

Celent.



# Know Your Customer Systems: Adverse Media Screening

Technology Capabilities Matrix and 2026 XCelent Awards


Neil Katkov, PhD, Rani Kalani and Srinivasa Ram Theja Chunduri

July 2, 2026

# Executive Summary

This report evaluates the providers of systems that support adverse media screening, pKYC, and risk assessment of individual and corporate customers, as an essential component of know your customer compliance operations at banks and other financial institutions.

Celent profiles adverse media systems from 25 vendors and evaluates them according to Celent’s **Technology Capabilities Matrix (TCM)**. Our assessment of the solutions is shown in Figure 1 below. Within each of the five categories, solutions are not ranked but listed alphabetically.



Source: Celent

Celent's Technology Capabilities Matrix categories are as follows:

- **Luminary:** Excels in solution capabilities and generally has a leading market presence.
- **Technology Standout:** Excels in technology modernity, although often without the same depth of features as leading competitors. Frequently newer, these solutions have chosen a focused set of functions with which to begin their journey.
- **Functionality Standout:** Excels in functionality and likely to have a large installed base. Often more established, these solutions may have built out a robust set of features over many years.
- **Noteworthy Solution:** Potential challengers to the more established competition. They may occupy a niche place in the market, whether by targeted use case, sector-leading features, client size, or geography.
- **Developing Solution:** Typically, a vendor new to the market or an established solution that has not kept pace with the market. They may have the potential to mature into a market challenger.

Highlights of each solution are provided below to give a quick overview of the vendor landscape. For more information on each vendor, see the Vendor Profiles section of this report.

# Castellum.AI: AI-Powered Adverse Media

## Company and Product Snapshot

Castellum.AI is a private company headquartered in New York, New York. Castellum.AI’s adverse media solution is AI-Powered Adverse Media.

| Table 32: Company Snapshot: Castellum.AI |                    |
|------------------------------------------|--------------------|
| Year Founded                             | 2019               |
| Headquarters                             | New York, New York |
| Number of Employees                      | 20 - 50            |
| Revenues (USD)                           | Not disclosed      |
| Financial Structure                      | Private            |
| VendorMatch Link                         | —                  |
| Source: Vendor RFI                       |                    |

The AI-Powered Adverse Media solution was initially released in 2023 and is currently in version 2.0.

| Table 33: Product Snapshot             |                                                                           |
|----------------------------------------|---------------------------------------------------------------------------|
| Name                                   | AI-Powered Adverse Media                                                  |
| Year Originally Released/In production | 2023/2023                                                                 |
| Current Release (Date of Release)      | 2.0 (2025)                                                                |
| Revenue Derived from the Product       | Not disclosed                                                             |
| R&D Expense                            | Not Disclosed                                                             |
| Notable Clients                        | Accelaron Bank, Lead Bank, top 5 US bank, top 5 global payments processor |
| Source: Vendor RFI                     |                                                                           |

## Solution Overview and Celent Opinion

Castellum.AI provides real-time adverse media screening and monitoring, with notifications when an entity’s risk profile changes. The platform processes large volumes of negative news content from 200,000 media sources. The solution uses translation, filtering, tagging, and enrichment to produce structured outputs intended for compliance review. Risk signals are prioritized by severity and match strength.

Key capabilities include frequent data refreshes (the vendor cites updates every hour), AI-generated article summaries, and a “victim-perpetrator-bystander” model that attempts to distinguish an entity’s role in reported events to reduce irrelevant matches. Users can tune screening thresholds and filters using more than 140 risk categories to align with the institution’s risk appetite and

exposure. Castellum.AI also highlights source-quality controls as well as automated deduplication to reduce redundant articles.

The vendor reports reduced alert volumes and fewer false positives through continuous monitoring, filtering, summarization, and deduplication, resulting in lower review effort for compliance teams.

**Table 34: Celent’s View of Castellum.AI AI-Powered Adverse Media**

| Strengths and Differentiators                                                                                                                                                                                               | Caution Points                                                                    | Coollest Feature                                                                    |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"><li>• AI article summaries</li><li>• Proprietary dataset of 200,000 media sources with deduplication of redundant articles</li><li>• Tunable screening with 140 risk categories</li></ul> | <ul style="list-style-type: none"><li>• Limited reporting functionality</li></ul> | <ul style="list-style-type: none"><li>• Adverse media feed updates hourly</li></ul> |

Source: Celent analysis

## Functionality

Celent's assessment of Castellum.AI's adverse media functionality is shown in the figure below.

