CASE STUDY

Logistics Operator

From Customer PDFs to target TMS Excel

How a Cloud28 connector turns manual container-level copy-and-paste into a 24/7 real-time order-processing workflow, with Codex accelerating the build, validation, and continuous improvement loop.

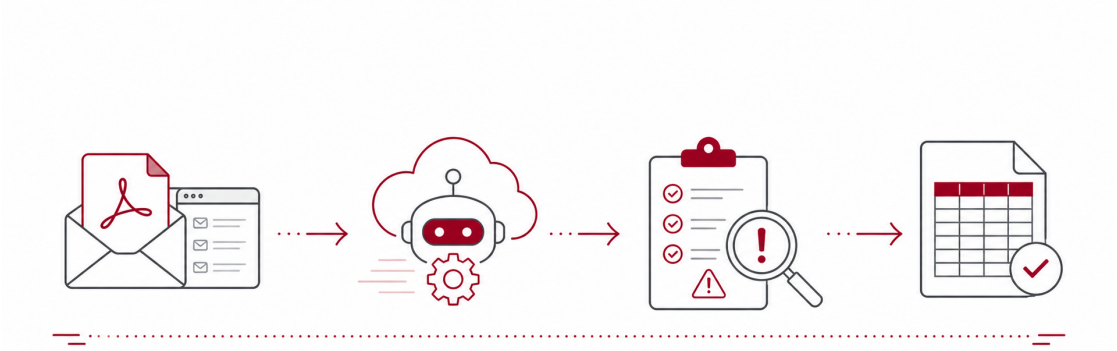

CLIENT	Logistics Operator
WORKFLOW	Customer PDF order intake to target TMS Excel import
SCALE	10+ orders per day; 1-60 containers per customer PDF
ENGAGEMENT	Cloud28 connector build with 24/7 shared mailbox monitoring, PDF extraction, port-code enrichment, Excel generation, folder upload, and exception routing
DEPLOYED	Five-day implementation path from kickoff through production test
OUTCOME	Manual transcription becomes real-time exception review; unknown customer PDFs are caught immediately and can be updated within hours on Cloud28

01 - OVERVIEW

Executive Summary

Logistics Operator receives new and change orders through customer PDF files from customers such as customer and customer These PDFs contain the practical operating instructions that move containers from ports to warehouses or rail yards, but the information does not arrive in a target TMS structure. Today the operations team must open each PDF, interpret the customer layout, identify the order fields, and manually copy every container into Excel.

The volume makes the workflow expensive even when each individual step looks simple. The team processes 10 or more orders per day, and each customer PDF can contain anywhere from one to 60 containers. Because each container copy-and-paste step takes approximately two to three minutes, a busy day can turn into hours of repetitive transcription before the order data even reaches target TMS

The Cloud28 connector turns this into a controlled exception-management process that runs continuously. The system monitors the shared mailbox 24/7, extracts order and container details from each PDF, enriches records with port-code lookup, stores the data in a digital twin database, generates the required Excel output, and uploads the completed file to the designated target TMS folder path. Staff stay involved where judgment matters, but the repetitive row-by-row movement of data is automated.

This changes the service level of the operation. The traditional process depends on people acting during business hours, typically 9:00 AM to 6:00 PM with rotational Saturday coverage. If a PDF arrives after hours, during a busy queue, or before the assigned user has time to convert it, the order waits. With the automation agent monitoring constantly, the order can move in real time as soon as the PDF lands in the mailbox.

The quality model also improves. Known customer PDFs can be processed automatically with a target error rate below 1%. Unknown or new customer PDF formats are not allowed to create silent downstream errors; they are caught instantly as exceptions, routed for review, and can be updated within hours using the Cloud28 platform. This makes the connector operationally resilient rather than brittle.

This use case is important because it creates both immediate operating relief and a long-term data foundation. The immediate gain is less manual work, real-time order handling, fewer copy/paste mistakes, and clearer exception routing. The longer-term gain is that Logistics Operator begins building a structured digital twin of order activity that can later support executive dashboards, auditability, analytics, and a phased modernization path beyond target TMS when the company is ready.

24/7
constant mailbox monitoring by the automation agent

<1%
target error rate for supported customer PDF formats

eal-time update
target TMS is updated within minutes of receiving ers

02 - CURRENT OPERATION

The manual workflow is container-by-container transcription.

The current operating flow depends on people reading every customer PDF and transforming unstructured or semi-structured order information into the exact Excel shape target TMS expects. It is a workflow that rewards care and repetition, but it also creates avoidable exposure: every copied field is a chance for the wrong container, customer, port, rail, warehouse, or delivery value to be entered.

The work is also interrupt-driven. New and change-order PDFs arrive through the shared mailbox, which means staff must monitor email, switch context, open attachments, and decide whether each order is routine or requires special handling. When a PDF includes many containers, the team repeats the same extraction and copy process row after row until the target TMS import file is ready.

The timing problem is just as important as the typing problem. A traditional operations schedule only covers business hours, typically 9:00 AM to 6:00 PM, with rotational Saturday shifts. Orders arriving outside that window, or during a period when staff are already busy, wait until someone notices the email and begins conversion. That delay is built into the human-driven workflow.

