

## PRIIMAL TECHNOLOGIES AS

A Vantage product

## WHITE PAPER · PRIIMAL CLAIM

# The SIU Detection Gap

Why Organised Claims Fraud Survives Conventional Investigation

Priimal Technologies AS · October 2026 · A Vantage product

**ABSTRACT**

Insurance fraud detection has made meaningful progress over the past two decades. Special investigation units are better resourced, fraud analytics platforms are more sophisticated, and regulatory expectations are higher. And yet, the fraud that costs insurers the most continues to go largely undetected. The reason is structural, not operational: organised fraud survives because the network patterns that identify it are distributed across dozens of individual claim files in systems that were never designed to be cross-referenced simultaneously. This paper describes that structural detection gap, explains why conventional SIU triage cannot close it, and demonstrates what cross-claim, full-population scoring makes visible that nothing else can.

**Document type:** Public, Tier 1: no registration required**Audience:** SIU directors, claims investigation leads, heads of claims, actuaries, reinsurance professionals**Product line:** Priimal Claim: A Vantage product

## The fraud that never gets investigated

Every insurer with a functioning special investigations unit catches some fraud. Staged accidents with implausible injury narratives. Repeat claimants with obviously inconsistent loss histories. Repair invoices from garages that do not exist. These cases reach the SIU because they are individually conspicuous: a single claim handler, a billing exception rule, or a complaint triggers a referral. The SIU investigates. The claim is denied or settled at reduced value. The case is closed.

This is not the fraud that is costing insurers the most. The fraud that generates the largest aggregate financial losses operates at the opposite end of the visibility spectrum: individually unremarkable claims, each within normal parameters, each filed by a different apparent claimant, each handled by a different claims processor, none of which triggers a referral, none of which ever reaches an SIU desk. The connection between them is invisible in any single claim file. It is visible only when the full claims population is scored simultaneously against the patterns that identify coordinated activity.

The Coalition Against Insurance Fraud and the NICB estimate that up to 10% of all insurance claims contain some element of fraud, representing billions of dollars in annual losses. Research by Shift Technology indicates that 2 to 3 percentage points of a typical insurer's combined ratio can be attributed to fraudulent claims. In a market where many European non-life insurers are operating with combined ratios above 100%, as documented by the German Insurance Association (GDV) and the Insurance Information Institute, those 2 to 3 points are not a rounding error. They are the margin.

*The organised fraud that generates the largest losses is not visible in any individual claim file. It is visible only when the full claims population is scored against network patterns simultaneously. Until that happens, those losses are not being investigated. They are being paid.*

## How conventional SIU triage works, and what it misses

---

### The referral system

In the majority of motor and health insurers, claims reach the SIU through one of three routes: referral from a claims handler who notices something unusual; a flag from a rules-based exception system that identifies claims above a threshold or matching a defined pattern; or an external report from a fraud bureau, law enforcement body, or third-party information source. Each of these mechanisms has the same structural limitation: they identify claims that are individually suspicious. They cannot identify claims that are individually normal but collectively organised.

A staged accident ring operating across thirty claims files does not need any single claim to be suspicious. The ring coordinator ensures each claim is individually plausible: the injury is soft-tissue, the damage is consistent with the reported incident, the medical provider bills at standard rates, the repair invoice is from a legitimate-looking garage. The individual claim handler who processes each file sees a normal motor claim. The rules engine does not flag it. No referral is made. The ring collects thirty settlements and moves on.

### The unstructured data problem

Conventional fraud detection platforms are built primarily on structured data: claim amounts, claimant demographics, policy dates, incident locations, repair costs. These fields are consistent, comparable across claims, and straightforward to score against historical fraud patterns. They are also, on their own, insufficient to detect the most sophisticated organised fraud.

Industry analysis consistently finds that unstructured data, principally adjuster notes, first notice of loss narratives, repair invoice descriptions, and medical report free text, contains between one-third and one-half of the most reliable fraud signals in a typical claims dataset. A SAS Institute analysis cited by the Coalition Against Insurance Fraud found that only 40% of insurers use text mining as an anti-fraud tool. The remaining 60% are making fraud detection decisions based on roughly 20% of the available information, according to analysis cited in insurance fraud technology literature.

The adjuster note that records that a claimant described the accident in nearly identical language to a claim from six months ago, handled by a different adjuster in a different branch, involving a different vehicle and a different reported location, does not trigger a rules-based flag. It is free text. It is not compared against other free text records. The pattern it contains is invisible until both notes are scored simultaneously against language similarity signatures.

### The threshold problem

Rules-based exception systems are designed to catch claims that exceed defined thresholds: repair costs above a benchmark, medical bills beyond an expected treatment protocol, injury claims without police reports. These thresholds are calibrated against historical fraud patterns. Organised fraud rings are calibrated against these thresholds. The ring coordinator knows what a claim needs to look like to avoid flagging: repair costs just below the benchmark, treatment protocols that follow standard documentation, police reports obtained for every claim regardless of the incident's actual nature.

The result is a detection system that is structurally optimised to catch the previous generation of fraud. The fraud that has adapted to the threshold system goes undetected, because the threshold system was not designed to detect adaptation. It was designed to detect excess.

