



Castellum.AI

Adverse Media: Buyer's Guide

How to assess screening features for speed, accuracy and real-time risk in 2025 and beyond.

www.castellum.ai

Executive Summary

Adverse media screening is a must for identifying Reputationally Exposed Persons (REPs) and mitigating associated financial crime risks. For most compliance teams, however, it creates more problems than it solves.

Poorly configured, legacy solutions generate endless false positives and bury real risks under irrelevant alerts. Compliance teams waste 100+ hours annually clearing these low-quality alerts, resulting in delayed decision-making and inflated operational costs.

The right screening solution changes this. In this guide, we'll outline the key features to look for in an adverse media screening solution and guide you through the vendor evaluation process so you can find the best fit for your compliance needs.



Table of Contents

Comprehensive and Timely Data Coverage

Why broad source coverage matters: global, local, sector-specific sources	5
Importance of real-time monitoring	6


Key Features that Improve Screening Accuracy

Negative news categorization and advanced filtering	8
Source Reliability Controls	9
Data enrichment and standardization	9
Relevancy scoring and result deduplication	10
Configurable risk thresholds	10
Clear context for faster alert disposition	11
Complete audit trail	11

Final Consideration


Cost vs. Value: What's worth paying for	11
Questions to ask during vendor evaluation	12





Comprehensive and Timely Data Coverage





Why source coverage matters: global, local, sector-specific sources


When evaluating an adverse media solution, look for comprehensive data coverage to ensure relevant risks don't go unnoticed. The solution should pull from a broad range of structured and unstructured data sources, including global news outlets, regulatory lists, government press releases and proprietary databases to ensure a more complete view of potential exposure.

It's also important to remember that many risks first emerge in jurisdiction-specific media or industry-focused reports before making their way to major international headlines. That's why coverage of local-language sources and regional media is essential. Accurate translations and access to historical archives further ensure no context is lost and equip compliance teams with complete visibility into how risk changes over time.

Castellum.AI in action

Castellum.AI contracts with media vendors to secure negative news articles from over 225,000 global sources covering 20+ languages across local outlets, global news publishers, government and law enforcement press releases and more. More than 250,000 relevant, adverse articles are added daily to our database of ~2 billion articles.





Importance of real-time monitoring


Adverse media risk is never static. A screening solution that only runs periodic checks creates blind spots, leaving organizations exposed to risks that surface after the last check.

The best solutions ingest live media data in real time and trigger automated alerts when there's a change in risk, so organizations can respond immediately to new developments. Continuous, real-time negative news monitoring replaces scheduled or manual searches. As a result, compliance teams stay proactive instead of playing catch-up.

Castellum.AI in Action

Our news feed updates every 5 seconds to deliver the latest information and flag changes in risk status instantly, eliminating 99% of the recurring alerts.





Key Features that Improve Adverse Media Screening Accuracy





Negative news categorization and advanced filtering

High volumes of adverse media data often makes it difficult for compliance teams to prioritize risks. Systematic classification of negative news, such as financial crimes, corruption or human rights violations, along with advanced keyword filtering allow organizations to customize risk parameters based on regulatory requirements and operational exposure. Thus, improving the accuracy of screening results.

Castellum.AI in Action


Castellum.AI helps you adjust screening parameters using 100+ adverse news categories and filters. This keeps screening focused on what your organization needs.

100+ Risk Categories and Subcategories

Castellum.AI offers broad coverage across all key risk categories, with highly specific, customizable subcategories for precise risk assessment. For example:

Financial Crimes: **Fraud**
Financial Crimes: **Tax Evasion**
Financial Crimes: **Embezzlement**
Financial Crimes: **Insider Trading**

Legal Proceedings and Outcomes: **Fine**
Legal Proceedings and Outcomes: **Arrest**
Legal Proceedings and Outcomes: **Civil Law**
Legal Proceedings and Outcomes: **Extradition**





Source reliability controls

Vendors must be transparent about their source vetting criteria to ensure unreliable sources don't compromise the screening outcome. Trustworthy content, from credible outlets with editorial oversight, should only make it into the screening process. It's also essential that negative news data is both accurate and timely. Outdated or incomplete information can result in poor decision-making.

Castellum.AI in Action

Castellum.AI verifies source credibility and implements editorial filters to remove unreliable outlets. We remove duplicate coverage, defaulting to highly credible outlets, which eliminates up to 48% of redundant articles.

Data enrichment and standardization

Data enrichment and standardization are key to improving precision in adverse media screening. Enrichment enhances search results with additional identifiers—such as dates of birth or geographic location—to bring more context to the search result. Meanwhile, data standardization ensures consistent data formatting across diverse sources, which makes it easier for analysts to compare and evaluate results.

Castellum.AI in Action

Castellum.AI's enrichment process leverages Large Language Models (LLM) to identify perpetrators from victims or bystanders and surface only the most relevant results, filtering out articles where a subject is merely a victim, or a bystander, or the author of the article. Our patented data process takes it a step further and automatically standardizes unstructured media information with critical metadata (like article severity and specific adverse topics) to enable smarter filtering and more precise screening outputs.





