



Castellum.AI

FORTUNE
50
GLOBAL
CORPORATE

How a Fortune 50 Reduced Alert Review Time by 80% with Arbiter AI Agents

80%

Reduction in time spent
reviewing alerts

4,000+

Analyst hours saved annually

Case Highlights

Client: Fortune 50 Global Corporate

Use case: Sanctions, PEPs, Adverse Media Alert Reviews
across Third Party Risk Management (TPRM)

Challenge: A global Fortune 50 enterprise, managing
thousands of suppliers and 30,000+ screening alerts
annually, struggled with the growing cost and operational
burden of manual alert review.

Solution: Castellum.AI's Arbiter agents are integrated within
the client's existing screening and case management system
to instantly review and adjudicate L1 alerts.



Client Overview

A global Fortune 50 enterprise with thousands of international suppliers, vendors and distributors faced increasing pressures on its third party risk management (TPRM) program. High false positive rates for sanctions screening, global reputational risk and adverse media monitoring overburdened the risk and compliance team.

The organization was processing more than 30,000 alerts annually, with multiple full-time analysts dedicated to L1 alert triage. Rising alert volume and high false positive rates were driving unnecessary review time and cost.

By implementing Castellum.AI's Arbitrator Alert Resolution Agents within its existing screening and case management systems, the Fortune 50 achieved:

- 80% reduction in time spent reviewing alerts
- Real-time adjudication of routine alerts
- 4,000 analyst hours recovered annually
- A scalable compliance model without expanding headcount

Challenge: Rising L1 Alert Rates

Within this large-scale compliance environment, three structural pressures emerged:

- **Manual L1 Review Dominated Resources:** A significant portion of analyst time was devoted to reviewing and classifying L1 alerts, rather than higher-risk investigative analysis.
- **Operational Sensitivity to Volume Growth:** Increased alert volumes would have required additional headcount to maintain service levels under the existing model.
- **Strict Governance Expectations:** Any efficiency gains needed to align with established risk appetite, preserve decision traceability and ensure human review remained central for consequential cases.

The objective was controlled efficiency within an enterprise compliance framework, not automation for automation's sake.



The Solution

Arbiter: AI Agents for Structured Resolution

The organization deployed Arbiter AI Agents to instantly review, investigate and triage sanctions, PEPs and adverse media alerts. Arbiter was integrated within the client's existing TPRM system that onboards and monitors the organization's global supply chain.

By embedding directly into the client's existing risk and compliance workflow,

Arbiter AI Agents:

- Evaluate alerts using contextual logic across multiple decision factors
- Classify alerts at the structured L1 stage
- Escalate higher-risk or ambiguous cases to human reviewers
- Produce audit-ready, explainable decision narratives

Customization to Risk Appetite

Unlike static rules engines, Arbiter adapts to institutional policy inputs and applies structured reasoning, bridging manual review and automation. In this case, Arbiter was calibrated to reflect the organization's existing escalation thresholds and decision logic.

Resolution thresholds were aligned with:

- Internal governance policies
- Escalation frameworks
- Oversight and audit expectations

This ensured that AI-assisted resolution reduced repetitive workload without compromising true positive handling.



Modular Deployment

Castellum.AI's modular architecture allowed the enterprise to deploy Arbiter in a configuration aligned to its existing operations.

The organization implemented Arbiter as a plug-and-play AI resolution layer within its existing screening and case management systems, requiring minimal implementation effort and no infrastructure changes.

The Results

Following implementation of Arbiter, the Fortune 50 enterprise achieved measurable reductions in alert volume, review time and operational spend.

- 1 80% Reduction in Time Spent Reviewing Alerts:** Arbiter eliminated 80% of false positives, dramatically reducing the compliance team's L1 alert review workload.
- 2 Over 4,000 Analyst Hours Recovered Annually:** By resolving alerts with Arbiter, the organization freed thousands of hours previously spent on adjudications, allowing analysts to focus on investigative work.
- 3 Real-Time Alert Resolution:** Routine alerts are now adjudicated instantly within existing screening and case management systems, eliminating backlog accumulation and stabilizing review throughput.
- 4 Regulator-Aligned, Human-in-the-Loop Decisioning:** Every automated resolution is supported by structured, explainable reasoning aligned to internal adjudication policies and human oversight built-in, ensuring audit readiness and regulatory confidence.



About Castellum.AI

Castellum.AI helps AML teams eliminate alert overload. With AI agents that resolve alerts across customer onboarding, payments and investigation workflows, we enable teams to process investigations 6x faster while eliminating 94% of false positives.

Our Arbiter agents are trained on each institution's own internal policies and procedures, ensuring every decision reflects that institution's specific risk appetite. There's no one-size-fits-all model: our agents work the way your compliance team works.

Arbiter agents can be deployed as a modular AI layer within existing workflows, or as part of a consolidated Castellum.AI solution that includes global risk data – sanctions, PEPs, adverse media and more – with real-time screening and monitoring.

For more information, visit Castellum.AI or follow us on LinkedIn.

Trusted By



Get in Touch

For more information: [Book a demo](#)

Contact: sales@castellum.ai