Mirage Metrics

CASE STUDY

# How OFRET Groupe automated quoting across two business lines with Mirage

OFRET Groupe runs two distinct operations under one roof: OFRET for full truckload transport, SWILL for groupage. Mirage deployed a tailored Quoting Agent for each, both connected to the group's shared TMS.

PUBLISHED JULY 24, 2026. UPDATED NOVEMBER 5, 2026.

miragemetrics.com/case-studies/ofret

**2**
tailored Quoting Agent environments, one per entity

**90%+**
of standard RFQs auto-handled

**1**
shared TMS, zero extra software layers

## About OFRET Groupe

OFRET Groupe is a freight forwarding group built around two entities. OFRET handles full truckload transport, moving complete shipments directly from origin to destination. SWILL runs groupage, consolidating multiple clients' smaller shipments into shared loads. Same group, same TMS, two different pricing logics and two different carrier networks to manage.

Both entities faced the same problem: a constant stream of quote requests to price by hand, and freight documents to read and re-key into the system for every shipment. Done manually, that work does not scale, and it does not scale the same way for full truckload as it does for groupage.

## Two entities, two environments

OFRET and SWILL quote differently: different carrier networks, different document types, different rules for what counts as a "standard" request. So Mirage did not deploy one generic agent across the group. The Quoting Agent was set up as two separate, purpose-built environments, one tuned to OFRET's full truckload workflow and one tuned to SWILL's groupage workflow.

Both environments plug into the same group TMS and share the same underlying platform, so an improvement made for one, like a new verification rule or a better extraction pattern, can be adapted for the other without starting from zero.

## Where it started

OFRET Groupe already ran its operations on Logitude World, its TMS. Rather than add another software layer on top, Mirage connected directly to the Logitude World API and built the automation inside the system both entities already used.

Mirage's forward-deployed engineers worked on-site with OFRET's and SWILL's own documents and workflows, not a generic template. That embedded approach, building against real carriers, real document formats, and each team's real exceptions, is what made it possible to stand up two tailored environments instead of one compromise.

## The Quoting Agent

At the center of both environments is a Quoting Agent. When a client emails a request to move a shipment, the agent reads the RFQ, extracts what it needs to price it (origin, destination, weight, volume, incoterms, deadlines), and checks it against the entity's own carrier and subcontractor network to prepare a quote.

For most requests, the agent prepares the response automatically, cutting the time between an incoming RFQ and a client getting a quote from hours to minutes. Anything that does not fit the standard pattern is routed for review instead of guessed at.

## From documents to full RFQ automation

The deployment grew in stages instead of arriving fully formed. Mirage started with document automation: reading freight documents across the TMS and turning them into structured, usable data. Once that extraction was reliable, it became the foundation for RFQ automation, reading incoming quote requests and preparing priced responses directly. The same path for another forwarder's desk, from the email request to an approved quote in the TMS: freight quote automation, step by step.

As trust in the agent's output grew, Mirage added a verification stage: rules that check a quote or an extraction against what "normal" looks like for that entity, plus an exception layer that catches anything failing those checks and notifies an operator instead of letting it through. Each stage was built on the one before it, so the automation keeps evolving instead of stopping at a one-time delivery.

## Human in the loop

The agent handles the routine. People handle the judgment calls. A case goes to a human operator before anything goes out when:

- Data is missing on a shipment document
- The extraction confidence is low
- Pricing falls outside an entity's standard rate patterns
- A request does not match either entity's usual shipment profile

Because Mirage's engineers were embedded with OFRET's and SWILL's teams while building this, the exception rules reflect real cases the teams had actually run into, not a generic guess at what might go wrong. That relationship stayed in place after launch, so the rules keep getting refined as new edge cases surface.

## What changed

Neither OFRET's nor SWILL's teams re-key documents or price standard RFQs by hand anymore. Routine quoting that used to consume the day now runs automatically inside the group's own TMS, tailored to each entity's way of working. People focus on the shipments that actually need a decision.

## Where it goes next

- Customs documentation
- Shipment tracking
- Automated client communication
- Collections
- Accounting processes
- Warehousing tracking