**Table 35: Adverse Media Features Availability**

| Function                           | Features                         | Availability |
|------------------------------------|----------------------------------|--------------|
| Adverse Media Screening & Alerting | Screening                        | ●            |
|                                    | Risk Detection                   | ●            |
|                                    | Risk Categories                  | 140          |
|                                    | Risk Scoring                     | ⊙            |
|                                    | Testing & Simulation             | ●            |
| Customer Due Diligence             | Periodic Review                  | ⊙            |
|                                    | Profile/Risk Score Updates       | ⊙            |
|                                    | Periodic Monitoring              | ●            |
|                                    | Perpetual Monitoring             | ●            |
| Analytics                          | Alert Generation                 | ⊙            |
|                                    | False Positives Reduction        | ⊙            |
|                                    | Natural Language Explanation     | ●            |
| Data and List Management           | Automated Data Downloads/Updates | ●            |
|                                    | Internal Watchlists              | ⊙            |
| Case Management                    | Workflow Support                 | ⊙            |
|                                    | Assignments/Queueing             | ●            |
|                                    | Process Automation               | ●            |
|                                    | Process Automation - GenAI       | ⊙            |
|                                    | Investigations Support           | ●            |
|                                    | Integration with AML Value Chain | ●            |
|                                    | Pre-Built Integrations           | 3            |
| Reporting                          | Risk Dashboard                   | ○            |
|                                    | Prepackaged Reports              | ○            |
|                                    | Ad Hoc Reporting                 | ●            |
|                                    | SARs Reporting                   | ●            |
|                                    | Audit Trail                      | ●            |

- |                           |                                                                            |
|---------------------------|----------------------------------------------------------------------------|
| ● Most features supported | ● Features supported with integration to separate module from vendor       |
| ⊙ Some features supported | ● Features supported with 3rd party integration                            |
| ○ Not available           | ● Features supported, mix of out of the box and available with integration |

Source: Vendor RFI

## Analytics

Analytics used for risk scoring and false positives reduction are outlined in the table below. The vendor did not disclose the average false positive rates achieved by its adverse media screening solution.

**Table 36: Support for Analytics**

|                   | Text Search/Analysis | Scoring/Decisioning | False Positives Reduction |
|-------------------|----------------------|---------------------|---------------------------|
| Pre-defined Rules | ✓                    | ✓                   | ✓                         |
| Machine Learning  | ✓                    | □                   | ✗                         |
| LLMs              | ✓                    | ✓                   | ✓                         |
| Agentic AI        | □                    | ✗                   | ✓                         |

Legend: ✓ = In production; □ = Supported but not in production; ✗ = Not supported  
Source: Vendor RFI

In addition to English, AI-Powered Adverse Media supports adverse media screening in 20 languages. Non-Latin scripts supported include Chinese - traditional, and simplified, Chinese Commercial Code, Japanese, Korean, Arabic, Cyrillic, and Thai.

## Artificial Intelligence

The solution leverages artificial intelligence as described in the table below.

**Table 37: Support for Artificial Intelligence**

|                       | Availability | Use Cases                                                                                               |
|-----------------------|--------------|---------------------------------------------------------------------------------------------------------|
| Supervised Learning   | ✓            | Standardization and enrichment of global risk data (sanctions, PEPs, adverse media) within ETL pipeline |
| Unsupervised Learning | ✗            |                                                                                                         |
| Deep Learning         | ✗            |                                                                                                         |
| NLP                   | ✗            |                                                                                                         |
| LLMs                  | ✓            | Alert review and adjudication                                                                           |
| Agentic AI            | ✓            | Alert investigation                                                                                     |

Legend: ✓ = In production; □ = Supported but not in production; ✗ = Not supported  
Source: Vendor RFI

## Data


Castellum.AI provides its own proprietary adverse media datasets, which include 200,000 media sources globally. The company provides its own proprietary watchlist data, which includes 1,600 sources.

Castellum.AI leverages adverse media data from one media aggregation provider.

## Customer Base

Castellum.AI has a total of 39 clients worldwide in production with their adverse media screening system. The breakdown by global region, industry sector, tier size, and deployment mode is shown in the figure below.

Figure 7: Client Base by Region, Sector, Size, and Deployment Option



Source: Vendor RFI

Castellum.AI has implemented its adverse media screening system in the countries shown in the table below.

