

The 2026 Media Signal Map

Which outlets are training the AI agents investors use for research?

A deep-dive intelligence report for institutional investors, compliance officers, and fintech product leaders navigating the post-scraping era of agentic AI research.

Executive Summary

The era of unlicensed scraping has ended for high-stakes financial research. As of 2026, institutional investors rely on **Agentic AI Orchestration**—multi-agent systems that use specific media outlets as deterministic signal feeds. The central shift: **licensed provenance**. Agents trained on open-web data are now a compliance liability under the EU AI Act's 2026 transparency mandates.

The New Standard

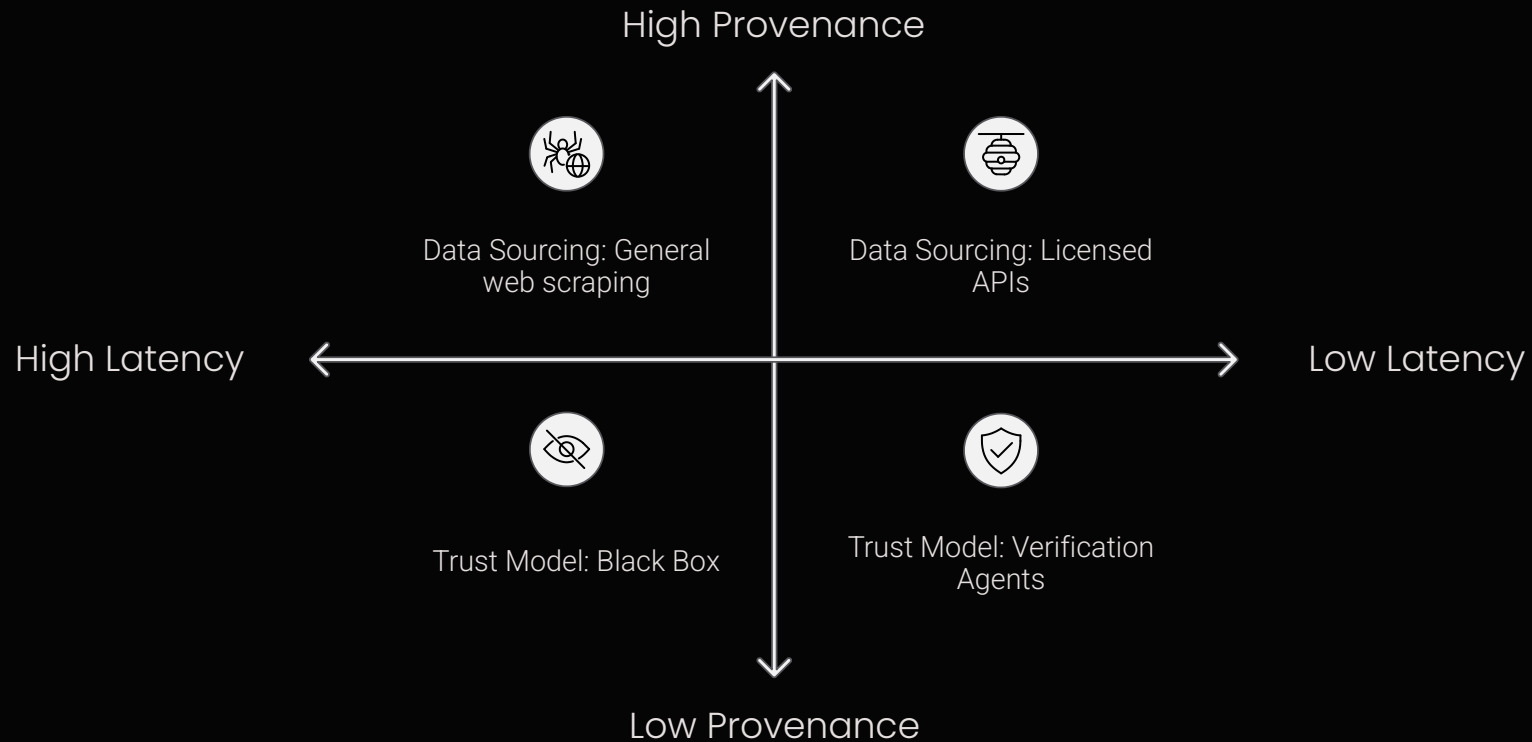
Top-tier outlets (News Corp, FT, AP) have become infrastructure providers—offering RAG-ready APIs that serve as the **"truth layer"** for financial models with 100% audit trails.

