pwc

From fragmented monitoring to risk intelligence:

How unified platforms and AI agents can redefine surveillance

When you receive healthcare from multiple specialists who don't share records, it increases the risk of the true diagnosis being missed.

Today, that's how many financial institutions treat market abuse surveillance. It is common to have trade and communications surveillance, archiving, transcription, analytics, and case management operate across separate platforms with distinct data models, workflows, and governance controls. This requires investigators to reconstruct context separately, duplicating effort across tools while lacking a unified view of market abuse risk, resulting in unbalanced cost ratios, high rates of false positives, and avoidable regulatory exposure.

A modern surveillance function, by contrast, requires integration: unified data ingestion, recording, retention, retrieval, investigative analytics, and case review operating within a coordinated control framework. Without this, surveillance remains constrained to isolated alert processing rather than contextual risk detection.

As data volumes expand, communication channels proliferate, and market abuse techniques become more coordinated and technologically sophisticated, institutions are confronting a structural inflection point. Surveillance can no longer operate as a layered collection of specialized tools – it needs to evolve into an integrated control architecture capable of scaling with complexity.

However, achieving that evolution is neither simple nor immediate. Legacy systems, vendor dependencies, and governance silos can make structural integration difficult. However, two structural forces are beginning to make this transition more attainable:

Platform convergence	Strategic investment, private equity activity, and targeted acquisitions could yield providers capable of offering integrated, end-to-end surveillance platforms.
AI-led automation	In parallel, agentic AI systems are emerging as the first line of investigative triage. These systems can increasingly perform Level 1 alert disposition, assemble cross-system context, wade through false positives, and apply consistent decision logic at scale.

Together, these forces can create the potential to move surveillance from a fragmented compliance obligation to a strategic enterprise risk capability: integrated, predictive, and as essential to the institution's safety as coordinated care is to a patient's recovery.

The current market landscape: Two parallel realities, one fragmented environment

To understand why surveillance remains fragmented, it is important to examine how its core components evolved. Trade surveillance and communications surveillance[1] were built under different regulatory drivers and technological constraints. Trade surveillance matured around market abuse typologies (e.g., spoofing, layering, wash trades, front-running) and consumes high-volume, structured data from order management systems, execution venues, and market tapes. Communications surveillance, by contrast, evolved from supervision and recordkeeping mandates and ingests semi-structured and unstructured data: email, chat, voice, mobile messaging, and collaboration platforms. Trade surveillance engines run on millisecond-level signals and price/volume movements while communications tools parse language, context, and sentiment over hours or days. Each could have its own vendors, data models, storage footprints, and end-user communities.

As a result, it is common for financial institutions to rely on multiple vendors depending on the risk scenarios or products being surveilled. No single provider dominates both domains across the surveillance lifecycle at global-bank scale, so banks shoulder the burden of interoperability, identity resolution, and case management across systems. According to a survey of trade and market surveillance executives, 56% of industry respondents face challenges in their technology infrastructure, working with fragmented, minimally standardized systems.

The impact of this fragmentation becomes visible in several recurring pain points:

Alert fatigue

Large financial institutions often deal with thousands of surveillance alerts each day. The 2023 PwC (UK) Market Abuse Surveillance Survey found that fewer than 0.02% of the approximately one million alerts generated annually by the average firm, are progressed beyond Level 1 review. While this does not mean the remaining alerts lack merit, it underscores the structural imbalance between alert volume and actionable risk.

Manual reconciliation

Disconnected systems can mean context gets rebuilt manually. A simple example: a trader's unusual order placement triggers a trade alert. Separately, a spike in messages using evasive language triggers a communications alert. Without automated correlation on the person, desk, instrument, and timestamp, these two seemingly unrelated alerts may remain a low-priority instead of being flagged as a high-priority case.

Even when alerts are escalated, investigators still need to reconstruct context across multiple systems, often through manual effort. This can include:

- Searching for related alerts across tools and queues;
- Reconstructing conversation and trade timelines;
- Reconciling identity mismatches (nicknames, multiple emails/handles, shared lines, ring-downs); and
- Writing duplicative narratives to satisfy different audit trails.

These manual processes extend time-to-disposition, introduce inconsistency in investigative quality, and increase the probability that coordinated misconduct patterns remain undetected.