## What cross-claim scoring makes visible

---

The patterns that identify organised fraud are not in any individual claim. They are in the relationships between claims: shared participants who appear across incidents with different vehicles and different injuries but consistent contact details or device signatures; repair facilities that invoice for the same parts at the same prices across unrelated incidents; medical providers whose billing patterns show templated treatment protocols inconsistent with the stated injuries; solicitors whose client referral patterns suggest a coordination function rather than individual client relationships.

These network patterns are present in the claims data that every insurer already holds. The structural reason they are not being identified is that the data exists in individual claim records that are examined sequentially, one at a time, by handlers who have no visibility of what other handlers have processed. The information needed to identify the network is there. The architecture to surface it is not.

| Detection approach                             | What it identifies                                                                                                                                                                                | What it misses                                                                                                         | Detection lag                                 |
|------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|
| <b>Handler referral</b>                        | Individually conspicuous claims: implausible narratives, known claimants, obvious document irregularities                                                                                         | Organised rings where individual claims are normal; network patterns across handlers and branches                      | Weeks to months after payment                 |
| <b>Rules-based exception flags</b>             | Claims exceeding defined thresholds on structured fields: cost, frequency, geography                                                                                                              | Claims calibrated below the threshold; soft fraud in free-text fields; network patterns not visible in structured data | At point of exception; misses adapted fraud   |
| <b>Periodic SIU audit</b>                      | Sample of claims reviewed for patterns; occasional network discovery                                                                                                                              | Everything between audit cycles; patterns that span long time windows or many claims                                   | Months to years after accumulation            |
| <b>External bureau alerts</b>                  | Known fraudulent identities, vehicles, and providers on shared registers                                                                                                                          | First-time fraud; rings using new identities; provider collusion not yet on the register                               | Only after registration, which lags discovery |
| Full-population AI scoring including free text | Network patterns across all claims simultaneously; language similarity in free-text fields; provider and participant clustering; adapted threshold fraud; soft fraud signals in unstructured data | Nothing that exists in the data at time of scoring                                                                     | At submission, before payment                 |

Sources: Coalition Against Insurance Fraud; Aite-Novarica Group, *AI-Based P/C Claims Fraud Detection Solutions* (2022); SAS Institute, *insurance fraud analytics publications*; Shift Technology, *Claims Fraud Benchmark Report* (2023).

## The combined ratio dimension

The financial case for closing the detection gap is not expressed in fraud statistics. It is expressed in combined ratios. Shift Technology research indicates that 2 to 3 percentage points of a typical insurer's combined ratio are attributable to fraudulent claims. In the German non-life insurance market, the GDV estimates that insurance fraud generates over EUR 6 billion in annual claims expenditure, approximately 10% of total claims spend. ZEB consulting analysis suggests that effective fraud management can reduce an insurer's combined ratio by up to 3 percentage points.

In practical terms: for a motor and health insurer with EUR 200 million in non-life gross written premium and a 65% loss ratio, gross claims paid are approximately EUR 130 million annually. A 7% fraud rate on that base implies EUR 9.1 million in annual fraud losses. If 40% of that fraud is identified and investigated through conventional triage, EUR 5.4 million is going uninvestigated each year, compounding with each claims cycle. At a 50% recovery rate on investigated claims, the gap between current detection and full-population scoring represents approximately EUR 2.7 million in annually recoverable value for an insurer of that size.

That is not a fraud statistic. That is a combined ratio problem. And the insurers closing that gap are not doing so by hiring more SIU investigators. They are doing so by changing the architecture of what gets scored and when.

*Every combined ratio point that comes from undetected fraud is a point that could have been recovered. It was not a question of whether the data existed. It was a question of whether all of it was scored.*

## Priimal Claim: A Vantage product

---

Priimal Claim is an AI-scoring and case management platform for motor and health insurance claims integrity. Insurers upload their existing claims data, including free-text adjuster notes, first notice of loss descriptions, emails, and call transcripts. Vantage scores every claim simultaneously across six fraud flag categories: Fraud Ring, Staged Accident, Provider Collusion, Inflated Repair, Duplicate Claim, and Soft Tissue. Every claim is scored, not a referral-based sample: the platform processes the full claims population in each scoring cycle, including free-text fields, which is where the organised fraud signals reside.

Each case in the output queue includes a risk score from 0 to 100, a confidence level, a flag breakdown, and a plain English explanation of why the case was flagged. Scores are decision support for investigators, not automated decisions: a case is reviewed by the SIU team before any action is taken. The case management workflow includes threaded investigation notes, audit trail, field investigation dispatch, status tracking, and a branded case export with a subrogation indicator built to stand in legal and regulatory contexts.

Available under Core Value, Priimal's no-cure-no-pay subscription: the insurer pays a share of documented recovered value only, capped at the flat-rate package equivalent, and nothing if nothing is recovered. Requires a signed Value Agreement; not self-serve. To discuss a Claim Flash diagnostic using your own claims data, contact [connect@priimal.com](mailto:connect@priimal.com).

*Priimal Technologies AS · [priimal.com](https://priimal.com) · [connect@priimal.com](mailto:connect@priimal.com) · A Vantage product. Published October 2026. Figures derived from publicly available sources. This document does not constitute legal, actuarial, or investment advice. © 2026 Priimal Technologies AS. All rights reserved.*