The Stakes

Total enterprise spend on AI data and services is projected to reach **\$2.5 trillion globally by end of 2026**—a 44% increase from 2025. The market for verified signal access has never been larger.

Source: **Al Jazeera / Gartner**

Legacy AI vs. 2026 Agent Requirements



The shift from open-web scraping to licensed, real-time feeds is not optional—it is now a regulatory and operational baseline for any fiduciary AI system deployed in 2026.

CHAPTER 1

Competitive Saturation

The 3 Power Hubs Controlling AI Signal Supply

Three Signal Hubs Dominate Investment Research AI

1

OpenAI / Microsoft Ecosystem

Via **News Corp & Axel Springer**.

Exclusive, multi-year rights to real-time feeds from *The Wall Street Journal*, *Barron's*, and *Business Insider*. Investors using ChatGPT or Copilot consume these privileged signals by default.

2

Financial Times (Direct-to-Enterprise)

The FT has bypassed broad aggregators, licensing its "**Gold Standard**" dataset directly to private bank LLMs—JPMorgan, Goldman Sachs—via proprietary RAG interfaces with structured Knowledge Graph access.

3

Microsoft Publisher Content Marketplace (PCM)

A 2026 industry-standard platform aggregating **500+ publishers**, including *The Washington Post* and *The Associated Press*, with built-in royalty tracking for AI training token purchases.

The OpenAI / News Corp Axis: Default Signal for Consumer Research

The \$250M Master Deal

Through their multi-year, **\$250 million+ partnership**, OpenAI agents have direct, low-latency access to the archives and real-time feeds of *The Wall Street Journal* and *Barron's*. This is the dominant "default" signal layer for consumer-grade investment research tools deployed at scale.

For institutional investors relying on ChatGPT or Microsoft Copilot, the underlying news intelligence is sourced almost exclusively through this privileged channel—making awareness of this dependency a governance imperative.