[1] This paper focuses on trade and communications surveillance for market abuse and conduct risk, opposed to financial crimes. While many of the same challenges and innovations could be applied, there are typically different vendor offerings for the latter.

Higher costs

The bill for fragmentation shows up across different areas:

Run-the-Bank	Overlapping licenses; dual ingestion pipelines; duplicative storage/egress for the same content (e.g., voice stored in one system for archive and re-stored/transcoded for surveillance); multiple case tools; parallel training and access management.
Change-the-Bank	Each policy update, channel onboarding, or regulatory enhancement ripples through several vendors; testing/validation needs to be repeated; integration work gets redone.
Hidden costs	Transcription/translation computing, data sovereignty workarounds, and model risk governance often sit on separate budgets but are in fact part of the total surveillance cost stack.

Compliance blind spots

Parallel systems create genuine coverage gaps, not just inefficiencies:

1	**Cross-product monitoring**	This capability is still not significantly present within the industry. Firms often struggle to normalize heterogeneous data streams, which can impair the ability to detect multi-leg or cross-venue manipulation.
2	**Off-channel communications**	If mobile and collaboration tools are not captured and supervised consistently (or if usage policies are not enforced), firms can face recordkeeping and supervision exposure.
3	**Voice reality gap**	Voice recordings are often retained but not fully surveilled due to transcription reliability, language coverage, or cost constraints, leaving a high-risk channel under-monitored.

