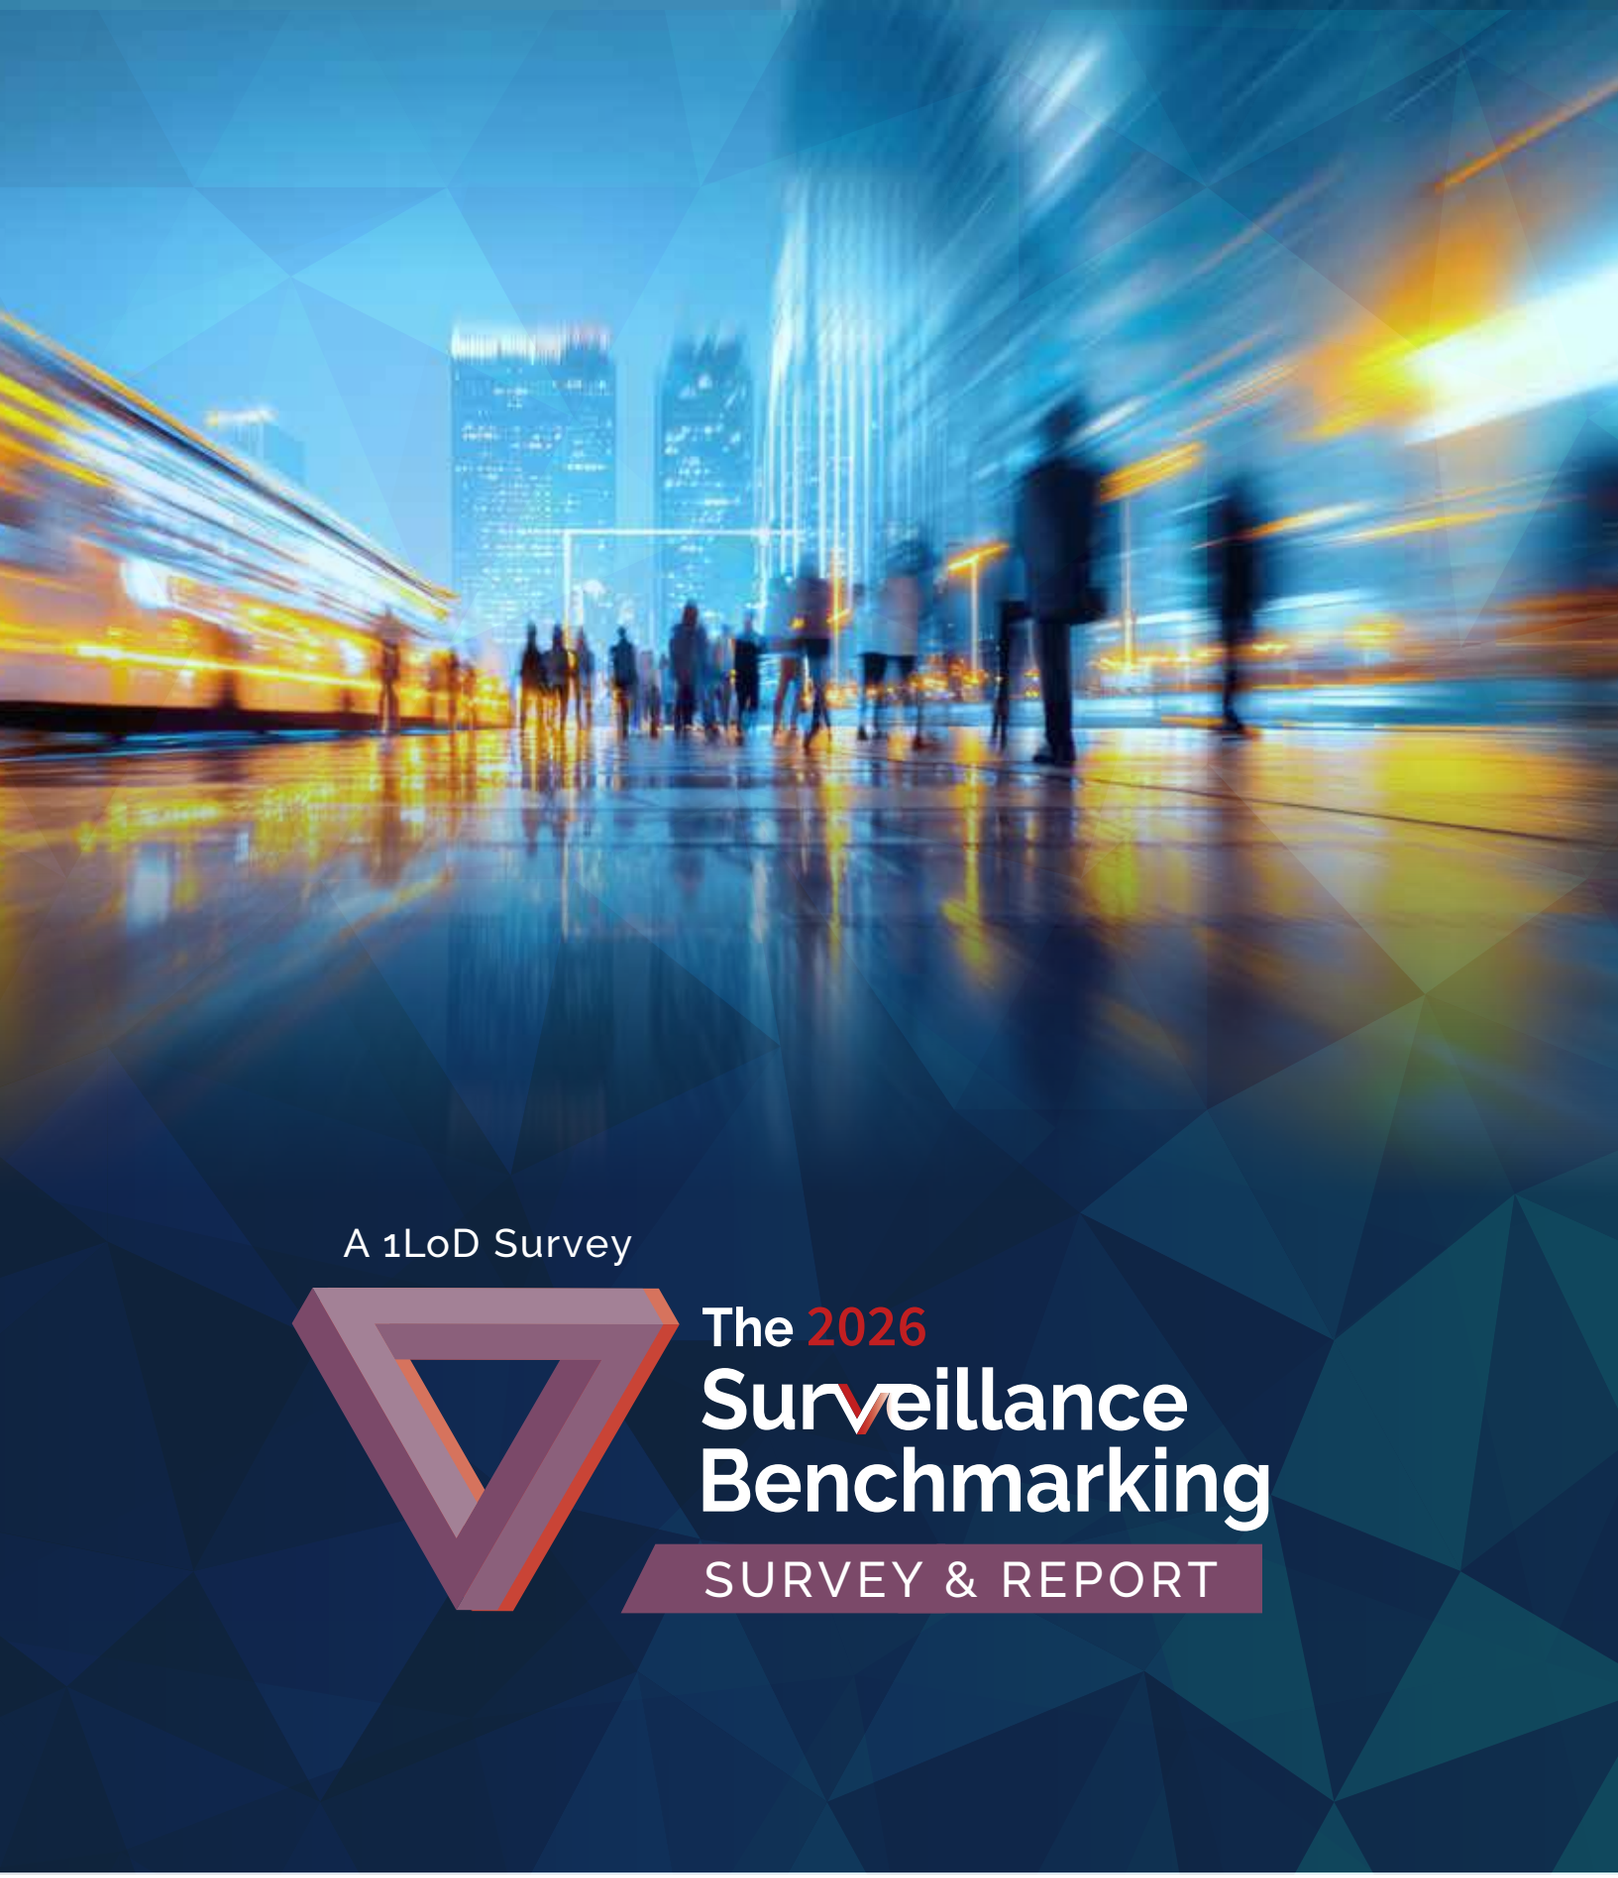

A 1LoD Survey

The 2026 Surveillance Benchmarking

SURVEY & REPORT

CO SPONSOR*

MCO
MyComplianceOffice

*Bespoke version created for MCO

Subscribe here
to receive new content and event information directly to your inbox

Introduction

The calm before the storm? How the next three years could reshape market abuse prevention and detection

These are the results of 1LoD's completely updated version of its 2024 Surveillance Benchmarking Survey & Report. Since then, we have had thousands of conversations with surveillance leaders across the industry, including detailed consultations with the members of our Surveillance Leaders' Network, who represent the most experienced surveillance leaders in the market.

This survey is the result of those conversations and it incorporates a host of new datapoints designed to help you benchmark your own trade, e-comms and voice surveillance processes against your peers.

The survey is divided into the following sections:

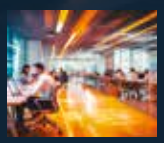

i. People

ii. Coverage

iii. Technology

iv. The future of surveillance

Copyright 2026, 1LoD. All rights reserved

Key takeaways

Only **37%** of firms are satisfied with their **trade surveillance technology**

52% of trade surveillance solutions are **in-house or hybrid builds** (versus 26% in e-comms and 29% in voice)

67% of firms have sustained or targeted investment plans for **e-comms surveillance technology**

89% of firms use **communications surveillance tools** to proactively monitor culture

82% of firms describe the state of cross-product market abuse surveillance at their firm as "Developing" or "Basic"

67% of banks say applying surveillance across multiple communication channels is a major challenge

44% of firms see their budgets for market abuse surveillance rising in 2026

48% of firms still have not linked any of their surveillance controls, raising questions about the progress of hollistic surveillance

93% of firms do not use trader behaviour analysis to filter their alerts

For **37%** of firms, legacy or outdated surveillance systems are a highly significant challenge to maintaining or advancing their surveillance capabilities

For **71%** of firms, fragmented, unstandardised and inconsistent data quality are the biggest hindrance to the effectiveness of their surveillance systems

For **37%** of firms, 'limited access to quality data' is a highly significant challenge to maintaining or advancing their surveillance capabilities

Foreword for 2026 Surveillance Benchmarking Survey & Report

I'm delighted to have been invited to write the foreword for 1LoD's 2026 Surveillance Benchmarking Survey & Report. With the Market Abuse Regulation's (MAR) 10th birthday just a few months away, it's a timely opportunity to reflect on the current landscape and good practices in the market.