Source: ***Mediaweek***

News Corp Archives

WSJ, Barron's — macro & real estate signals

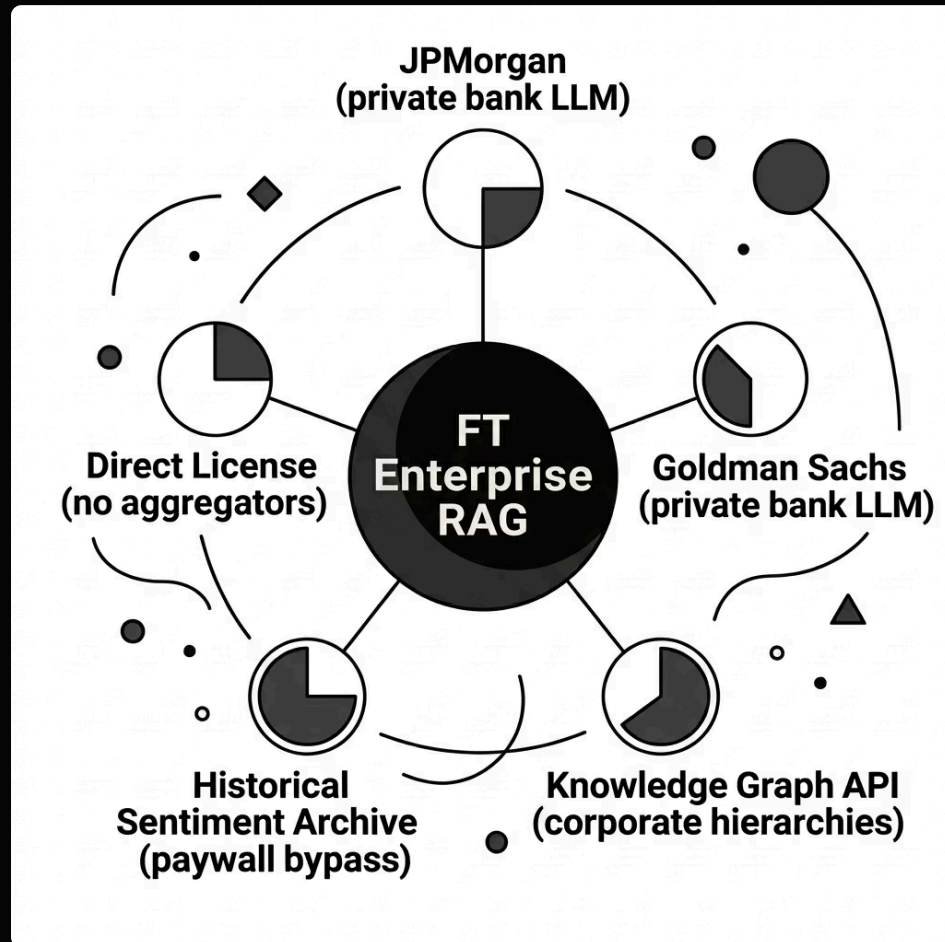
Axel Springer

Business Insider — consumer & tech sentiment

Meta Add-On

\$50M/yr for US & UK news archives via social-agent research

The Financial Times: RAG-Native Knowledge Infrastructure



Beyond Text — Structured Intelligence

The FT's 2026 offering goes beyond raw article feeds. Their API provides **Knowledge Graphs** enabling agents to navigate complex corporate hierarchies and historical sentiment shifts without model re-training.

Enterprise cost: \$100,000–\$500,000+ per institution annually. Designed exclusively for deep due diligence workloads where paywall bypass and structured provenance are non-negotiable.

- ❏ The FT has deliberately avoided broad aggregators, preserving premium positioning and margin control by licensing directly to Tier-1 financial institutions.

Meta's \$50M Bet: Social-Agent News Intelligence

As of **March 2026**, Meta pays up to **\$50 million annually** to News Corp for access to its data, specifically to power social-agent research capabilities focused on US and UK news archives. This represents a third demand vector—alongside OpenAI and direct enterprise licensing—for the same underlying News Corp signal pool.

\$250M+

OpenAI / News Corp

Multi-year master deal for WSJ & Barron's
real-time access

\$50M

Meta / News Corp

Annual payment for social-agent US & UK
archive access

500+

PCM Publishers

Aggregated in Microsoft's Content
Marketplace with royalty tracking

Sources: **Mediaweek (Meta deal)** · **Digiday AI Licensing Timeline**

CHAPTER 2

Access Tiers & Pricing

What Institutional AI Access Actually Costs in 2026

2026 Media Access Tiers for AI Agents

Media Provider	Access Model (2026)	Est. Annual Enterprise Cost	Primary Use Case
News Corp	Exclusive Microsoft / OpenAI API	Included in \$250M+ Master Deal	Macro-sentiment & real estate signals
Financial Times	Proprietary "Enterprise RAG"	\$100k – \$500k+ per institution	Deep due diligence & paywall bypass
Reuters / LSEG	AI Mirror / Workspace API	\$25,000 – \$120,000 per agent	Real-time ticker-linked verification
Associated Press	Content Marketplace (PCM)	\$1,000 – \$10,000/mo (usage-based)	Ground-truth verification & local signals

Consumption-Based Pricing: The 2026 Standard

From Flat-Rate to Usage-Based

2026 marks the decisive shift away from flat annual licenses toward **per-conversation, per-query pricing**. Most media APIs now charge **\$0.01 to \$0.50 per agent conversation**, scaled to the depth of archive accessed.

For multi-agent orchestration systems running thousands of verification queries daily, cost management at the data-ingestion layer has become a primary engineering and finance concern.

Source: [**SearchUnify, 2026**](#)

Low-Depth Query

\$0.01–\$0.05 per conversation. Recent headlines, ticker-linked news, real-time AP ground-truth checks.