Relevancy scoring and result deduplication

A high-performing adverse media screening solution must deliver exceptional search accuracy to surface risks that truly matter. Advanced AI-driven search capabilities are non-negotiable, as they account for variations in spelling, name formats and contextual relevance to eliminate noise and duplicates.

Castellum.AI in Action

Castellum.AI filters, taxonomizes and enriches source materials using its patented [Jgram algorithm](#), reducing false positives by 94% compared to legacy screening vendors. Relevancy scoring assesses additional information—such as name match, entity role (perpetrator vs. victim), article severity, confidence level, sentiment and subject—to intelligently score and rank results. The outcome is highly accurate alerts, minimal duplication and dramatically improved search performance.


Configurable risk thresholds

Effective screening depends on the system's ability to separate high-priority threats like financial frauds, high-risk geographies or serious convictions from low-risk mentions like unproven allegations and outdated cases. Configurable parameters allow teams to fine-tune risk based on real exposure.

Castellum.AI in Action

Castellum.AI enables precise tuning through customizable risk categories and targeted configurations. Users can easily adjust screening to reflect internal risk priorities and policies.





Access to clear context for faster alert disposition

Lengthy adverse media articles often slow down investigations and make it difficult for analysts to determine whether an alert is relevant. Without clear context, teams waste time sifting through repetitive or tangential information before they can close an alert or escalate it. Access to concise, reliable summaries helps surface the key facts needed for quick decision-making and reduce review fatigue.

Castellum.AI in Action


For each alert, Castellum.AI provides AI-generated summaries in English to instantly highlight the most critical details from each article. Analysts can understand context at a glance without reading entire stories, allowing them to resolve alerts faster and focus their time on true risks.

Complete data trail for auditability

Regulators want assurance that your adverse media processes are thorough, consistent and defensible. That means every screening decision, such as alert escalation or dismissal, needs to be backed by a clear, auditable record.

The right solution makes this easy by providing automated audit trails that track the entire decision-making process, including timestamps, risk rationale and historical data access. For instance, if a regulator questions why a corruption-related alert wasn't escalated, your screening system should quickly retrieve prior case reviews, supporting documentation and adjudication notes to eliminate manual backtracking.






Cost vs. Value: What's worth paying for

Not all adverse media screening systems are created equal. While low-cost options may seem appealing, they often rely on outdated, incomplete or unreliable datasets, limited source coverage and weak filtering mechanisms. The result? A flood of false positives, wasted analyst hours and increased risk of overlooking critical threats.

On the other hand, the right solution can save you money in multiple ways:

- **Lower false positives = Higher analyst productivity:** Advanced filtering, real-time monitoring and contextual analysis reduces the amount of false positives. More accurate results mean analysts spend less time clearing irrelevant alerts and more time on high-priority tasks.
- **Fewer regulatory risks = Reduced fines and penalties:** Comprehensive coverage, risk alignment and audit-ready reporting strengthen compliance practices and mitigate enforcement risk.
- **Scalability:** A screening solution must be able to scale as your business grows without continuously increasing headcount or operational costs. This requires flexible architecture and API integration capabilities that can handle high screening volumes and adapt to new risk factors as needed.





Questions to Ask During Vendor Evaluation

Data quality and coverage

- How do you categorize articles?
- How often is your database updated?
- How do you handle unstructured data?
- How do you treat articles in other languages?
- What adverse subjects or topics do you cover?
- How do you determine if an article is "adverse"?
- How many years of historical data do you have?
- What mechanisms are in place to filter out irrelevant or outdated media content?

System performance and customization

- What is the average response time?
- What customization options are available?
- How does the system ensure real-time monitoring and alert accuracy?

Trial, integration and compatibility

- Do you offer trial access or proof-of-concept testing?
- How easily does the tool integrate with existing compliance systems?
- Can adverse media screening be paired with other AML and risk data, like sanctions, PEPs, and OFAC 50 Rule screening?





The new standard in adverse media screening

Adverse media screening is no longer just about checking a box. It's about making quick decisions in a world where information moves faster than compliance teams can keep up. Traditional tools struggle to keep pace with the sheer volume and complexity of global news. That's where AI changes the game.

Castellum.AI embeds AI throughout the entire adverse media workflow to give compliance teams a strategic edge. From ingestion to scoring, our system is built for accuracy, speed and transparency:

- During ingestion, AI determines article relevance, identifies entities and their roles (e.g., victim vs. perpetrator vs. bystander), extracts key context, and generates concise English summaries from any source language.
- For scoring, AI classifies crime severity, flagging critical risks like money laundering while down-ranking minor issues like traffic fines, to help teams focus on what actually matters.

The result? Fewer false positives and faster reviews. Learn more at [Castellum.AI](#)

Get in Touch

For a personalized walkthrough: [Connect with the team](#)