Clean, fair markets promote growth and wider participation. Fighting financial crime is central to the Financial Conduct Authority's strategy, and firms play a critical role. Our collective success here will deliver better outcomes for markets and the people who rely on them. As markets and technology evolve, compliance teams face growing challenges in keeping pace to protect the marketplace from harm. The 2026 Surveillance Benchmarking Survey & Report helps firms recognise common problems and potential solutions. We hope it will also help them understand where they are now, and importantly, where they need to be.

On 4 July 2016, MAR (now UK MAR) came into force. It brought with it significant change and firms needed to adapt. Since then, markets have become more fragmented and complex, while monitoring tools have seen growing sophistication.

But in the surveillance world, some things remain the same. Strong fundamentals were always important, but the evolving landscape has, if anything, made them vital. Without the core elements of an effective programme, the foundations, firms will be poorly placed to tackle market abuse as it shifts with the changing environment.

Ten years into MAR, a suitably granular market abuse risk assessment – one that considers execution methods, types of market abuse, business areas and other factors – is not a nice-to-have, it's an essential. As are surveillance controls mapped to risks, venue inventories and strong data governance, amongst other things. We now expect all firms to have fundamentals such as these nailed down. And where we identify firms where they are not in place, we will consider regulatory action.

On top of the fundamentals, there are the challenges arising from evolving factors like new products, trading mechanisms, market structure, technology, plus geopolitical unpredictability and increased market volatility. Then there are growing market abuse risks, such as pump and dump, and using dominant positions to manipulate commodities. Firms need to be alive to them – the old risks, the new risks… the new versions of the old risks.

We are also looking more at those areas that have been more challenging to monitor, e.g. cross product/venue market abuse, and instruments traded over-the-counter. We appreciate surveillance here is often far from simple, but doing nothing is not an option. Firms with unmonitored risks – where there is no attempt to mitigate them, and no plans for strategic solutions – are a long way from meeting our expectations.

In our experience, some of the best

Copyright 2026, 1LoD. All rights reserved

Subscribe here
to receive new content
and event information
directly to your inbox

work being done, with the best outcomes, is at firms who strive for continuous improvement. That means ensuring the day-to-day monitoring is at a good standard, while scanning for new risks and working with vendors or in-house teams to develop and refine surveillance systems. We are keen to see new technological solutions being deployed, whether that involves AI, or the use of more sophisticated alerts using tailoring or contextualising. But we are also conscious that human expertise and judgement remain vital elements – technology is an enhancer, not a replacement.

As you look forward then, think of your surveillance programmes as needing continuous development and improvement. This will require senior management support and the maintenance of investment. We know this isn't easy, but we also know that it's essential as the challenges will keep coming. Standing still will inevitably mean going backwards. And as we have all seen in the past, ultimately, today's challenges end up being tomorrow's fundamentals.

Richard Littlechild
Head of Secondary Market Oversight - Enforcement & Market Oversight
Financial Conduct Authority

Firms with unmonitored risks – where there is no attempt to mitigate them, and no plans for strategic solutions – are a long way from meeting our expectations.

Section i: **People**

Copyright 2026, 1LoD. All rights reserved

Deceptively stable

Across many areas of financial services, the past few years have been characterised by rapid technological change. Artificial intelligence (AI), cloud data platforms, and increasingly sophisticated analytics tools are transforming how institutions approach risk, compliance, and operational oversight. Yet the results of this year's surveillance survey reveal something striking: despite this broader technological momentum, many core surveillance practices appear relatively stable.

At first glance, this stability might suggest that the surveillance industry has plateaued at a level of maturity that requires just minor tweaking year to year. Most institutions now operate established surveillance programmes capable of monitoring trading behaviour and communications across multiple channels. Governance frameworks have largely stabilised. Regulatory expectations are well understood, and most banks have implemented systems capable of identifying a wide range of market abuse scenarios.

However, a closer examination of the survey data suggests that the apparent stability instead represents a transitional phase—a period in which the industry is approaching a structural inflection point.

The survey findings highlight a persistent operational reality: surveillance programmes remain heavily constrained by the quality and structure of the data that underpins them. Respondents consistently identify fragmented data environments, inconsistent identifiers, and poor data quality as the primary factors limiting the effectiveness of surveillance systems, especially newer and more sophisticated surveillance engines incorporating advanced analytics, natural language processing, and behavioural models. These capabilities can only operate effectively when supported by high-quality, integrated data.

In many institutions, the underlying data infrastructure required to support this level of analytical capability is not there yet. In this environment, improvements in analytical sophistication alone are unlikely to deliver significant gains in surveillance effectiveness. Instead, the next phase of surveillance evolution will likely depend on progress in a different domain: enterprise data architecture.

Financial institutions are increasingly investing in large-scale data modernisation initiatives designed to break down internal data silos and create unified behavioural datasets. These initiatives—often driven by broader AI strategies—have the potential to fundamentally reshape surveillance capabilities.

Yet the survey results also suggest that many institutions have not yet completed the foundational work required to enable this transformation. Data fragmentation, inconsistent identifiers, and limited communications integration remain widespread challenges.

This is the calm before the surveillance storm: the underlying technologies capable of reshaping surveillance—cloud data platforms, AI, and large-scale behavioural analytics—are now widely available. The question facing financial institutions is no longer whether these capabilities exist, but how quickly the organisational and data infrastructure required to support them can evolve quickly enough to unlock their full potential.