Mid-Depth Query

\$0.05–\$0.20 per conversation. Historical archive traversal, FT Knowledge Graph navigation.

High-Depth Query

\$0.20–\$0.50 per conversation. Full due diligence pipelines with multi-source cross-verification.

CHAPTER 3

The Compliance Wall

Regulatory Mandates Reshaping the AI Data Supply Chain

Four Regulatory Regimes Now Govern AI Signal Use




SEC / FINRA Care Obligation

**Accuracy of
datasets,
Reasonable Basis
documentation**

EU AI OFFICE Article 53(1)(d) GPAI

**Mandatory public
summary of
copyrighted training
data for GPAI models**



TEXAS TRAIGA Section 5 AI Disclosures

**All AI agents used
in high-stakes
financial decisions**

ASIC AUSTRALIA 2026 Key Issues

**Agentic AI risks and
autonomous
decisions without
audit trails**



Each jurisdiction targets a distinct failure mode: the SEC focuses on data accuracy, the EU on copyright transparency, Texas on disclosure, and ASIC on autonomous decision-making without human oversight.

EU AI Act: The Global Benchmark for Provenance Compliance

February 2026 Enforcement

AI providers must now publish a **public summary of all copyrighted training data** used in general-purpose AI models. Using "shadow data" for financial agents carries fines of up to **€15 million or 3% of global revenue**—whichever is higher.

Source: *European Union Official*

What "Transparency Phase" Means in Practice

- Every AI agent deployed for financial research must have a documented training data summary on file
- Robots.txt and ai.txt protocol compliance is now legally required—not optional
- Synthetic or AI-augmented data must be disclosed separately
- Firms licensing data through aggregators must verify the aggregator's own compliance status

SEC / FINRA: Explainable AI Is Now a Fiduciary Requirement

The **SEC's 2026 Oversight Report** specifically flags "reliance on limited or outdated datasets" as a breach of fiduciary duty. If an agent recommends a trade based on unverified media signals, the firm is liable for "**Data Integrity Failure**" under Rule 3110 and Care Obligation standards.

1 Reasonable Basis Documentation

Every AI-assisted trade recommendation must include a documented "reasonable basis"—traceable to a licensed, timestamped source. Black-box inference chains are no longer defensible.

2 Explainable AI for Fiduciary Decisions

Regulatory bodies now require interpretable outputs for any AI agent that influences portfolio allocation. Verification agents must cross-check and log their signal sources in real time.

3 Shadow AI: #1 Compliance Risk

As of March 2026, using unverified, scraped tools for research is the top-cited compliance risk in SEC and FINRA examinations of investment advisers.

Sources: **FINRA 2026 Regulatory Oversight Report** · **Plante Moran / SEC 2026 Priorities**

TRAIGA: Texas Sets the De Facto U.S. State Standard

Texas Responsible AI Governance Act

Effective **January 1, 2026**, TRAIGA's Section 5 requires disclosure of all AI agents used in high-stakes financial decisions—including research, modeling, and trade recommendations.

Because Texas is home to a significant concentration of financial services firms, compliance with TRAIGA has effectively become a **de facto national standard** for provenance logs in research reports.

Source: **Baker Donelson**

1

Disclose Agent Identity

Name every AI model used in a financial decision workflow

2

Log Data Sources

Document the provenance of every signal input to the agent

3

Human Sign-Off Required

High-stakes outputs must include a human-in-the-loop validation record

ASIC 2026: Agentic AI in Focus Down Under

"Autonomous decisions made by AI agents without sufficient audit trails represent a systemic risk to market integrity. Firms must demonstrate that human accountability is embedded in every agentic workflow."

– *ASIC 2026 Key Issues & Outlook*

Australia's Securities and Investments Commission has made **agentic AI governance** a primary 2026 examination focus. The regulator is specifically targeting the absence of audit trails in autonomous research and trading decisions—pushing Australian institutions toward the same licensed provenance standards emerging globally.

