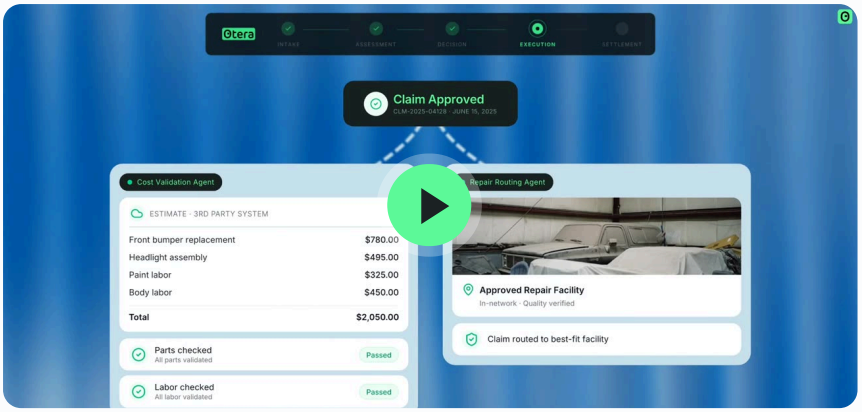
Motor Claims Runs Itself.

> **An autonomous operation cannot run on a probabilistic guess. It requires deterministic code: the same decision every time, enforced, with the exact clause behind it. Hold every option you evaluate to the thirteen criteria inside.**



SEE THE AGENT RUN

One motor claim, end to end

FNOL to settlement, with every decision logged and the estimate checked before approval.

Scan or click to watch

IN PRODUCTION AT CED | COVERED BY Gartner FORRESTER IDC HFS

The challenge, and the solution

The process looks digital. The work is still manual. Claims arrive through portals, inboxes, phone calls, and paper, and people still read, classify, and route them by hand. Every team runs its own version of the process, and a decade of OCR, rules engines, workflow tools, and RPA made the steps faster without ever deciding anything. The decision itself still rests on a probabilistic judge in the middle: an AI model guessing, or a person interpreting, and neither can prove its answer.

Replace that judge with deterministic code, and three things fall out of one move:

- **One connected system.**
  Every step reads and writes one graph. Same input, same decision, enforced by compiled logic. A new line or region is configuration.
- **Intelligence capital you own.**
  A complete mathematical model of your business, yours to own and export, with core logic immune to model churn.
- **Control with a dial.**
  Every decision cites its clause with a calibrated confidence score, and you steer strictness from the business seat.

What it changes on the P&L

TOPLINE

- **Win more at the point of sale.**
  Claims service becomes the differentiator, not price.
- **Launch faster.**
  New products and new regions go live typically in under six weeks, not as a yearlong program.
- **A significant decrease in churn.**
  The claim is where loyalty is decided; settled in days, churn typically drops by 10% or more.

BOTTOM LINE

- **A fully elastic workforce, at a fraction of today's cost.**
  Capacity flexed from 500 to 50,000 cases a day on the same system, overnight.
- **Leakage recovered.**
  Every file settles exactly to the policy terms, typically recovering 1 to 3 points of claims spend.
- **Scale the team you have.**
  Provably correct decisions need no second reviewer; capacity scales 10x or more without seasonal hiring.

The Motor Claims Agent: eleven sub-agents, one decision graph

The Motor Claims Agent runs a motor claim end to end, from first notice of loss to subrogation recovery, through eleven sub-agents on one decision graph. **Precision AI translates your policies, T&Cs, contracts, SOPs, and documented tribal knowledge into deterministic code; code makes every decision.** Before the first claim, it reads your policy wording for the six kinds of gap no decision can get past, from contradicting clauses to evidence no customer can obtain, and proposes the rewrite.

AI READS · CODE DECIDES · IN PARALLEL · SUB-AGENTS ACT

1	2	3	4	5	6	7	8	9	10	11
FNOL Intake	**Document Processing**	**Coverage / Proof**	**Fraud Detection**	**Liability**	**Damage Assessment**	**Estimation**	**Total Loss**	**Settlement**	**Rental & Repair**	**Subrogation & Recovery**
App, web, phone, email; every motor sub-type and language.	Accident statements, police reports, photos, invoices; 100+ types, multilingual.	Coverage, exclusions, and territorial limits decided deterministically, with the clause.	Claimant history, repair-shop patterns, timing; alongside your fraud system.	Negligence frameworks, multi-vehicle and cross-border; the split established.	Photos to components, severity, and repair vs. total loss.	Audatex, DAT, CCC, or Mitchell; parts and labor rates validated.	Actual cash value, comparables, salvage, owner retention.	Payment workflows, authority rules, reserves, regulatory timing.	Rental authorization, duration monitoring, repair-timeline alignment.	Recovery opportunities, demand letters, cross-border coordination.

Runs on top of the operation you already have. It connects through your APIs, your database, or an RPA bridge, writes every outcome back, and goes live with review in place. Nothing migrates. Core systems include:

Guidewire · Duck Creek · SAP · In-house & mainframe-era

In production, not in pilots

In production in 12 weeks: not a pilot, not a POC, live autonomous operations processing real claims. The platform behind it runs at scale in production today:

TOP-FIVE GLOBAL INSURANCE GROUP

~99.7% end-to-end decision accuracy

~80% straight-through processing

A business unit runs FNOL to payment with no human in the middle, multi-country, multi-language, multi-line. It started exactly like this: **one line, one contained pilot.**

CED GROUP

600,000+ annual claims

Same-day resolution

Backlog permanently eliminated; 80% of manual claims staff redeployed to complex cases. 500 to 50,000 cases a day **with nothing rebuilt.**

Mid-sized operations typically reach production faster, with less organizational complexity in the way.

SOC 2 Type II · ISO 27001 · GDPR compliant · EU AI Act addressed · DORA addressed

Thirteen criteria for an agentic platform

Thirteen criteria, assembled from what Gartner, Forrester, and the EU AI Act already score. They compare every option on the same terms, including us.

A · Decision quality

1. **Decision consistency.** Rerun it as often as you like; the decisions are consistent.
2. **Human-AI delegation.** Every case is either fully automated or stops for a specific, cited reason.
3. **Decision traceability.** Every decision carries the criteria and the facts that produced it.

B · Governance, security, lifecycle

4. **Agentic governance.** Change the underlying logic or policies, and every decision follows.
5. **Enterprise readiness.** It passes your security review.
6. **Observability and drift.** You know when quality moves, because it is measured against ground truth.

C · Architecture and independence

7. **Neuro-symbolic architecture.** Repeatable decisions come from the architecture, not the model.
8. **Interoperability and openness.** Your tools can call the platform, and the platform can call your tools.
9. **Model independence and upgrade safety.** A model change cannot set you back.
10. **Data and AI sovereignty.** It runs where you require, and your data stays within the boundaries you set.

D · Fit and delivery

11. **Operating-model encoding and ownership.** Your policies decide your cases; the intelligence they encode stays your IP.
12. **Speed to production.** Each subsequent use case costs a fraction of the previous one.
13. **Track record and references.** It should already be running in production, at a reference you can call.



See it run, end to end

A demo walks one motor claim from FNOL to payment. Bring the thirteen criteria and judge what you see against them. From there, we scope what a first use case could look like on your book.

Request a demo
otera.ai/get-demo

Results from production deployments, as published by Otera. Deployment timing is scoped per operation. © 2026 Otera · otera.ai