

Document AI for Financial Services: Split PDF Case Packets and Classify Each Document in Salesforce

Intelligent document processing that splits multi-document packets, classifies and renames each document, and attaches it to the right record

CloudFiles Document AI is part of CloudFiles, Document Automation for Salesforce. This guide describes a pattern we see at financial services firms that run cases in Salesforce: lending, account servicing, claims and client onboarding teams that receive their documents as one combined packet per case. Document AI runs AI document processing on each packet before case work starts.

What is in a financial services case packet?

A case packet is several documents saved as one PDF. A single packet might hold an application form, account statements, a signed agreement, proof of address, tax forms, correspondence and a supporting letter, in whatever order the client or the branch scanned them. Some packets arrive by email, some through a client portal, some from a scanner in a branch office.

Before anyone can work the case, someone has to open the packet, find where each document starts and ends, cut it into separate files, work out what each one is, rename it so the next person can find it, and attach it to the right Case, Account or Opportunity. Many teams ask branches to put a separator sheet before every document so a tool can find the breaks. The separator sheets get forgotten, and the packet goes back to manual splitting.

How does Document AI split the packet?

CloudFiles: Split Document using AI takes the packet and splits it into one file per document. It reads the content to find where each document begins, so no separator sheets are needed. It names each output file from its content, for example the document type, the client name and the date, and it saves the pieces as Salesforce Files.

This is split PDF work done by reading the documents rather than counting pages. A packet with an extra statement, a missing form or pages in an unusual order is split by what it contains. When the split finishes, a document-split event comes back, read with CloudFiles: Get Event Details, and the Flow continues with each piece.

How is each document classified?

Each split file is classified against the firm's own document types. There are three ways to do it, and a firm picks the one that fits the process:

- **A Classify step in a Doc AI Workflow**, when the classification drives a sequence of further steps.
- **A CloudFiles: Query Document/Dataset prompt in Flow**, which asks the document what it is and returns the answer to the Flow.
- **A Document Type field in Document Mapper**, when the same run also writes fields from the document.

The document types are described in plain language: "a bank or brokerage account statement", "a signed loan or account agreement", "a government tax form". Document classification by description means a new statement design from a bank the firm has not seen before is still recognized as a statement.

A Decision step then routes each document by type.

How is each document attached to the right record?

Routing by type decides what happens next. A statement can be attached to the Account and have its closing balance and statement date extracted. An agreement can be attached to the Opportunity and its signature date checked. A tax form can be attached to the client's Contact record. Workflow find and search steps locate the right record from the client name, account number or case number in the document, by field match, fuzzy match or a match picked by AI.

With CloudFiles Document Management, classified documents can also be moved into the right folder in SharePoint, OneDrive, Google Drive, AWS S3, Azure Blob, Box or Dropbox with the Move Resource action, so the firm's document repository stays organized by client and document type.

What gets extracted after classification?

Once each document is classified, the process can extract data from it by type. A Doc AI Workflow's Extract fields step returns typed fields and lists of items, such as the transactions on a statement. Document Mapper creates or updates a Salesforce record plus child records in one run. Writes go through standard Salesforce DML, so validation rules and field-level security apply, and a blank value never wipes a field in a workflow.

Can a person review before anything is written?

Yes. Workflow review policies decide what is written without a person's sign-off: "Straight through", a hold on the first write, or a policy that stops at every write. A single step can be set to "Require approval for this write", for example the step that updates an Account from a statement, while attaching files runs on its own. Paused runs show "Awaiting review". In Document Mapper, runs can be auto-approved or held in the review queue as Pending Review, where a reviewer edits cells and chooses "Approve & Write to Salesforce" or Reject. A lookup with no match or several matches forces review.

What is recorded for each packet?

Each run is recorded in Salesforce as a CloudFiles DocAI Run record with its status: Processing, Completed, Pending Review, Failed or Rejected. Workflows keep a step-by-step history of every run, and every pause, approval and rejection is recorded. Reprocessing creates a new run and keeps the old one.

Question from the case team	Where the answer lives
Which documents were in this packet?	The split files, named from content, on the record
What was each document classified as?	The workflow's step-by-step history
Who approved the write?	The workflow's recorded approvals
Did processing finish?	The CloudFiles DocAI Run record status

How are client documents handled?

CloudFiles does not train or fine-tune AI models on your documents. Document AI uses large language models and OCR. Your files are stored in the data-residency region you choose: United States, European Union, United Kingdom or Australia. Security documents are available under NDA from the CloudFiles Trust Center.

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 is offered, so the case team can split and classify its own packets first.