Source: **ASIC 2026 Outlook**

CHAPTER 4

GTM Intelligence

Entry Points & Strategic Opportunities for 2026

The Agentic Orchestration Gap

Most firms are actively migrating away from single-agent research. The dominant emerging architecture uses a **Lead Coordinator Agent** that manages specialized downstream agents—each trained on a different licensed signal source.



This architecture creates a significant market opportunity for tools that act as the orchestration layer—managing signal routing, source verification, and compliance logging across multiple licensed feeds simultaneously.

Source: **Deloitte TMT 2026**

The Verification Layer Opportunity

Bloomberg Signal + AP Ground-Truth

A critical gap exists for agents that act as "**Verification Layers**"—checking a Bloomberg signal against an AP ground-truth report before surfacing a recommendation. No single incumbent currently owns this cross-source validation role at scale.

The winning product in this space will need to integrate with **at least two licensed signal sources**—one for speed (e.g., Reuters real-time), one for accuracy (e.g., AP's PCM feed)—and produce a machine-readable provenance certificate with every output.

Speed Layer

Reuters / LSEG AI Mirror:
real-time ticker-linked news
at \$25k–\$120k/agent/year

Accuracy Layer

AP Content Marketplace
(PCM): usage-based
ground-truth at \$1k–
\$10k/month

Provenance Certificate

Machine-readable audit trail with unique hash, timestamp,
and source URL

Regional Alpha: The Data Gap the Big 3 Miss

The Mid-Market Signal Void

While Tier-1 outlets are fully licensed and saturated with demand, there is a structural **data gap in regional and mid-market news**. The "Big 3" models (OpenAI, FT, PCM) are optimized for large-cap, English-language, US/EU market signals.

Agents tapping into the **News/Media Alliance & Prorata consortium** can access unique local signals that no incumbent model currently surfaces at institutional quality.

Where the Alpha Lives


- **Small-cap earnings coverage** from regional business journals not in the PCM
- **Municipal bond signals** from state-level news outlets
- **Non-English language feeds** for EM-focused funds underserved by current licensed datasets
- **Prorata consortium pricing**: royalty-tracked micro-licensing for long-tail publishers

CHAPTER 5

Data Architecture

The Bronze–Silver–Gold Lakehouse Standard

How Verified Data Reaches a Portfolio Manager's Dashboard



The Bronze-Silver-Gold Lakehouse has emerged as the 2026 standard architecture. Raw licensed data is progressively cleaned, enriched, and verified before it ever influences a portfolio decision—each layer producing its own audit metadata.

Source: [**OvalEdge Data Lineage Best Practices**](#)

Bronze Layer: Raw Ingestion from Licensed Sources



RAG (Real-Time) Feeds

Data pulled at query-time via licensed API. Zero-latency access to live news as published. Includes News Corp, Reuters, AP, FT enterprise endpoints.



Static / Fine-Tuned Weights

Data embedded in model training weights from 2025–2026 training runs. Must be documented in the EU AI Act public training data summary.



Copyright Opt-Out Verification

Every source must confirm crawler compliance with robots.txt and ai.txt protocols—now legally required under the EU AI Act for all financial AI systems.

Silver Layer: Processing & Transformation

The Cleaning Logic

At the Silver Layer, raw signals are filtered and structured before verification. The key transformation decisions made here define the quality and defensibility of the final output.

- **Exclusion filters:** Remove social media sentiment, unverified blogs, non-SEC-filing news
- **Synthetic data disclosure:** Flag any signal generated or augmented by another AI model
- **Bias mitigation metadata:** Document whether the dataset was balanced for regional/small-cap vs. large-cap coverage

📄 **2026 Transparency Requirement**

Any synthetic or AI-augmented signal introduced at the Silver Layer must be separately disclosed in the agent's provenance log. Failure to disclose constitutes a violation under EU AI Act Article 53(1) (d) and TRAIGA Section 5.

The Silver Layer is where most compliance failures originate—primarily through undisclosed synthetic augmentation or inadequate social media exclusion filters.

Gold Layer: Verification & Governance

Human-in-the-Loop (HITL)

A human analyst must verify and sign off on any signal that directly influences a portfolio allocation. Required for FINRA Rule 3110 compliance.