Technology inflection point

When viewed collectively, the technology results in this year's survey suggest that market abuse surveillance may be approaching a genuine inflection point. For more than a decade, the architecture of surveillance has evolved incrementally: broader communications capture, larger data volumes, improved lexicons, and more sophisticated—but still largely deterministic—rule-based detection models.

The latest results indicate that several underlying structural transitions are now converging at once. Rather than representing incremental improvement, the technology landscape appears to be moving through three simultaneous shifts: an infrastructure transition, an analytical transition, and an operational transition. Together, these changes have the potential to reshape how surveillance functions actually detect misconduct.

The past decade has focused on establishing the core infrastructure necessary to meet regulatory expectations for market abuse detection. The coming decade is likely to focus on something more ambitious: transforming surveillance from a compliance monitoring tool into a broader, preventive and integrated intelligence capability.

Copyright 2026, 1LoD. All rights reserved

Beneath that big picture is the convergence of three related themes

1. Crunch time for surveillance
2. Agentic AI: the biggest transformative force in surveillance
3. Coverage is still king

Crunch time for surveillance

A structural tension is evident in both the data and in conversations with surveillance leaders: to maintain current programmes, and what is still in most cases a compliance-driven operating model, present levels of investment are sufficient to give banks a very high degree of confidence that their surveillance functions are effective enough (93% of banks are confident or very confident that their surveillance function is effective enough to detect potential future cases of market abuse **[Chart 23]**.

However, banks are also still mostly only partly satisfied with their trade, voice and e-comms surveillance solutions and 41% say that budget and resourcing constraints are a highly significant challenge to maintaining or advancing their surveillance capabilities surveillance coverage. This reveals a surveillance function at a crossroads.

Regulatory expectations around coverage, data integrity and preventive controls are expanding; there is still significant growth in surveillance data volumes and communication channels, but operational teams remain relatively small.

In addition, AI, cloud data platforms, and increasingly sophisticated analytics tools are transforming how institutions approach risk, compliance, and operational oversight and are creating pressure on surveillance teams to implement new technology to cut costs, to improve efficiency and to be able to develop the next generation of preventive controls regulators want to see.

However, even the current generation of sophisticated surveillance tools, incorporating advanced analytics, natural language processing, and behavioural models can only operate effectively when supported by high-quality, integrated data. And as the survey shows, data is the Achilles heel of surveillance with 71% of firms saying that the combination of fragmented data across silos, lack of standardised data formats/identifiers and poor or inconsistent data quality are the biggest hindrance to the effectiveness of their surveillance systems **[Chart 38]**. And for 37% of firms, 'limited access to quality data' is a highly significant challenge to maintaining or advancing their surveillance capabilities **[Chart 37]**.

The next generation of AI tools will require even more in the way of good data, and in many institutions, the underlying data infrastructure required to support this level of analytical capability requires significant investment.

If surveillance continues to be structured and resourced traditionally, as a manual-process-heavy, noise-dominated, compliance function at current levels, then it is difficult to see how it can make the leap to a modern, automated, risk-management function. And right now, the survey shows a function that is under pressure to transform faster than resources allow. For example, in e-comms surveillance, 69% of firms say either no-one or between 1% and 10% of staff are working on transformation projects. In trade that is 48% and in voice it's 76% **[Chart 28]**.

What the survey and interviews show though, is that banks do understand the necessity of adopting new technology, re-structuring surveillance skillsets and getting data right. The transformation is coming.

People

2

Agentic AI: the biggest transformative force in surveillance

Almost every conversation with surveillance chiefs reveals AI to be the most important development in the surveillance space. That said, banks are at very different points in their AI adoption journeys, and it is not always the case that the biggest and richest are the most advanced.

To a surprising degree, AI adoption in surveillance still seems to rely on somewhat ad hoc implementations that exist as much through the specific enthusiasm of individuals as through a genuinely formal process of use case evaluation, return on investment (ROI) analysis and rigorous proof of concept (POC). This process runs in parallel with somewhat vague enterprise-level exhortations to use generic AI tools such as Microsoft Copilot.

Even where this is not the case, worries about regulatory approval, concerns over the need for and usefulness of human-in-the-loop oversight, and the requirement for model-risk governance teams to approve AI-augmented processes create a series of bureaucratic hurdles that are clearly getting in the way of surveillance transformation. It seems strange that in an era of fully AI-autonomous weapons systems, surveillance teams are having to justify why an AI tool capable of analysing data in quantities and at speeds far faster than any human team, should be overseen by humans – surely this simply moves the bottleneck one step further down the process rather than removing it?

Vendors already offer AI-enabled managed services solutions that include the replacement of Level-1 analysts with AI agents, and banks – mostly mid-sized but not all – are adopting them. This suggests that in fact, the true obstacles to change are institutional, and internal, not exogenous.

At the leading banks, large language models (LLM) are already old hat. They are an everyday tool for automated narrative generation, investigative workflow support, summarising alerts, contextual communications analysis and similar tasks.
For these banks, and indeed a wider group of survey respondents, agentic AI will be the most transformative development in surveillance in the last decade.

As one head of trade surveillance says, "We will see the industry moving from using AI to support analysts to using agentic AI that can perform parts of the surveillance workflow itself."

These agentic systems can gather information across datasets, assemble relevant context around alerts, draft narratives explaining the activity observed, and even perform initial quality checks on the analysis. In effect, they function as digital surveillance analysts, capable of executing the first stages of investigative workflows.

This shift matters because the majority of cost and effort within surveillance programmes sits in L1 alert review. Surveillance teams across the industry

Copyright 2026, 1LoD. All rights reserved

3 Coverage is still king

are struggling with growing volumes of alerts, high false-positive rates, and increasing expectations from regulators around documentation and evidencing of investigative reasoning. Agentic AI directly targets this bottleneck by automating the hard graft of the surveillance process: data collection, contextualisation, and first-pass analysis.