Table 38: Implementations by Country

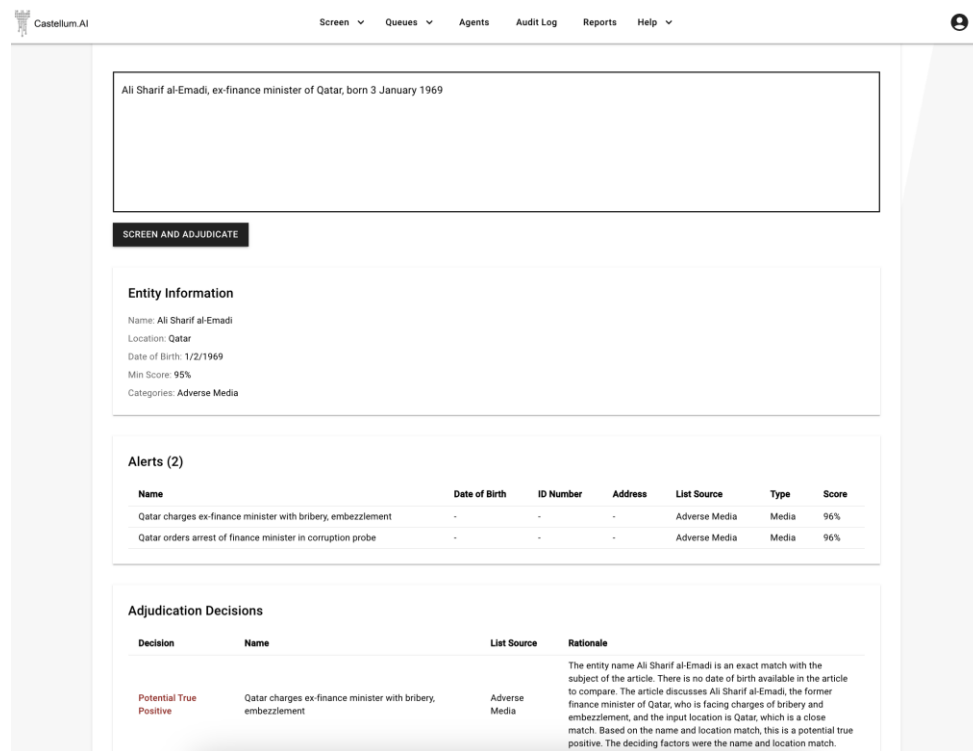
| Region          | Countries                                                       |
|-----------------|-----------------------------------------------------------------|
| North America   | United States                                                   |
| Europe          | Germany, Netherlands, Switzerland, United Kingdom, Other Europe |
| Middle East     | United Arab Emirates                                            |
| Africa          |                                                                 |
| Asia-Pacific    |                                                                 |
| Central America |                                                                 |
| South America   |                                                                 |
| Caribbean       |                                                                 |

Source: Vendor RFI

# Technology

Castellum.AI KYC Platform can be deployed on Private Cloud.

**Figure 8: Screenshot — Castellum.AI KYC Platform**



Source: Castellum.AI

Technology details for Castellum.AI AI-Powered Adverse Media are provided in the table below.

**Table 39: Technology Options**

| Technology Options  | Responses                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Code Base           | GoLang (Go): 40% ; Python: 40% ; React: 20%                                                                                                                                                                                                                                                                                                                                                                                                   |
| Databases Supported | Postgresql                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Integration Methods | API (Restful or SOAP); Web hooks; Flat files                                                                                                                                                                                                                                                                                                                                                                                                  |
| Scalability         | - 1.5 million adverse media screens/hour                                                                                                                                                                                                                                                                                                                                                                                                      |
| API Details         | <div> <div>✓</div> <div>External systems can trigger an event in the system which can be responded to by a workflow or business rules system</div> </div> <div> <div>✓</div> <div>The system allows API publishing in SOAP, REST, JSON, and XML style services as APIs</div> </div> <div> <div>✗</div> <div>API version management is available</div> </div> <div> <div>✓</div> <div>Training in extending the system is offered</div> </div> |

**Legend:** ✓ = Available; ✗ = Not available

Source: Vendor RFI

Architecture and technology features that support SaaS deployment of the solution are shown in the table below.

**Table 40: SaaS Support**

| Elements                                                             | Availability |
|----------------------------------------------------------------------|--------------|
| Containerized application (can be on-premise or cloud)               | ✓            |
| Supports a multi-tenant architecture                                 | ✓            |
| Proportion of the system architected as microservices                | 100%         |
| Proportion of the system architected as APIs                         | 100%         |
| Supports automation of development and deployment processes (DevOps) | ✓            |

Legend: ✓ = Yes ✗ = No

Source: Vendor RFI

## Partnership

Castellum.AI's implementation and fintech/regtech partners are shown in the table below.

**Table 41: Implementation and Support**

| Type of Partnership      | Partner Vendor                                                           |
|--------------------------|--------------------------------------------------------------------------|
| System Integrators       | —                                                                        |
| Fintech/Regtech Partners | Jack Henry, Hummingbird, Certa, Antuar, Yanez, LexisNexis Risk Solutions |

Source: Vendor RFI

## Implementation, Pricing, and Support

Support for implementation and service is indicated in the table below.

**Table 42: Implementation, Support, and Pricing Models**

|                                     |                                                                                                                                                                                        |
|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Typical Implementation Team Size    | 5                                                                                                                                                                                      |
| Resource Breakdown                  | Client: 80%; Vendor: 20%; Third party: 0%                                                                                                                                              |
| Average Time to Implementation      | <u>Initial Implementation:</u> Platform: 1 day, API integration: 2-4 weeks<br><u>2nd and subsequent line of business:</u> Same<br><u>2nd and subsequent states/jurisdictions:</u> Same |
| Preferred Implementation Approach   | API integration                                                                                                                                                                        |
| Pricing Models                      | Term license, Enterprise license, Subscription based license                                                                                                                           |
| Annual Maintenance Fee              | Not Applicable                                                                                                                                                                         |
| Core code/development modifications | —                                                                                                                                                                                      |

Source: Vendor RFI