Port-code lookup adds another layer of manual effort. If the port mapping is known, staff can proceed; if it is unclear, they must search public sources or internal target TMS reference data, then decide which value belongs in the file. Without a reusable mapping library, the same lookup work can recur across future orders.

The business risk is not only time. Quality control often happens late, after the Excel file is prepared or after target TMS import issues surface. That means the team may spend effort creating a file and then spend additional effort correcting missing, inconsistent, or mis-entered information.

Step	Manual action today	Operational friction
1	Watch the shared mailbox for customer order PDFs during staffed hours.	New and change orders wait until a user notices and acts on them.
2	Open each PDF and identify customer, pickup, delivery, warehouse, rail, and container details.	Different customer layouts create interpretation work and missed-field risk.
3	Copy each container into an individual Excel row.	Each container takes about 2-3 minutes; a PDF can contain up to 60 containers.
4	Look up or confirm port codes.	Unclear mappings slow order processing and create inconsistent data.
5		Manual naming, saving, and placement can vary by user.
6	Correct import issues or missing fields after review.	Quality control happens late, after manual effort is already spent.

9-6 + Sat	After-hours delay	Manual queue
traditional staffed coverage with rotational Saturday shifts	orders wait when PDFs arrive outside staffed coverage	incoming PDFs sit until a user opens, reads, and converts them

03 - SOLUTION

Cloud28 turns the order inbox into a structured operating flow.

The proposed connector does not treat this as a generic PDF extraction exercise. It treats the email inbox, the PDF, the port-code mapping, the Excel template, the target TMS folder path, and the exception process as one operating flow. That matters because the value is created when the entire chain works together, not when a single PDF is parsed in isolation.

Cloud28 monitors the shared mailbox continuously and routes incoming PDFs into an extraction workflow as soon as they arrive. The extraction agent converts customer order content into structured fields at the order and container level. The port-code lookup agent then enriches those records using configured public sources or target TMS reference data, and confirmed mappings are stored for future reuse.

The result is real-time order handling instead of user-dependent conversion. A PDF no longer waits for the next available staff member to open the email, read the attachment, and copy rows into Excel. The automation agent begins the work immediately and only asks for help when the content is unknown, ambiguous, or outside the trained customer formats.

The digital twin database is the center of the use case. It stores extracted order records, container rows, port and destination fields, processing status, and mapping history. This gives Logistics Operator traceability for each output row and creates a foundation for future reporting dashboards, operational analytics, and longer-term TMS modernization.

The final operating step is the generation of the target TMS Excel file and upload to the agreed folder path. When confidence is high, the file moves through automatically. When confidence is low, a mapping is unresolved, or an unsupported layout appears, the system notifies the operations team and asks for review instead of forcing silent automation.

Capability	What it does	What changes for Logistics Operator
24/7 mailbox reader	Detects incoming customer PDFs from the shared inbox at all hours.	Orders enter the workflow immediately instead of waiting for staffed coverage.
PDF extraction agent	Parses order, container, pickup, delivery, and shipment fields.	Staff stop copying every row from the PDF manually.
Digital twin database	Stores order, container, port, destination, processing status, and mapping data.	The team gains traceability and the foundation for reporting.
Port-code lookup	Checks configured sources and stores confirmed mappings.	Repeated mappings become reusable instead of rediscovered.
Excel generator	Creates the agreed target TMS-ready Excel output, one row per container.	target TMS receives consistent files in the expected format.
Exception notifier	Emails the team when extraction confidence is low, lookup cannot be resolved, or a new customer PDF format appears.	Unknown formats are caught instantly and can be updated within hours on Cloud28.

Cloud28 Platform role in delivery

Cloud28 Platform acts as the full-stack service layer for this workflow, leveraging agentic engineering to build and maintain connector code, extraction tests, schema definitions, validation rules, Excel mappings, exception-notification templates, and repeatable QA checks. The result is a maintainable operating connector with a measurable improvement loop, rather than a one-off script that only works for a narrow demo.

04 - REALIZATION PLAN

A narrow, production-oriented slice proves the use case quickly.

The use case is designed to become real through a short, focused implementation path. The first goal is not to automate every possible customer PDF forever; it is to prove the workflow on representative PDFs, confirm the target TMS Excel output, validate the exception process, and give the operations team a connector they can trust.

Day 1 focuses on kickoff, sample review, and the first extraction connector. This is where Logistics Operator confirms the target Excel format, representative PDFs, and the data fields that must be present for target TMS processing. By the end of the day, the project should have a working conversion path for sample PDFs into structured Excel rows.

Days 2 and 3 turn the prototype into an operating flow. The digital twin database is built to store order, container, port, destination, processing status, and mapping records. The 24/7 shared mailbox connector, Excel generator, folder output step, and exception-notification routing are connected so the workflow resembles the production path rather than a disconnected demo.

Days 4 and 5 are about trust. UAT tests the connector across representative customer customer and other priority customer PDFs, including single-container and multi-container cases. The production test then confirms live or production-like processing, folder output, team escalation, and go-live readiness.