At present, early implementations appear to favour communications surveillance (e-comms and voice) rather than trade surveillance. The reason is largely technical: agentic systems currently perform better when analysing qualitative data such as text and conversations, compared with the highly structured quantitative data associated with trading activity. As a result, initial deployments often focus on areas where AI can interpret language and context more effectively, before expanding into more complex trade-based analytics.

Nevertheless, the longer-term trajectory suggests convergence between communications and trading surveillance. Future systems may combine communications context with trading behaviour, enabling agents to correlate conversations, instructions, and trade execution patterns within a single investigative workflow. For example, identifying potential front-running behaviour often requires linking trading activity with communications around orders or client instructions. Agentic AI systems could eventually perform this type of cross-dataset analysis automatically. Are we about to enter the age of the synthetic surveillance analyst?

Although discussions of innovation—particularly around AI and agentic workflows—have become increasingly prominent in the surveillance community, the interviews conducted for this report and the data suggest that the day-to-day reality for most surveillance teams remains far more operational. Across multiple institutions, practitioners repeatedly returned to the same set of challenges: ensuring complete venue coverage, expanding communications surveillance across new channels, supporting additional languages, and maintaining the integrity of the underlying data feeds that make surveillance possible in the first place.

One of the most persistent themes is the ongoing challenge of achieving complete trading venue coverage. The enforcement actions that followed the widely discussed surveillance failures of recent years have reinforced the expectation that firms should know exactly where they trade and ensure those venues are captured within their surveillance frameworks. Yet practitioners note that achieving this level of completeness is operationally complex and remains a continuous exercise rather than a one-off remediation effort. As one participant observed, surveillance coverage is not something that can simply be "fixed"; it is a process that must mature continuously over time.

Alongside venue coverage, firms continue to devote significant effort to data completeness and control frameworks. Several surveillance heads in Europe and the US described expanding internal

governance functions dedicated specifically to monitoring the health of surveillance data feeds, identifying outages, and ensuring that trading and communications data is arriving in surveillance systems complete, accurate and on time.

The expansion of communications surveillance coverage is another recurring operational theme. While technological developments such as AI-driven analytics attract significant attention, many surveillance teams are still grappling with the practical challenge of capturing and monitoring an ever-expanding range of communications channels. New platforms are increasingly being incorporated into official communications workflows, requiring surveillance teams to extend recording and monitoring capabilities into environments that were not originally designed for regulatory oversight. For some institutions, the addition of a single new channel represents a significant operational undertaking, requiring new recording infrastructure, testing processes, and integration with existing surveillance workflows.

Language coverage has also emerged as a major operational priority. Several banks describe significant investment in expanding lexicon and model coverage across multiple languages. For these firms, the challenge is not simply one of translation and transcription but ensuring that detection logic remains consistent across languages and cultural contexts. AI transcription and translation models have improved hugely, but banks still report issues with language coverage and basic recording quality in voice comms.

Another operational theme is the growing need to extend surveillance beyond internal trading activity to include client and market access activity. In particular, the surveillance of direct market access (DMA) and prime brokerage clients is receiving greater attention as regulators scrutinise firms' ability to detect misconduct originating from external participants using their trading infrastructure. For some institutions this represents an expansion of surveillance responsibilities beyond traditional employee monitoring, requiring additional datasets, controls, and analytical capabilities to understand activity conducted by external clients.

In practice, all this means that change-the-bank initiatives such as AI adoption often run in parallel with a large volume of run-the-bank activity focused on coverage and control. The interviews suggest that this operational workload remains the dominant reality for many surveillance teams. Innovation may shape the future direction of surveillance technology, but the present is still defined by the ongoing effort required to ensure that surveillance coverage is complete, reliable, and regulator-ready across an increasingly complex trading environment.

Copyright 2026, 1LoD. All rights reserved

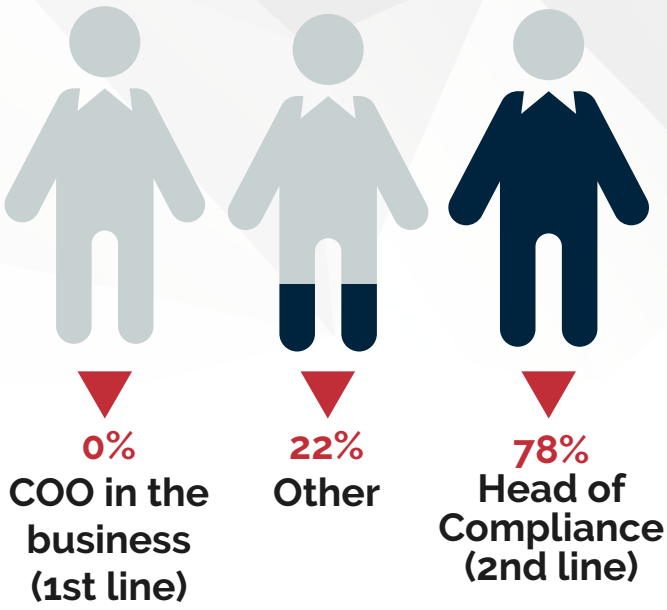

CHART 25
As the Head of Market Abuse Surveillance in your banks, who do you report to?

A team embedded close to the trading environment is often better positioned to understand how products behave in volatile conditions.

People, who, where, how many

Surveillance reporting lines remain firmly anchored in the 2LOD with 78% of heads of market abuse surveillance reporting into the head of compliance **[Chart 25]**. This reporting structure reinforces the independence of the function from Front Office revenue generation. However, it also frames surveillance primarily as a compliance-managed control environment rather than as a 1st line owned risk-detection capability embedded within business operations.

In practice the picture is more nuanced, with different controls being operated by the 1st and 2nd lines, with quality assurance (QA) and testing split across the two lines and shared escalation responsibilities depending on the type of abuse referenced by alerts.