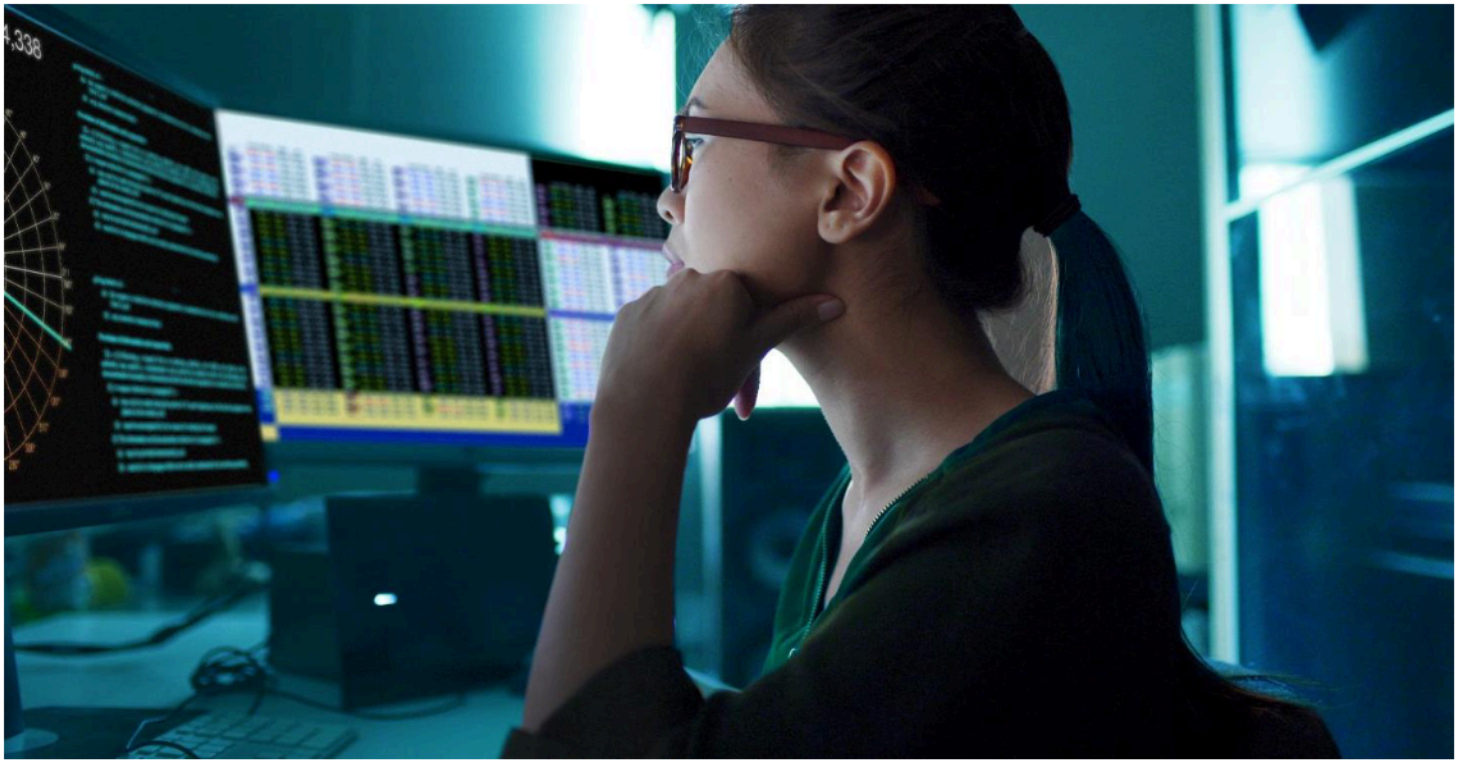

The future market landscape: AI integration and ecosystem evolution

Unified surveillance has been an industry aspiration for more than a decade. Periodically, some institutions have explored integration initiatives, often prompted by regulatory events, technology refresh cycles, or cost pressures. Many of these efforts delivered incremental improvements but fell short of broad architectural convergence, constrained by legacy infrastructure, vendor limitations, data interoperability challenges, and debatable ROI when factoring in time and investment.

What distinguishes the current moment is not simply renewed interest, but a shift in feasibility. Advances in platform architecture, cloud scalability, and AI-enabled analytics are helping to reduce barriers that previously made integration impractical at scale.

In addition, the strategic rationale for convergence has expanded. Integration is no longer viewed solely as a cost-efficiency initiative. It is increasingly framed as a prerequisite for detecting cross-channel misconduct, improving investigative consistency, and generating enterprise-wide behavioral insights.

The business case for convergence and AI integration rests on four structural advantages:

1. **Risk detection**: Connecting trade data and comms patterns to catch insider trading and collusion allows firms to detect potential risks *before* they occur.
2. **Reducing false positives**: Integrated systems and AI applications assist with reducing false positives, so compliance can focus on true issues. Contextual analysis prevents wasted efforts. Several vendors tout next-gen AI technologies as being able to reduce false positives by over 90% when there's been a move from rules-based lexicons to AI-driven contextual analysis.
3. **Operational efficiency**: Consolidated systems can create simplified workflows, and lower IT complexity.
4. **Data insights beyond compliance**: Unified surveillance empowered with AI allows firms to unlock behavioral analytics and proactive risk management.

The strategic case for convergence is increasingly clear. The more difficult question has historically been feasibility. Two developments could materially shift that reality: evolutions across the surveillance vendor ecosystem, and the rapid maturation of AI-enabled workflows.

Ecosystem evolution: capital, consolidation, and platform expansion

The push toward unified surveillance is helping reshape the vendor landscape. Surveillance compliance is somewhat unique when it comes to financial institutions working with third party providers, with most of the industry opting for "buy" over "build" technology. While this model has contributed to a fragmented ecosystem of specialized solutions, it also creates opportunity within the market. Fragmentation leaves space for capital deployment, consolidation, capability expansion, and the emergence of more integrated, enterprise-scale platforms.

Recent dealmaking in the market abuse technology sector reflects both consolidation among established providers and increased investment in next-generation platforms. Importantly, these two dynamics are themselves beginning to converge, with innovators steadily demonstrating greater enterprise readiness and trustworthiness, while incumbents are expanding their portfolios with more advanced analytical and AI-enabled capabilities. Together, these dynamics are shaping a more competitive ecosystem.

There are several other dynamics driving investment across the surveillance technology sector as they test the limits of institutions' reliance on legacy infrastructure: 1) sustained regulatory expectations around monitoring communications and trading activity and 2) emerging asset classes – including digital assets and prediction markets –introducing new surveillance considerations as they intersect with traditional financial services.

The rise of AI agents

AI has the potential to reshape the future of the surveillance industry for financial institutions through its ability to analyze vast amounts of data and identify trends that humans may miss, enabling faster, data-driven decisions. Rather than replacing human oversight, AI is being deployed to augment decision-making, reduce subjectivity in alert handling, and surface behavioral patterns that rule-based approaches often miss.

A notable development is the emergence of agentic AI – systems designed not merely to generate insights, but to perform defined investigative tasks within surveillance workflows. These systems can operate as structured digital assistants, executing repeatable processes while escalating higher-risk scenarios for human review.