Day	Focus	Customer-readable deliverable	Acceptance signal
1	Kickoff and PDF-to-Excel connector build	Confirmed field map, sample PDF review, initial extraction connector	Sample PDFs convert into structured Excel rows.
2	Digital twin database build	Order, container, port, destination, status, and mapping schema	Extracted PDF data is stored with container-level traceability.
3	24/7 email connector and folder output	Constant mailbox monitoring, Excel generation, target TMS folder upload, exception routing	New order PDFs create Excel files in the agreed folder path in real time.
4	UAT and full test	Representative customer and customer PDF tests with tuned extractio	Logistics Operator confirms correct orders or clear
5	Production test	Live or production-like processing, operating notes, go-live readiness confirmation	Production test completes successfully and the team is ready to use it.

05 - EFFECTIVENESS OVER TIME

The proof is measured before rollout, during pilot, and after production.

The effectiveness model starts before automation. Logistics Operator should capture a one-week baseline using real orders: number of PDFs, containers per PDF, manual minutes spent, port lookup issues, target TMS import corrections, and rework events. That baseline becomes the control group for the pilot.

During the pilot, the connector should be measured against the same order types and customer PDF formats. The most important question is not whether every edge case disappears on day one; it is whether supported PDFs move from email to target TMS Excel in real time, with materially less manual effort and clearer exception handling.

After production launch, effectiveness should be reviewed as a trend. Exception rates should decline as the team confirms mappings and the connector learns recurring customer layouts. Manual copy/paste time should drop because staff only intervene when confidence is low, a lookup is unresolved, or a new format requires controlled enhancement. For supported customer PDFs, the target operating error rate should be less than 1%.

Unknown or new customer PDF layouts should be measured differently from normal processing errors. They should be treated as instantly detected exceptions, not silent failures. Once caught, Cloud28 gives the team a rapid update path so a new format can be mapped, tested, and moved into the supported automation set within hours.

The digital twin database makes the improvement visible over time. Instead of relying on anecdotal feedback, Logistics Operator can track source PDFs, generated rows, unresolved mappings, processing status, upload timestamps, reviewer actions, and rework. That operating evidence supports executive reporting and gives leadership a factual basis for deciding which follow-on automation or target TMS modernization phase should come next.

per day, an average of 15 containers per order, and 2-3 minutes of manual work per container, the current workflow consumes a meaningful block of employee time every month.

Metric	Baseline to capture	Target signal after rollout
Manual effort	2-3 minutes per container copied into Excel; measure one full week by order and container count.	70%-90% reduction in manual copy/paste time for supported PDF formats.
Cycle time	Time from email receipt to TMS-ready target TMS Excel file during 9-6 staffed operations.	Real-time generation from 24/7 agent monitoring unless an exception requires review.
Error and rework	Wrong container, customer, port, rail, warehouse, delivery, or missing-field corrections.	Less than 1% target error rate for supported PDF formats through structured extraction and validation.
Port lookup	Manual lookup effort and unresolved mapping count.	Repeated mappings become automatic through the stored mapping library.
Unknown PDF handling	Initial share of new or unknown customer PDF layouts that require human review.	Instantly caught as exceptions and updated within hours using Cloud28.
Digital twin coverage	No complete structured database for extracted order/container data in the initial manual flow.	100% of processed PDFs stored with source traceability and processing status.

Monthly employee hours returned to operations

This does not require laying off staff. The business value is that employees stop wasting 110-165 hours each month on repetitive PDF conversion and can instead focus on exceptions, customer service, dispatch coordination, quality checks, and other work that actually needs human judgment.

Measure	Calculation		
Order volume	10 orders/day x 22 workdays/month	220 orders/month	Baseline workload entering the manual PDF-to-Excel process.
Container rows	220 orders x 15 containers/order	3,300 container rows/month	Rows that no longer need to be manually copied into Excel.
	3,300 rows x 2-3 minutes/container	110-165 hours/month	About 0.63-0.95 of one full-time employee month returned to the business.

06 - OPERATING CONTROLS AND NEXT STEP

Automation is useful because the uncertain work is visible.

The connector should not hide uncertainty. Its purpose is to remove repetitive transcription while making uncertain work easier to identify, route, and resolve. Low-confidence extraction, unresolved ports, unsupported layouts, unusual change orders, and new customer PDF formats become explicit exceptions for staff review.

This is what allows the automation to run 24/7 without creating uncontrolled risk. The agent can process known formats immediately, while unknown formats are caught at the edge of the workflow and sent into a rapid Cloud28 update cycle. Logistics Operator gets real-time processing for normal orders and fast correction for the small number of exceptions.

The control model keeps the current manual workflow available during pilot and early production while the automated workflow earns trust. Access to the mailbox, database, output folder, and reference data should be limited to approved users and service accounts.

Each generated row should be traceable to source email, source PDF, extraction status, mapping status, generated file, upload timestamp, and reviewer action. That audit trail lets Logistics Operator inspect what happened without reconstructing the work from inboxes and spreadsheets.