The debate is as old as surveillance but respondents to the survey interviewed on this topic still think it is important and banks are still moving certain surveillance activities back and forth between the two lines.

Advocates of 1st line surveillance argue that proximity to the business is essential if surveillance is to be effective. Trading behaviour, market structure, and product dynamics evolve constantly, and understanding these changes requires close interaction with the individuals actually conducting the business.

When surveillance teams sit within or near the 1st line, they are often better able to observe shifts in trading patterns, understand the mechanics of specific markets, and interpret unusual activity. Several participants emphasised that regular dialogue with traders, sales staff, and desk heads can provide valuable context that improves the quality of surveillance analysis. Business teams may proactively explain changes in strategy, product mix, or execution venues, enabling surveillance teams to distinguish between legitimate commercial developments and potentially suspicious behaviour.

This proximity can also strengthen the design and calibration of surveillance controls. Effective surveillance depends on correctly interpreting market behaviour and tuning alerts to capture genuine risks while minimising noise. A team embedded close to the trading environment is often better positioned to understand how products behave in volatile conditions, how liquidity shifts across venues, and how legitimate trading patterns may evolve over time.

In practice, surveillance is rarely a static control environment. New products, regulatory expectations, and market conditions constantly require enhancements to models, thresholds, and monitoring logic.

Several survey participants argued that this continuous development cycle aligns more naturally with a 1st line environment, where teams can work directly with the business to implement improvements rather than acting purely as external reviewers.

There is also a pragmatic dimension to this argument. Surveillance functions often require sustained investment in technology, data infrastructure, and specialised staff. When surveillance sits too far from the business, it can be perceived as a remote compliance cost centre. Budget requests may appear abstract or disconnected from commercial priorities.

By contrast, when surveillance leaders can clearly explain to business heads why particular capabilities are needed—drawing on direct knowledge of trading activity—they may find it easier to secure funding and operational support.

Finally, 1st line placement may allow surveillance teams to act more effectively in driving change. Compliance functions traditionally identify issues and escalate them, but the business ultimately owns remediation. When surveillance sits closer to operational decision makers, it may have greater influence over how controls are implemented, calibrated, and improved.

Copyright 2026, 1LoD. All rights reserved

The risks of 1st line surveillance

Despite these advantages, locating surveillance within the 1st line raises persistent concerns about independence. The core purpose of surveillance is to detect and investigate potential misconduct within the business itself. If the function sits too close to the population it monitors, there is a risk—real or perceived—that conflicts of interest may arise.

Survey participants described practical examples of this tension. In some cases, when surveillance sat within the 1st line, investigations became difficult because the business resisted sharing information or attempted to shape how inquiries were conducted. Situations involving clients or sensitive trading strategies could become particularly contentious, raising doubts about whether the surveillance function could operate with full independence.

Regulators and boards also remain sensitive to these structural risks. Even if 1st line surveillance operates effectively in practice, institutions must still demonstrate that credible oversight exists. In many jurisdictions, supervisors expect to see clear mechanisms through which independent challenge can be exercised. Without such safeguards, a purely 1st line surveillance model may struggle to satisfy regulatory expectations.

For these reasons, even proponents of 1st line surveillance generally acknowledge that some form of independent oversight remains necessary. This often results in additional compliance review functions or 2nd line assurance teams, complicating the organisational model.

The case for 2nd line surveillance

The traditional argument for 2nd line surveillance is grounded in governance and independence. Compliance functions exist specifically to provide challenge to the business and ensure that conduct risks are appropriately managed. Locating surveillance within compliance therefore creates a clear structural separation between the monitoring function and the activity being monitored.

This separation can be particularly important during investigations and regulatory reporting decisions. Determining whether activity should be escalated to regulators, classified as suspicious, or formally investigated carries significant legal and reputational implications. A 2nd line function may be better positioned to make these determinations without undue influence from business interests.

2nd line placement also aligns more naturally with the conventional three lines of defence model. Many institutions and regulators remain committed to the principle that the 1st line owns risk, the 2nd line oversees and challenges, and the 3rd line provides independent audit. Within this framework, surveillance sitting in compliance is conceptually straightforward and easy to explain to external stakeholders.

Moreover, 2nd line surveillance can strengthen formal challenge mechanisms. Compliance functions often play a key role in reviewing control effectiveness, assessing regulatory risk, and ensuring that escalation processes are followed correctly. A surveillance team embedded within compliance may therefore integrate more seamlessly into broader governance structures.

The squeezed middle

But the question remains: given the obvious benefits of a 1st line surveillance function, and the existence of audit and model risk management, is 'independence' the only real argument for keeping reporting lines into compliance? And what is the value of an independent check-and-challenge function if it does not understand what are extremely complex and fast-evolving businesses and markets?

CHART 9
Does your surveillance function have a dedicated testing and / or quality assurance (QA) resources?

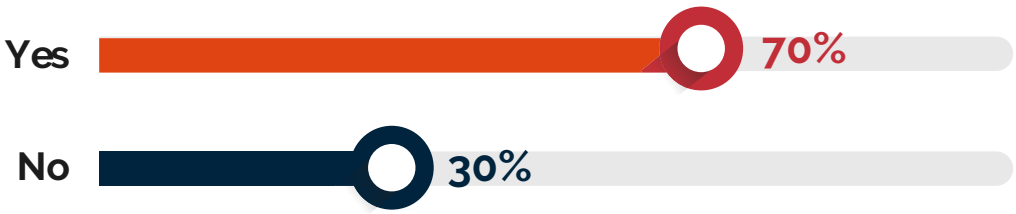

These questions are becoming more pressing for a number of reasons: if model risk management (MRM) teams are closely involved in model evaluation and even business as usual (BAU) parameter recalibration, and data integrity is increasingly managed by specialised data teams, then that takes another potential function away from a generalist surveillance team. Also, a generalist 2nd line function is more likely to slow innovation in surveillance technology and methodologies while it gets up to speed – and technology agility will be a key driver of surveillance success in the next few years.

