


Your orders wait somewhere.

Now you see where.

Each step an order, a purchase or a warehouse transfer takes is recorded the moment it happens: who moved it, and when. One screen then shows where the time goes, the paths documents really take, how often work is done twice, and how much ran with no person at all. Every number opens the documents behind it.



Six questions it answers

- **Where the time goes:** each pair of steps, its median wait and its slowest tenth, the worst one marked
- **Paths:** the ten most common sequences of steps, and how many documents leave them
- **Rework:** orders cancelled and confirmed again, invoices taken back to draft
- **Who moves it:** a person, a shop or marketplace, a scheduled job, a rule or the AI
- **Untouched:** the share that finished with no person at all
- **Three flows:** sales orders, purchases and the warehouse floor

Why it is different

- **Recorded, not reconstructed.** Each step is written in the same database transaction as the change, into a log nothing can rewrite. No extraction project.
- **It knows who acted.** A marketplace order is the channel's, a scheduled run is the job's, a request through the AI door is the AI's, so untouched means untouched.
- **Yours.** In your system, on your server, with your code. No second licence; it comes with the system, priced per business flow, not per user.

Specialist process-mining tools sell this picture as a separate product, after a project to extract it from your ERP. In an a-systems system **it is recorded where the work happens, from the first day.**