Deterministic Source Link

Direct URL and timestamp to the original media source must be stored alongside every inference. No inference without a traceable origin.



Audit Trail ID

Unique cryptographic hash for each inference, stored for a minimum of **5 years** per 2026 FINRA compliance standards. The non-negotiable backbone of "Explainable AI."

CHAPTER 6

The Audit Template

2026 AI Data Lineage Audit Framework for Institutional Research

AI Data Lineage Audit: Section 1 – Provenance & Acquisition

Target Use Case: Institutional Investment Research & Fiduciary Reporting

Field	Required Entry	Compliance Reference
Source Identifier	E.g., "News Corp RAG API" or "FT Enterprise Feed"	EU AI Act Art. 53 / TRAIGA Sec. 5
Licensing Type	Direct Enterprise License vs. Marketplace Token (PCM)	SEC Rule 3110
Collection Method	RAG (real-time) or Static (fine-tuned weights)	FINRA 2026 Oversight Report
Copyright Opt-Out Status	Confirm robots.txt / ai.txt protocol compliance	EU AI Act Transparency Phase

AI Data Lineage Audit: Section 2 — Processing & Transformation

Field	Required Entry	Compliance Reference
Cleaning Logic	List filters applied (e.g., "Exclude social media," "SEC-verified only")	SEC Data Integrity / FINRA
Synthetic Data Disclosure	Was any signal generated or augmented by another AI model?	EU AI Act Art. 53(1)(d)
Bias Mitigation Metadata	Was dataset balanced for regional/small-cap vs. large-cap coverage?	TRAIGA Sec. 5

- ❏ Synthetic data augmentation is the most commonly missed disclosure field in 2026 audit reviews. Even partial AI-assisted summarization of a licensed feed constitutes synthetic augmentation under current EU guidance.

AI Data Lineage Audit: Section 3 – Verification & Governance

Field	Required Entry	Compliance Reference
Human-in-the-Loop (HITL)	Did a human analyst verify and sign off on this signal?	FINRA Rule 3110 / ASIC 2026
Deterministic Source Link	Direct URL and timestamp of the original media source	SEC Examination Priorities 2026
Audit Trail ID	Unique hash for this inference; store for minimum 5 years	FINRA 2026 / TRAIGA Sec. 5
Verification Agent Log	Record of any secondary agent that checked the primary signal	EU AI Act / SEC Explainability

Full Audit Requirements by Regulatory Body

Regulatory Body	Key 2026 Requirement	Audit Focus
SEC / FINRA	Rule 3110 / Care Obligation	Accuracy of datasets; documentation of "Reasonable Basis" for trades
EU AI Office	Article 53(1)(d)	Mandatory public summary of copyrighted training data (GPAI models)
Texas (TRAIGA)	Section 5 – AI Disclosures	Disclosure of all AI agents used in high-stakes financial decisions
ASIC (Australia)	2026 Key Issues	"Agentic AI" risks; autonomous decisions without audit trails

Sources: **FINRA** · **EU Commission** · **Baker Donelson** · **ASIC**

CHAPTER 7

Strategic Synthesis

What This Means for Investors, Compliance Officers & Fintech Leaders

For Institutional Investors: Know Your Signal Stack

→ Audit Your Research Tool's Data Provenance

If your firm uses ChatGPT, Copilot, or any LLM-based research tool, determine which media licenses power its underlying signals. The News Corp / OpenAI \$250M deal means WSJ and Barron's are the dominant default—but their geographic and asset-class biases should be understood and disclosed.

→ Demand Provenance Certificates from Vendors

Before procuring any AI research tool in 2026, require a written data lineage disclosure covering: source identifiers, licensing type, collection methodology, and copyright compliance status. This is no longer a "nice to have"—it is a fiduciary requirement under both SEC and EU AI Act standards.

→ Regional Alpha Requires Supplemental Feeds

No single licensed Tier-1 outlet covers small-cap, mid-market, or non-English signals at institutional quality. Firms seeking regional alpha must actively source supplemental data through the Prorata / News Media Alliance consortium or equivalent.