And what about the duplication – and not just with the 3rd line? 70% of banks surveyed have control testing and QA capabilities in the surveillance function **[Chart 9]**.

Expertise is everything

The answer surveillance heads come up with is straightforward: it doesn't matter who the function reports to, it matters how much subject matter expertise it has and how close its relationship with the business is.

While this would seem to favour a 1st line reporting structure, clearly most banks think it is possible to upgrade skills in the 2nd line while maintaining the independence they believe regulators demand.

In the meantime, many institutions have moved toward hybrid surveillance structures that combine elements of both lines. In these models, operational surveillance activities—such as alert triage, initial analysis, and interaction with the business—often sit within the 1st line. Analysts may be recruited from trading or sales backgrounds, enabling them to interpret market activity with a high degree of expertise.

At the same time, the 2nd line retains oversight responsibilities. Compliance may review escalations, approve control changes that could weaken the monitoring environment, and ultimately determine whether regulatory reporting obligations are triggered.

This structure attempts to capture the benefits of both proximity and independence. 1st line teams maintain close engagement with the business and possess the expertise required to interpret complex trading behaviour, while 2nd line teams provide governance oversight and independent challenge.
In practice, such models often involve multiple layers of review within the surveillance process. Initial alert triage may be conducted by operational analysts, followed by deeper investigation by subject-matter specialists before escalation to compliance. This staged approach ensures that compliance receives well-developed cases while preserving its authority over final regulatory decisions.

Copyright 2026, 1LoD. All rights reserved

CHART 14
How many covered persons is your surveillance function responsible for surveilling?

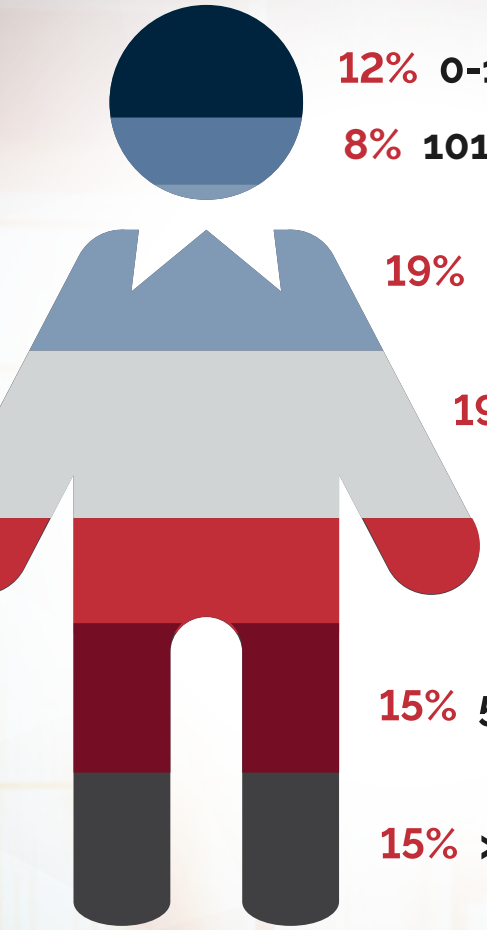

12% 0-100

8% 101-500

19% 501-1000

19% 1001-2000

12% 2001-5000

15% 5001-10000

15% >10000

Scaling surveillance is a big issue

The survey covers a range of institution types with varying surveillance needs. 20% of respondents had under 500 individuals to surveil, but beyond this group, the surveillance perimeter is large and the scale of monitoring is essentially industrial. 38% of firms surveil between 500 and 2000 individuals and 42% need to monitor more than 2,000 people. Almost a third of the respondents oversee functions that have more than 2,000 covered persons with 30% surveilling between 5,000 and more than 10,000 people.

CHART 26
How many full-time equivalent (FTE) headcount in your team are performing alert reviews on a daily basis?

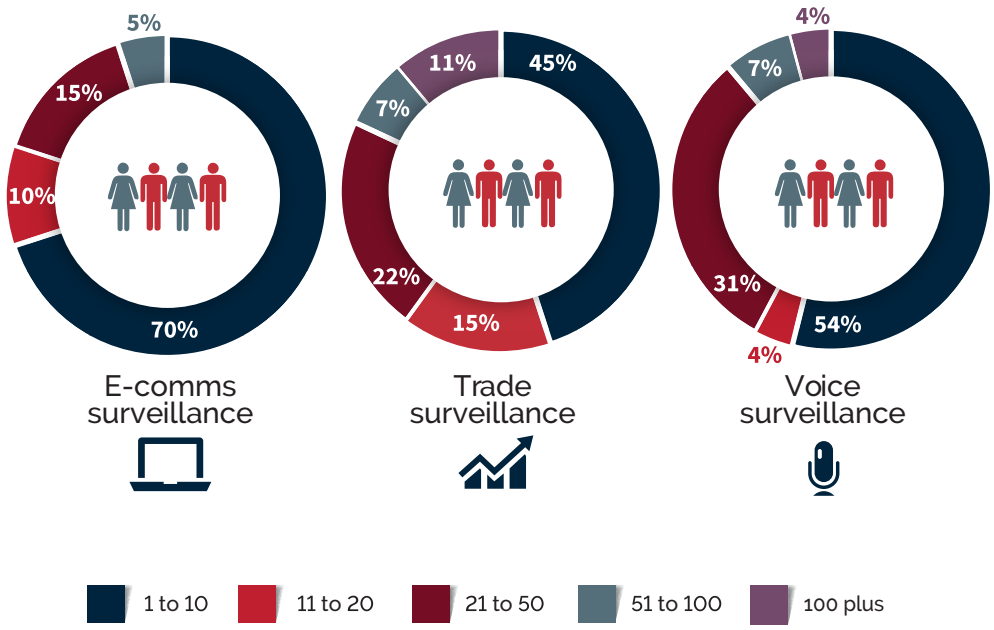

Surveillance heads often take these numbers for granted, but given the complexity and granularity of surveillance requirements, this represents an industrial-scale monitoring obligation, particularly in the context of fragmented execution environments and proliferating communications channels.

