

Document AI for Accounts Payable: Matching and Filing Vendor Packets in Salesforce

OCR, matching and filing for scanned vendor packets, with an Unmatched folder for anything that needs a person

CloudFiles Document AI is part of CloudFiles, Document Automation for Salesforce. This guide describes a pattern Document AI supports for accounts payable teams that pay vendors by check run and keep vendor records in Salesforce: scanned vendor packets are read, matched to the check run, renamed and filed without anyone sorting them by hand. Each stage is intelligent document processing inside Salesforce: read, extract, match, file.

What is a vendor packet, and why does it take so long?

After a check run, accounts payable is left with paper. Each vendor packet holds the invoice, a voucher or remittance advice, sometimes a copy of the check stub, and sometimes supporting documents such as a purchase order or a delivery receipt. The packets are scanned, often in one batch, and land as PDFs with names like the scanner's date stamp.

Someone then opens each scan, reads the voucher or invoice reference, looks it up in the check-run export, works out which vendor it belongs to, renames the file so it can be found later, and saves it in the vendor's folder. Anything that does not match gets set aside in a pile. It is invoice processing work that produces no new data, only order, and it still takes a person's full attention.

How does Document AI read a scanned vendor packet?

Scanned pages are sent to OCR, and any native PDF text is read directly. A Doc AI Workflow, started with CloudFiles: Run Doc AI Workflow, reads the packet and uses its Extract fields step to pull the fields the match depends on: the voucher number, the invoice number, the vendor name, the invoice date and the amount. The fields are described in plain language, so a packet from a vendor whose invoice layout changed last month is read the same way as one from a vendor who has used the same layout for years.

If one scan holds several vendors' packets, a split PDF step with CloudFiles: Split Document using AI first splits it into one file per packet, names each output file from its content and saves the pieces as Salesforce Files. Each piece then goes through the workflow on its own.

How is the voucher or invoice reference matched to the check run?

The check-run export is the source of truth for what was paid. In this pattern it is brought into Salesforce as records, and the workflow's find and search steps match the reference from the packet against it:

- **Field match** on the voucher or invoice number when the reference is printed clearly.
- **Fuzzy match** when a scan drops a character or a vendor formats the reference its own way.
- **A match picked by AI** when the workflow has to weigh the vendor name, date and amount together to find the right line.

The same export can also be combined with the packet files in a dataset built with CloudFiles: Create Dataset, so CloudFiles: Query Document/Dataset can answer questions across both, such as which paid vouchers have no packet on file.

How is the packet renamed and filed?

A Decision step follows the match.

- **Matched.** The packet is renamed from its content, for example the vendor, the voucher number and the payment date, and CloudFiles Document Management's Move Resource action files it into that vendor's folder.
- **Not matched.** The packet goes to an Unmatched folder instead, and a task or a list view shows the accounts payable team what is waiting.

Vendor folders can live in SharePoint, OneDrive, Google Drive, AWS S3, Azure Blob, Box or Dropbox through Document Management, or stay in Salesforce Files. A scan dropped into a connected intake folder can trigger the whole process, so the team's only job at the start is to scan.

Can the team review matches before files move or records change?

Yes. The workflow's review policy decides what happens without a person's sign-off. "Straight through" runs every step. A hold on the first write pauses the run before it changes any record, and a stricter policy stops at every write. A single step, such as the one that updates the vendor's payment record, can be set to "Require approval for this write", while filing runs on its own. Paused runs show "Awaiting review", and every pause, approval and rejection is recorded.

Writes go through standard Salesforce DML, so validation rules and field-level security apply. A blank value never wipes a field.

Can the same process write the invoice to Salesforce?

Yes. Once a packet is matched, CloudFiles: Document Mapper can read the invoice in it and create or update an Invoice record with its line items as child records. Lookups resolve the vendor per field with Get Existing, Create New or Call Flow Action. Held runs show Pending Review, and a reviewer edits cells,

adds rows, and chooses "Approve & Write to Salesforce" or Reject. This turns the filing process into full invoice processing when the team wants it.

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. The workflow keeps a step-by-step history of every run, so the team can see which reference was read, which check-run line it matched and which folder the packet went to. Reprocessing creates a new run and keeps the old one.

Packet outcome	What the team sees
Matched and filed	Renamed file in the vendor's folder, completed run
Not matched	File in the Unmatched folder, run available to review
Held for approval	Run showing Awaiting review

Do vendor documents train the AI model?

No. CloudFiles does not train or fine-tune AI models on your documents. Document AI uses large language models and OCR.

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 accounts payable can run its own vendor packets through the process first.