For Compliance Officers: The Shadow AI Exposure Checklist

Shadow AI: The #1 Risk Category

As of March 2026, the SEC's Examination Priorities report explicitly identifies "reliance on unverified, scraped AI tools" as a primary breach vector for fiduciary duty. Any firm whose analysts use consumer-grade AI tools for research without a formal provenance log is exposed.

Immediate Compliance Actions

1. Inventory all AI research tools currently in use across the investment team
2. Verify each tool's data provenance against the 2026 Audit Template fields
3. Establish a mandatory HITL sign-off protocol for all AI-assisted recommendations
4. Implement 5-year audit trail storage for all inference hashes
5. Confirm TRAIGA compliance if any agent is used for US financial modeling

For Fintech Leaders: Where to Build in 2026

Orchestration Layer

Build the **Lead Coordinator** that routes queries across licensed Sentiment, Fact-Check, and Verification agents. No incumbent owns this layer.

Verification-as-a-Service

Cross-source validation between Bloomberg speed signals and AP ground-truth—packaged as a compliance-ready API with machine-readable provenance certificates.


Regional Signal Aggregation

Aggregate and normalize Prorata/News Media Alliance feeds for institutional use. Fill the gap the Big 3 models structurally cannot address.


Audit Trail Infrastructure

5-year compliant storage and retrieval of inference hashes, source links, and HITL records. A foundational compliance product with recurring enterprise revenue.

(1) DATA INGESTION



(2) PROCESSING



(3) VERIFICATION



(4) OUTPUT



CHAPTER 8

The Signal Map

Outlet-by-Outlet Intelligence: Who Powers What

The Wall Street Journal & Barron's

Signal Role: Macro-Sentiment & Real Estate

The WSJ and Barron's are the **default signal layer** for the largest deployed AI research infrastructure in the world—the OpenAI / Microsoft ecosystem. Any institutional investor using ChatGPT Enterprise or Microsoft Copilot for research is consuming WSJ and Barron's signals as their primary macro-intelligence feed, whether they know it or not.

The \$250M News Corp master deal ensures this will remain the case through at least 2027, with Axel Springer's Business Insider providing supplemental tech and consumer sentiment coverage.

Access Model

Exclusive Microsoft / OpenAI API

Cost

Bundled in \$250M+ master deal

Primary Signal

Macro-sentiment, real estate, earnings

Coverage Gap

Limited regional / small-cap

The Financial Times

Access Model

Proprietary Enterprise RAG
(no aggregators)

Cost

\$100,000–\$500,000+ per
institution/year

Primary Signal

Deep due diligence, corporate
hierarchies, paywall-depth
historical sentiment

Key Differentiator

Knowledge Graphs — agents
navigate structure, not just
text

Signal Role: Deep Due Diligence & Knowledge Structure

The FT's deliberate avoidance of broad aggregators has preserved its positioning as the **premium, Tier-1 signal** for the most demanding institutional use cases. Its Knowledge Graph API—unique among major publishers—allows agents to traverse corporate hierarchies and historical sentiment sequences without model re-training.

Current enterprise clients include JPMorgan and Goldman Sachs. The FT's high price point (\$100k–\$500k+) is a feature, not a bug—it signals exclusivity and audit-grade reliability to compliance teams.

Reuters / LSEG: Real-Time Ticker-Linked Verification

Signal Role: Speed Layer for Market Events

Reuters, now deeply integrated with LSEG's financial data infrastructure, occupies the **speed layer** of the 2026 signal stack. Its AI Mirror / Workspace API provides real-time, ticker-linked news that agents can use to verify market-moving events against live price action.

At \$25,000–\$120,000 per agent per year, Reuters is the primary choice for high-frequency research agents that need to correlate news events with real-time market data without introducing latency.

Key Capability

Ticker-linked news verification: Every Reuters article is structurally tagged with the relevant equity, commodity, or currency ticker—enabling agents to cross-reference news events with live price data in a single API call.

