learning and Watson AI technologies
- Massive brand recognition

Aspirations

- Transform enterprise document capture
- Deliver end-to-end document processing

Opportunities

- Modernize and expand the IBM document capture business
- Challenge and replace non-IBM document capture installations

Results

- General release in December 2020
- Integration with IBM content management platforms

About Deep Analysis

Deep Analysis is an advisory firm that helps organizations understand and address the challenges of innovative and disruptive technologies in the enterprise software marketplace.

Its work is built on decades of experience in advising and consulting to global technology firms large and small, from SAP, Oracle, and HP to countless start-ups.

Led by Alan Pelz-Sharpe, the firm focuses on Information Management and the business application of Cloud, Artificial Intelligence, and Blockchain. Deep Analysis recently published the book "Practical Artificial Intelligence: An Enterprise Playbook," co-authored by Alan and Kashyap Kompella, outlining strategies for organizations to avoid pitfalls and successfully deploy AI.

Deep Analysis works with technology vendors to improve their understanding and provide actionable guidance on current and future market opportunities.

Yet, unlike traditional analyst firms, Deep Analysis takes a buyer-centric approach to its research and understands real-world buyer and market needs versus the "echo chamber" of the technology industry.

Contact us:

info@deep-analysis.net

+1 978 877 7915