

Accounts payable, on autopilot

A high-volume, purchase-order-driven AP operation processed **~289K invoices in six months** with two engines working together: an **AI that reads every field** on an invoice and keeps learning, and a **three-way PO matching layer** that gates risk so clean invoices flow straight through to the ERP — no human touch. This is what that looks like in production: the volume, the configuration, the funnel, the accuracy, the approval network, and the autopilot.

~289K

invoices processed in 6 months (peak ~80K in a single month)

97.7%

overall field-prediction accuracy (98.5% latest month, and rising)

87.9%

of invoices land on a clean PO match; the rest auto-route to a person

63.0%

of recent fresh invoices flow fully hands-off, end to end

01 The account

A large enterprise with high invoice volume and a vendor base that skews to rentals, parts, freight, and fuel — the kind of long-tail, PO-heavy spend that is hard to automate by template. Nearly every invoice is tied to a purchase order, so **matching quality is the whole game**.

Sector	Equipment & field services (heavy-asset, multi-branch)
AP model	~97% purchase-order based — three-way matched
Scale	~289K invoices / 6 months · ~80K in peak month
Spend profile	Long-tail vendor base; rentals, parts, freight, fuel

Live configuration

settings & gates

Autopilot

The models predict every field on the invoice — no manual data entry.

Three-way PO matching

Invoice matched against its purchase order and goods receipt before it can post.

\$5,000 approval line

Invoices under \$5k need no human approval; \$5k+ always route to an approver.

Auto-post


Approved invoices post to the ERP automatically, with no person in the loop.

Workday HRIS sync

Employees, roles, and org structure flow from Workday — approval routing stays current automatically.


02 Volume & throughput

Volume scaled fast — from a few thousand invoices in January to nearly 80K in March — and posting throughput chased it down. The gap between the bars is the backlog forming each month; it narrowed all spring and **crossed over in June**, when the system posted **88,143** against 57,318 received.



03 From received to posted — the lifecycle funnel

Of 289,156 invoices received, **crossed over in June**, within seconds, and **79% reach posted**. The one place time accumulates is the step from **matched** → **submitted** (invoices waiting on a PO fix, a receipt, or review) — once an invoice is submitted, approval and posting are effectively instant.



Matching outcome of every received invoice

78% clean • 16% needs a PO step • 6% mismatch



Clean = matched with no rule violation (225,355). Needs a PO step = a PO is present but not yet loaded / applied, or a true expense (47,141). Mismatch = a PO rule needs an approval (16,660).

04 AI acumen — it reads invoices accurately, and keeps learning

Across every automated invoice, a person overrides the model on only a small fraction of fields — and the fields that matter most for payment (the total paid and the vendor) are among the most accurate. Every field is corrected less than 2% of the time. The models learn from the corrections people make, so accuracy climbs month over month.

97.7%

overall accuracy, full window

95.4% → 98.5%


overall accuracy, Feb → May (rising)

3.0% → 0.2%

vendor corrections, Feb → May

How often a person corrects each field

Snapshot



$x = \text{field}$ · $y = \% \text{ of invoices where a person changed the model's value}$ · full window · lower = more accurate.

Correction rate by field, by month

Over time



$x = \text{prediction month}$ · $y = \text{correction rate} (\%)$ · lower = better · May* is partial · elevated Feb figures reflect the initial ramp.

Accuracy improves as the models learn the account

Since the February ramp, overall accuracy rose from 95.4% to 98.5% and the correction rate on nearly every field fell — vendor corrections dropped from 3.0% to 0.2%, cost-category corrections from 3.3% to 0.8%. This is a model that gets measurably better the longer it runs on a customer's data.

05 PO matching — the skill that gates risk and unlocks no-touch

Reading an invoice accurately is only half the job. The other half is **matching it to the right purchase order and goods receipt** — and, just as importantly, knowing when **not** to. This is what makes hands-off posting safe.

Three-way matching at scale


Nearly every invoice is matched against its purchase order and goods receipt before it can post. 87.9% land on a clean match.

The PO safety net

If the model picks the wrong PO — or none — the invoice can't match, so it's automatically held for a person and never auto-posts. The highest-impact error is structurally prevented.

The PO safety net

If the model picks the wrong PO — or none — the invoice can't match, so it's automatically held for a person and never auto-posts. The highest-impact error is structurally prevented.



06 Workforce & approvals — synced from Workday

Hands-off posting only works if the human-approval layer behind it is rock-solid. This account runs it on a **Workday HRIS integration**: employees, roles, and reporting lines sync from Workday into the platform, and **role-based approval flows** route every invoice that needs a person to the right approver — automatically. As people are hired, change roles, or leave, routing updates with **no manual approver maintenance**.



Why HRIS-driven routing matters at this scale

Maintaining approver lists by hand across 8 entities and 9,000+ people is not realistic. Because roles come straight from Workday, the **\$5k+ invoices always reach a valid, current approver** and the sub-\$5k auto-approve path stays clean — the approval layer scales with the organization instead of becoming the bottleneck.

07 Autopilot in production

Autopilot — the models processing an invoice end to end with no manual entry — scaled from near zero to roughly half of all fresh invoices, and the payoff shows up as speed: recent cohorts post far faster, and the system now executes essentially all of the posting itself.

2% → 47%

of fresh invoices auto-processed
(autopilot rollout)

~99%

of posting executed by the system,
not a person

~63%

of recent fresh invoices flow fully
hands-off, end to end


Autopilot rollout — share auto-processed

Over time



Fast-post rate — % posted within 1 / 3 / 7 days

Over time



Daily posting volume

Over time · invoices / day



08 The result — high automation, small and contained risk

Automation only counts if the money stays right. On the hands-off cohort, a person changes any field on just 3.54% of invoices, and the validated cash exposure is a rounding error against the dollars flowing through — and what little exposure exists is concentrated in a handful of vendors that a few approval rules contain.

9,194

employees & roles kept
in sync from Workday

1,522

role-based approval
flows routing invoices
automatically

~530

people actively serving
as approvers today

8

legal entities under one
synced workforce

The two engines compound

The AI reads every field at 97.7% accuracy and keeps improving; PO matching structurally holds anything that doesn't match, making hands-off posting safe; and the Workday-synced approval layer routes exceptions to the right person. Together they delivered 87.9% clean matching, ~63% hands-off flow, and 0.05% net cash exposure — across ~289K invoices in six months.