Cost range: \$25,000–\$120,000 per agent/year, depending on archive depth and real-time feed volume.

Associated Press: The Ground-Truth Standard

Access Model

Content Marketplace (PCM) – usage-based

Cost

\$1,000–\$10,000/month

Signal Type

Neutral, wire-standard facts with strong regional and local coverage

Signal Role: Fact-Check Agent's Primary Feed

The AP is the **designated ground-truth layer** for verification agents in 2026's multi-agent architectures. Its neutrality, geographic breadth, and relatively low cost-per-query make it the natural anchor for Fact-Check Agents in the Deloitte-cited orchestration model.

AP's inclusion in the Microsoft PCM also means it benefits from the full royalty-tracking and provenance infrastructure of that platform—making it the most compliance-ready option for firms that need defensible audit trails at scale.

Washington Post & The PCM Long Tail

Beyond the Tier-1 outlets, the **Microsoft Publisher Content Marketplace (PCM)** aggregates 500+ publishers—including *The Washington Post*—under a single royalty-tracked licensing framework. This creates the infrastructure for AI agents to access a far broader signal universe than any single bilateral deal can provide.




PCM Publishers

Aggregated under a single royalty-tracking framework for AI training token purchases



EU Revenue Fine Cap

Maximum penalty for using unlicensed "shadow data" under the EU AI Act's 2026 enforcement phase



YoY Spend Growth

Projected increase in enterprise AI data & services spend from 2025 to 2026 (Gartner / Al Jazeera)

The 2026 Media Signal Map: Five Things to Know

01

Provenance Is Now Infrastructure

Licensed, auditable media feeds are no longer a premium add-on—they are the baseline for any AI system used in fiduciary research. Open-web scraping is a compliance liability, not a cost-saving measure.

02

Three Hubs Control the Supply Chain

OpenAI/News Corp, the Financial Times, and the Microsoft PCM are the dominant gatekeepers. Knowing which hub powers your research tool defines both your signal bias and your regulatory exposure.

03

Multi-Agent Orchestration Is the New Standard

Single-agent research is being replaced by Lead Coordinator architectures that route queries across licensed Sentiment, Fact-Check, and Verification agents in parallel.

04

Shadow AI Is the #1 Exam Finding

The SEC, FINRA, EU AI Office, and TRAIGA are all targeting the same failure mode: AI agents used in financial decisions without documented data provenance. The enforcement era is here.

05

Regional Alpha Remains Uncaptured

The Tier-1 licensed universe has a structural gap in regional, mid-market, and non-English signals. Firms and products that close this gap via the Prorata/NMA consortium will access differentiated intelligence.

Full Citation Index

Licensing & Deal Intelligence

[Timeline of Major 2025–2026 AI Licensing Deals – Digiday News Corp & OpenAI Multi-Year Partnership – Mediaweek](#)
[Meta–News Corp \\$50M AI Content Deal – Mediaweek](#)

Regulatory & Compliance

[EU AI Act: Transparency & Copyright Enforcement – European Union](#)
[EU Commission Training Data Summary Template – European Commission](#)
[FINRA 2026 Regulatory Oversight Report – FINRA](#)
[SEC 2026 Examination Priorities – Plante Moran](#)
[2026 AI Legal Forecast: TRAIGA & Compliance – Baker Donelson](#)
[ASIC 2026 Outlook: Agentic AI – ASIC](#)
[PwC 2026 Business Predictions: Agentic AI – PwC](#)

Market & Spend Data

[AI Spending to Total \\$2.5 Trillion in 2026 – Al Jazeera / Gartner](#)
[5 AI Trends for 2026 Every CFO Must Know – Lucanet](#)
[OpenAI 2026 Enterprise Focus & Pricing – Constellation Research](#)
[AI Agent Cost & Infrastructure Breakdown 2026 – SearchUnify](#)

Architecture & GTM

[Deloitte TMT Predictions 2026 – Deloitte](#)
[Data Lineage Best Practices for 2026 – OvalEdge](#)
[EU AI Act 2026 Guidelines – artificialintelligenceact.eu](#)