

AI OCR for Salesforce

Read scanned, faxed, photographed and native documents inside Salesforce, then run the process that follows

What is AI OCR for Salesforce?

CloudFiles is Document Automation for Salesforce, and CloudFiles Document AI brings AI OCR into the Salesforce org where your documents already arrive. OCR is the reading step. It turns a scanned page, a fax, a phone photo or a PDF into text a process can use. Document AI pairs OCR with large language models, so reading is followed by understanding: what the document is, which values matter, which Salesforce records they belong to, and what should happen next.

That answers the question buyers ask most often: is it an OCR solution or an AI solution? Document AI uses large language models and OCR. OCR reads the page. The AI works out what the page means and drives the rest of the document process in Salesforce Flow.

Which documents can AI OCR read?

- **Scanned documents.** Paper forms, signed applications, audit reports and vendor packets scanned to PDF.
- **Faxed documents.** Faxes that land in Salesforce as PDF attachments.
- **Photographed documents.** JPEG and PNG images, including photos taken on a phone.
- **Mixed PDFs.** Files where some pages are native text and some are scans.
- **Native PDFs.** Digital PDFs, where the text is read directly.
- **Other files.** CSV, TXT and JSON files through Process Document using AI, and Office documents through Doc AI Workflows.

Document AI sends scanned and mixed pages to OCR and reads native PDF text directly, so each page is read in the way that suits it. The Pages to process setting limits a run to the pages you choose, which is useful when the fields you need sit on the cover page or the summary page of a long file.

How does AI powered OCR extract data from uploaded documents?

AI powered OCR to extract data from uploaded documents works in two parts. First the page is read. Then Document AI takes the fields you described in plain language and extracts them from what was read.

You describe each field the way you would explain it to a colleague: "the invoice number", "the total amount due including tax", "each line item with its description, quantity and amount". There are no layout templates to draw and no zones to maintain. The same description works on the next vendor's layout, because it describes what the value means rather than where it sits.

Three Flow actions put the reading step to work:

- **CloudFiles: Process Document using AI** reads a PDF, JPEG, PNG, CSV, TXT or JSON file and returns its content to the Flow.
- **CloudFiles: Document Mapper** reads a document and creates or updates a Salesforce record plus child records from a saved mapper configuration. It reads JPEG and PNG images as well as PDFs.
- **CloudFiles: Run Doc AI Workflow** runs a workflow whose Extract fields step returns typed fields and lists of items. Workflows take PDFs and Office documents.

Why does template OCR break when plain-language fields do not?

Template OCR reads a value from a position on the page. When a vendor moves the invoice number, adds a column or changes a form, the value is no longer where the template expects it, and someone has to rebuild the template or go back to manual data entry. Every new layout is a new template to build and maintain.

Document AI has no layout templates. Fields are described in plain language, and large language models read what each value means on the page, so a moved field needs no new template. A new vendor, a revised form or a document your team has never seen before is read with the same field descriptions. That is the difference between OCR tools that capture text and AI OCR that runs a document process.

How do you automate document scanning and OCR for file attachments?

To automate document scanning and OCR for file attachments, start Document AI the moment a file arrives. Every trigger is a standard Salesforce entry point:

- **Record-triggered Flows**, on an asynchronous path.
- **CloudFiles events**, such as a Salesforce file attached or a file uploaded.
- **Email-to-Case attachments**, a common entry point for scans and faxes that arrive by email.
- **Screen Flows**, so a user uploads a document and processes it in one step.
- **Mapper buttons** on List Views and Detail Pages.
- **Connected folders**. With CloudFiles Document Management, a file dropped into a connected SharePoint, OneDrive, Google Drive, AWS S3, Azure Blob, Box or Dropbox folder can trigger processing.

Results and completion come back as CloudFiles events, such as document-processed, document-mapper-run and error-event, which your Flow reads with CloudFiles: Get Event Details and uses to decide the next step.

Can I use mobile document scanning and OCR in Salesforce apps?

Yes. A field user can photograph a document on a phone and attach the image to a record in the Salesforce mobile app. JPEG and PNG images work with Process Document using AI and Document Mapper, so a Flow can read the photo and write the values back to the record. For mobile document scanning and OCR for Salesforce apps, the phone's camera captures the page and Document AI does the reading, the matching and the writing inside the org.

Which OCR tools extract text from documents to objects?

OCR tools to extract text from documents to objects have to do more than read. They have to put each value on the right record. Document Mapper does this with a saved configuration that creates or updates a record plus its child records, for example an invoice and its line items, or an application and its services.

- Writes go through standard Salesforce DML, so validation rules and field-level security apply.
- Lookups are resolved per field by Get Existing, Create New or Call Flow Action.
- A per-field option protects values that are already on the record.
- Document sets let one run read several files.
- When automation runs finish, the result is either auto approved or held in the review queue. Held runs show Pending Review, and a reviewer can edit cells, add rows, then choose "Approve & Write to Salesforce" or Reject.
- A lookup with no match, or several matches, forces review.

What happens after the page is read?

OCR is where the process starts. After reading, Document AI can classify the document with a Classify step, a Query prompt or a Document Type field, route it with a Decision step, split a multi-document PDF bundle into one Salesforce File per document, match values against existing records, write records and child records, and hold runs for review. Document Generation can create documents from the extracted data in the same process, and Document Management's Move Resource action can file classified documents into the right folder.

Each run is recorded in Salesforce as a CloudFiles DocAI Run record, with a status of Processing, Completed, Pending Review, Failed or Rejected. Reprocessing creates a new run and keeps the old one.

What should the best OCR apps for automated document and form processing do?

What to test	CloudFiles Document AI
Reads scans, faxes and native PDFs	Yes: scanned and mixed pages go to OCR, native text is read directly
Reads phone photos	Yes: JPEG and PNG through Process Document using AI and Document Mapper
Works without layout templates	Yes: fields are described in plain language
Runs inside Salesforce Flow	Yes: every capability is a Flow action, also callable from Apex
Writes records and child records	Yes: through standard DML, with validation rules and field-level security
Holds results for review	Yes: review queue in Mapper, review policy in workflows
Keeps a record of every run	Yes: a CloudFiles DocAI Run record per run

On AI data handling: CloudFiles does not train or fine-tune AI models on your documents.

Get started

Install CloudFiles Document AI from AgentExchange and start a free 14-day trial or a Test Drive, or book a demo at cloudfiles.io. A free sandbox proof of concept on your own scans and forms is also offered.