

How AI proactively identifies, flags, and resolves problems across the platform

Sona asked us to show how an AI agent acts on its own across the platform. The agent runs a three-stage loop on every screen — **identify** → **flag** → **resolve**. Below is exactly where each stage shows up.

1 · IDENTIFY	2 · FLAG	3 · RESOLVE
<ul style="list-style-type: none"> • Scan ingested data for schema, currency, dedup issues • Match every invoice line vs contract and rate card • Detect 80/20 patterns across exception categories • Track 4-quarter compliance trend per vendor 	<ul style="list-style-type: none"> • Surface findings with confidence % and \$ impact • Color-code: red flag, amber suggest, green auto-resolved • Connect single invoices to systemic patterns ("+17 like this") • Always cite evidence — never act silently 	<ul style="list-style-type: none"> • One-click batch recovery for high-confidence patterns • Auto-fix routine issues; log for audit • Hand off to human at amber-colored pipeline pause • Recommend specific, dollar-anchored next action

WHERE THE AGENT SHOWS UP — SCREEN BY SCREEN

<p>SCREEN 1</p> <p>Compliance Overview</p> <p>Identify Cross-vendor pattern across 17 invoices with the same surcharge miscalculation.</p> <p>Flag AI insight panel surfaces \$48K recoverable + 84% confidence + evidence chain.</p> <p>Resolve Two CTAs: Open batch recovery (17) or Review one-by-one.</p>	<p>SCREEN 2</p> <p>Data Quality Report</p> <p>Identify Schema mismatches, currency inconsistencies, duplicates, vendor naming drift.</p> <p>Flag Agentic pipeline shows step-by-step status; pauses amber at vendor harmonization.</p> <p>Resolve Auto-resolves dedup + currency. Hands the ambiguous 6 vendors to a human.</p>
<p>SCREEN 3</p> <p>Invoice Detail</p> <p>Identify Line-level deviation: surcharge tier 3 applied where contract requires tier 1.</p> <p>Flag Exception card with deviation type, \$ impact, root cause, recommended action.</p> <p>Resolve Connects to 17 sibling invoices; offers batch recovery or cautious one-by-one.</p>	<p>SCREEN 4</p> <p>Root Cause Analysis</p> <p>Identify Pareto distribution of all 8 exception categories by frequency × \$ impact.</p> <p>Flag 3-tier insights: systemic (red), predictive (amber), quick-win (green).</p> <p>Resolve Drill-through to affected invoices; suggested rate-card update for owner.</p>

THE AGENT'S THREE ACTION MODES

<p>FLAG FOR HUMAN</p> <p>High impact + confidence ≥ 75%. Always presents evidence first; no auto-action.</p>	<p>SUGGEST ACTION</p> <p>Confidence 50–75%. Shows reasoning. Action button requires explicit approval.</p>	<p>AUTO-RESOLVE</p> <p>High confidence + low impact. Logged with rationale. User can audit at any time.</p>
---	---	--

Design rules for the agent · Transparency over magic — confidence and evidence shown before any action. · Anchored to dollars — every recommendation cites \$ impact. · Never silent — every action is logged and auditable. · Visually consistent — purple #6D28D9 marks every AI moment.