Agentic AI is more than just a time-saver, it can offer benefits such as:

Standardizing and streamlining Level 1 alert disposition	Agents can apply consistent investigative logic and decision criteria across high-volume, low-risk alerts, reducing reviewer subjectivity and variability while maintaining regulatory-defensible outcomes. By learning from historical dispositions and escalation patterns, agents can pre-classify alerts with confidence scores, enabling straight-through disposition or targeted human review.
Reducing time-to-disposition	By autonomously assembling relevant context, such as related communications, trade activity, historical behavior, and prior cases, agents can significantly compress investigation cycles and eliminate manual data gathering. This allows institutions to resolve alerts faster, reduce alert backlogs, and reallocate human effort toward complex, higher-risk investigations.
Automated case management	Agents can gather contextual data from systems to categorize alerts and develop risk reports, so humans can focus on key decision-making rather than data cleansing and compilation.

Further, AI agents aren't just theoretical; PwC supported a global technology company to transform its risk and control environment by enabling a unified risk management framework empowered by AI agents. PwC deployed agents to strengthen control descriptions, compare against enterprise standards and business operations, and flag ambiguities for remediation, cutting documentation errors in half, boosting regulatory alignment by 25%, and reducing manual effort by ~60%.

AI-led automation does not eliminate the architectural and data challenges facing surveillance functions. However, when deployed within integrated platforms, it can materially improve signal detection, investigative speed, and documentation consistency – reinforcing the transition from fragmented monitoring to intelligence-driven control.

What should firms do now?

A fully unified surveillance model is not yet a universal reality. Architectural fragmentation, vendor dependencies, and operating model silos remain embedded across many institutions. Yet the absence of a perfect end-state does not preclude progress. The structural forces reshaping the surveillance landscape – platform convergence and AI-enabled workflows – can create practical entry points for institutions willing to act deliberately.

The objective is not to pursue integration for its own sake, nor to replace one tool with another. It is to begin redesigning surveillance as a coordinated control architecture capable of generating cross-domain risk intelligence. Institutions that take measured, architecture-first steps today can be better positioned to scale as the ecosystem continues to mature:

1. Define a target surveillance architecture	Articulate a clear control model that integrates trade, communications, archiving, analytics, and case management. Identify structural seams, reduce duplicative workflows, and prioritize integration over incremental tool enhancement.
2. Reassess the vendor ecosystem	Evaluate current providers against the target architecture, focusing on interoperability, data portability, workflow integration, and governance transparency. Consider where consolidation, capability expansion, or platform convergence could reduce structural complexity and long-term integration burden.
3. Strengthen data foundations	Establish common identifiers, aligned timestamps, and normalized data models across surveillance domains. Invest in scalable ingestion pipelines and strong data governance to help support cross-channel analytics and regulatory defensibility.
4. Implement targeted cross-domain use cases	Start with focused initiatives such as correlating high-risk trading activity with related communications patterns or consolidating case workflows into a unified view. Use measurable improvements in detection quality and time-to-disposition to guide broader rollout.
5. Deploy AI within governance guardrails	Introduce AI in clearly defined workflows such as initial alert triage, clustering, and contextual assembly. Maintain model validation, explainability, monitoring, and human oversight to enable consistent and defensible outcomes.
6. Align operating models and accountability	Clarify clear ownership of surveillance outcomes across the surveilllance lifecycle. Standardize escalation pathways, documentation practices, and performance metrics across first- and second-line functions to help support integrated risk coverage.

Contact us

For additional information on the Our Take series or PwC's Risk and Regulatory Practice please contact:

Amanda Cox
Financial Services Risk & Regulatory Leader
773 456 5019
amanda.cox@pwc.com

Adam Gilbert
Global Senior Regulatory Advisor
914 882 2851
adam.gilbert@pwc.com

Martha Pampel
Managing Director
312 833 3385
martha.pampel@pwc.com

For additional information on **Surveillance in Financial Services**, please contact:

Scott Levine
Principal
954 240 6994
scott.a.levine@pwc.com

Josh Horowitz
Principal
201 406 4161
joshua.horowitz@pwc.com

Contributing Authors: Jason Maan, Sanam Shah, Jared Carrier, Ryane Thorne

pwc.com