

INSURANCE

ASIA PACIFIC

FRAUD, WASTE & ABUSE DETECTION

AI Accuracy at Rule-Level Cost

How a global life insurer achieved 94% fraud detection at 71% lower cost per claim with the Trust Intelligence Platform.

EXECUTIVE SUMMARY

The insurer's fraud detection was stuck: rules caught 58%, ML reached 71%, and LLMs hit 89% but at \$540K/year. With a CFO demanding 10:1 ROI, the math didn't work. Rotascale's Trust Cascade architecture delivered AI-level detection at rule-level cost: 94% fraud detection, 71% lower cost per claim, \$55K annual spend. The cascade paid for itself in the first quarter.

94%

DETECTION RATE

-71%

COST PER CLAIM

+\$2.8M

ADDITIONAL FRAUD PREVENTED

16

WEEKS TO PRODUCTION

AT A GLANCE

INDUSTRY

**Life & Health
Insurance**

REGION

Asia Pacific

COMPANY SIZE

Global Enterprise

PRODUCTS USED

**Guardian,
Orchestrate, APLS**

TIMELINE

16 weeks

THE CHALLENGE

The Math Didn't Work

The insurer's fraud detection system was a patchwork of rules and ML models. A recent LLM proof-of-concept showed promise, but projected costs made AI fraud detection unaffordable at scale.

Detection vs. Cost Trade-off

Approach	Detection	Challenge
Pure Rules	58%	Too many fraud cases missed
Rules + ML	71%	40% false positive rate
Rules + ML + LLM	89%	\$540K/year - unaffordable

The Business Case Gap

The CFO set clear requirements for any AI investment:

- ✓ Estimated \$12M annual fraud losses
- ✓ LLM solution would cost \$540K/year
- ✓ CFO demanded 10:1 ROI minimum
- ✓ Solution needed to cost <\$120K/year

Operational Problems

- Investigator Overload**
40% false positive rate meant investigators spent more time chasing false alarms than catching real fraud.
- Sophisticated Fraud Slipping Through**
Rule-based systems caught obvious patterns but missed complex, coordinated fraud schemes.
- No Explanations**
When claims were flagged, investigators couldn't understand why. Black-box decisions slowed review.
- CFO Skepticism**
Previous AI investments hadn't delivered promised ROI. The fraud team needed proven economics.

"Everyone told us we needed AI for fraud detection. Nobody told us we'd go bankrupt running it at scale."

— Chief Risk Officer

THE SOLUTION

Trust Cascade Architecture

Instead of choosing between cheap-but-limited rules and accurate-but-expensive AI, the Trust Cascade routes each claim to the right level of intelligence based on complexity and economic value.

Five-Level Cascade

1	Rules Engine Known patterns, velocity checks, duplicate detection	\$0.0001	68% of claims
2	ML Models Anomaly scoring, risk classification, pattern matching	\$0.001	22% of claims
3	Single Agent Complex pattern analysis, contextual reasoning	\$0.008	7% of claims
4	Agent Panel Multi-perspective review, specialist consultation	\$0.025	2% of claims
5	Adversarial Debate Prosecution vs. defense argumentation	\$0.045	1% of claims

ROI-Based Routing

Claims route based on economic value, not just complexity:

- ✓ **Claims <\$1K:** Max Level 2 (not worth AI cost)
- ✓ **Claims \$1K-\$10K:** Max Level 3 (single agent sufficient)
- ✓ **Claims \$10K-\$50K:** Max Level 4 (panel review justified)
- ✓ **Claims >\$50K:** Full cascade (adversarial debate)

Self-Learning Rules (APLS)

The system gets cheaper over time:

- ✓ When Levels 3-5 catch fraud, system extracts pattern
- ✓ Generates candidate rule for Level 1 or 2
- ✓ Human review and approval workflow
- ✓ Auto-deploys to lower, cheaper levels
- ✓ 127 new rules generated in first 6 months

PRODUCTS USED

Trust Intelligence Platform for Insurance FWA

Guardian monitors detection accuracy and model drift. Orchestrate coordinates the multi-level cascade with full reasoning capture. APLS (Automatic Pattern Learning System) converts AI insights into production rules.

Guardian

Orchestrate

APLS

RESULTS

AI Accuracy at Rule-Level Cost

16 weeks from kickoff to production. The cascade delivered detection rates better than full LLM deployment at a fraction of the cost.

Metric	Before	After	Change
Detection rate	71%	94%	+32%
False positive rate	40%	12%	-70%
Cost per claim (at 2M claims)	\$0.008	\$0.0023	-71%
Annual detection cost	\$180K	\$55K	-69%
Fraud prevented (estimated)	\$8.5M	\$11.3M	+\$2.8M

Implementation Timeline

- WEEKS 1-2**

Analysis

18 months of claims data. Fraud pattern mapping. Economic modeling.
- WEEKS 3-5**

Cascade Design

5-level architecture. ROI-based routing logic. APLS feedback loop.
- WEEKS 6-12**

Implementation

Rules integration. ML retraining. Multi-agent debate system.
- WEEKS 13-16**

Optimization

Production tuning. APLS activation. Team training.

Additional Outcomes

- ✓ **CFO approved expansion** to health claims based on proven ROI
- ✓ **Investigation productivity up 3x** - investigators focus on real fraud
- ✓ **127 new rules auto-generated** in first 6 months via APLS
- ✓ **System improving monthly** - detection migrating to cheaper levels

"Rotascale showed us how to get AI-level accuracy at rule-level cost. The cascade paid for itself in the first quarter."

— Chief Risk Officer

Facing similar fraud economics challenges?

Let's discuss how Trust Cascade can deliver AI detection at sustainable cost.

Request Demo

Insurance Solutions