ZARELVA · Fraud Intelligence & Risk Architecture

ZIC Engine — Investigator Canvas

Operating Guide

A practical guide to investigating fraud rings, resolving entities, and dispositioning cases on a single canvas

At a glance	
Product	ZIC Engine — Investigator Canvas
Guide version	1.0
Issued	26 June 2026
Audience	Fraud and risk analysts and their teams
Classification	Shareable — client and evaluation use

Structure Controls Risk.™

Contents

1. About ZIC Engine

ZIC Engine — Investigator Canvas is a single-screen console for fraud and financial-crime investigation. It takes the entities and transactions involved in a case and presents them as an interactive relationship graph alongside an investigator workspace, so an analyst can see suspected fraud rings, drill into individual parties, weigh an AI recommendation, verify evidence, and record a decision — without switching between tools.

It is built for money-mule and fraud-ring detection, entity resolution, transaction-flow review, and case disposition, and it runs in a standard desktop browser.

2. Who it is for, and where it runs

The canvas is designed for fraud and risk analysts and the teams that review their work. It opens in a current desktop browser (Chrome, Edge, or Firefox); the investigator experience runs locally, with case data and detection sourced from your environment in a production deployment.

3. What you can do

- **Investigate visually.** See rings, entities, and the connections between them on one canvas, and isolate any ring with a single click.
- **Detect rings automatically.** Surface coordinated clusters and the links between them from connectivity in the data.
- **Resolve entities.** Identify likely duplicate parties and merge them into a single, canonical entity.
- **Triage with AI assistance.** Read a recommendation, its model confidence, and the explainable risk factors behind a score.
- **Keep evidence tamper-evident.** Hash attachments with SHA-256 and verify integrity on demand.
- **Maintain a complete audit trail.** Every material action is recorded for review.
- **Disposition cleanly.** Record a standardised outcome with rationale, then save, report, and export the case.

4. The investigation workflow

A typical case follows nine steps, from first launch to a saved record.

4.1 Launch and orient

1. Open the canvas in your browser. Read the hero strip — fraud rings, Overall Case Risk, and the AI recommendation — for an instant read on the case, then scan the left rail for case health and the rings present.

4.2 Import case data

1. Load the case from a CSV of entities, a CSV of transactions, or a saved case bundle. Review the import summary, then confirm the data has loaded into the graph.

4.3 Detect rings

1. Run ring detection to group connected parties into candidate rings and surface the links between them. New rings appear in the legend and the quick-focus chips.

4.4 Investigate the graph

1. Click a ring to spotlight it and dim the rest of the network, then click individual entities to open their profiles — risk, signals, history, and linked parties. Where duplicates are flagged, review and merge them.

4.5 Review AI insights

1. Read the recommendation, its confidence, and the risk factors that explain the score, and confirm they match what the graph shows. Accept the recommendation or override it with your own reasoning.

4.6 Verify evidence

1. Attach supporting files, categorise them, and link them to the relevant entity or ring. Each file is fingerprinted with SHA-256; use Verify Integrity to confirm a file is unchanged before you rely on it.

4.7 Capture notes

1. Record observations, hypotheses, and next steps as you work, and save them with the case. Good notes let a reviewer follow the investigation on their own.

4.8 Record the decision

1. Choose a standard outcome — proceed to investigation, close as false positive, escalate, or add to watchlist — and submit it. The decision and its trail are recorded with the case.

4.9 Save and export

1. Save the case, generate a report, or export the full archive for hand-off. Saved cases can be restored later for continued work or review.

5. Understanding the metrics

ZIC Engine keeps four measures distinct so the picture stays clear. Each answers a different question.

Measure	What it tells you
Overall Case Risk	Case-level severity (0–100, banded), driven by the highest-risk entity — a prioritisation signal, not a verdict.
Model Confidence	How certain the AI is in its recommendation (0–100) — a measure of certainty, not of risk.
Ring detection confidence	Per ring, how strongly the cluster reads as coordinated — a prompt to review, not proof.
Risk factors	The itemised signals (0–100 each) that explain a score, so the reasoning is transparent.

Risk bands: Low (0–39), Elevated (40–69), and High (70–100) set the level of attention a case warrants.

6. Evidence integrity and audit trail

Two features make a case easy to defend in review. Evidence files are fingerprinted with SHA-256 on attachment, and Verify Integrity re-checks that fingerprint on demand — a match confirms the file is unchanged. Separately, the audit trail records every material action on a case, so the investigation can be reconstructed from the record alone.

7. Scope, assumptions, and responsible use

ZIC Engine is designed to augment analyst judgement, not to replace it, and it is not a system of record.

- **Evaluation vs production.** In evaluation builds, selected AI narrative and enrichment content is illustrative; in a production deployment these are configured against your own data, models, and intelligence sources.
- **Corroboration.** Findings should be corroborated against your authoritative systems of record before any regulatory filing or customer action.
- **What operates as described.** Ring detection, entity resolution, SHA-256 evidence verification, the audit trail, and case save/restore/export work as set out in this guide.

> **In short.** ZIC Engine accelerates investigation and keeps it transparent and reviewable; analysts remain responsible for verifying findings and for the actions taken on a case.

8. Glossary

Term	Definition
AML / KYC	Anti-Money-Laundering / Know-Your-Customer controls for detecting illicit funds and verifying identity.
Fraud ring	A connected cluster of entities acting together to commit or launder the proceeds of fraud.
Money mule	A person or account used to receive and move illicit funds for others.
Entity resolution	Determining that records refer to the same real-world party and merging them.
SAR	Suspicious Activity Report — a regulatory filing raised where activity indicates possible financial crime.
SHA-256	A cryptographic hash used as a tamper-evident fingerprint of an evidence file.

9. About Zarelva

Zarelva is a fraud intelligence and risk architecture practice. ZIC Engine is part of its work to make fraud investigation structured, transparent, and defensible — because structure controls risk.

To arrange a walkthrough or discuss a deployment, contact **hello@zarelva.com** or visit **zarelva.com**.