However, the human resources allocated to this task look more like the kind of numbers that would have made sense at the beginning of the market abuse regulatory revolution, when markets were entirely different and surveillance was genuinely viewed as a narrowly defined regulatory compliance activity.

For example, in e-comms surveillance, 80% of respondents have 20 or fewer full-time equivalent (FTE) performing daily alert reviews (with 15% having between 21 and 50 and only 5% at between 51 and 100 with – none at 100+). And yet banks are surveilling populations in the thousands, across dozens of communications channels in an expanding palette of languages.

It's a different picture in voice where only 58% of respondents have 20 or fewer FTE performing daily alert reviews and a much more significant 31% have between 21 and 50 FTE reviewers (versus the 15% of e-comms.) In voice too, unlike e-comms, the largest banks do have teams of 100 plus (4% of respondents).

But it is in trade surveillance where the difference is most stark: 60% still only have between one and 20 reviewers, but 22% have 21–50, 7% have 51–100 and 11% have 100+.

So, what do these findings actually tell us in the context of the kinds of alert volumes that one global bank estimated for 1LoD? That bank said that a similar institution to themselves might see 650,000 e-comms alerts in a year and only 10% of that in trade.

So, the fact that so many more human resources are applied to reviewing that much smaller number of trade alerts confirms that banks still see trade surveillance as their primary tool for the detection of (and perhaps also the prevention of) market abuse.

This is unlikely to change. Trade surveillance is directly tied to market abuse regulations and is therefore subject to the greatest regulatory scrutiny. Also,

Copyright 2026, 1LoD. All rights reserved

because these alerts relate to actual executed trading behaviour, they carry higher legal risk.

Trade alerts are clearly escalated into more complex investigations requiring market context analysis, order book reconstruction, cross-venue behaviour analysis and trader strategy evaluation and so a single alert investigation can take significantly longer than reviewing a suspicious email or chat. This explains why 11% of firms have teams larger than 100 reviewers—a scale not seen in e-comms surveillance.

Taking trade surveillance alerts seriously is still a process that needs people. Even if you use a managed service provider of L1 analysts or, as some firms have, you have automated that job, the investigative process is still manual and labour intensive. New AI-based solutions are just at the beginning of making an impact **[see Section iii: Technology]**.

In voice, banks still struggle not just with volume but also with more mundane issues such as recording quality from underlying squawk box and phone technology, as well as the challenging task of extracting meaning from what are often multiple-speaker overlaid and shouted conversations in context-dependent argot. Voice-to-text, especially across all of the languages used by banks, is getting better, but it will need to get better still to reduce the human resource load the survey reveals.

As one head of surveillance at a top US investment bank says, "When you're doing voice-to-text on top of brokers shouting out prices and fuzzy lines, it's difficult. The machines are great with good recordings, but with terrible recordings they're as bad as a human being."

By contrast, in e-comms current natural language processing (NLP) and NLP and lexicon-based detection systems are still generating huge alert volumes, but these are then being triaged by automated and AI-driven tooling that allows a relatively small team to run the process (again see Section iii: Technology for a deep dive into these topics). There is also a clear implication that there is heavy reliance on filtering solutions or risk-based sampling to keep alert volumes manageable. However, these results also reveal, perhaps, that e-comms is still not seen as a primary detective tool: the contrast with trade, and the fact that much of the alert handling has to be done automatically, makes that clear.

It also shows why a genuinely holistic approach (combining trade and comms data) and subsequently generating appropriate alerts is so challenging. The disparity in alert volumes, and in their implied value, makes blending the data sets hard even without the data quality and noise issues inherent in each data set, as well as the structured/unstructured data conflict **[see Section iv: The Future of Surveillance for more on this]**.

The underlying challenge

Underneath all that though, these charts reveal a fundamental challenge for both present and future: surveillance is currently scoped like a broad perimeter function (in some cases a third of banks are surveilling more than 5,000 people) but it is often staffed like a super-lean ops function. And all three surveillance functions need to be able to scale to meet the coverage challenges ahead.

The only solution here is technology, as survey respondents overwhelmingly report that their teams are at full stretch and budgets for adding staff are constrained. "People are at capacity and when additional work appears it creates significant stress on the team," says one global head of surveillance. "The problem remains the same: it's hard for the bank to quantify the benefits of what we [surveillance teams] do and the only really easy pitch [for us to get more budget] is to say that we will turn FTE into technology."

Where are the teams?

The past couple of years have seen rapid development in the capabilities of surveillance solutions. There have been advances in older 'AI-like' machine learning (ML) and NLP, as well as newer technologies in LLMs and behavioural analytics. Banks' own models have been improved, and vendors have engaged in an arms race to develop the next generation of surveillance solutions.

However, this has not yet translated into any significant shifts in the core surveillance operating model. Aggregating across the three surveillance buckets, **Chart 29** shows that the model is essentially about 50% hub, 30% offshore, and 15–20% nearshore. This is almost identical to the results from two years ago (53% hub, 31% offshore, 16% near shore). The fundamental geography of the surveillance workforce appears largely unchanged.

Where the newer chart becomes more interesting is in revealing meaningful variation between different surveillance disciplines. Trade surveillance remains the most hub-centric activity, with 58% of staff located in financial centres. This reflects the continuing importance of proximity to the business, market expertise and interaction with traders, surveillance specialists and model validation teams. Trade surveillance rules, calibration and investigations often require a nuanced understanding of market behaviour and product structures, which tends to anchor these roles in major financial centres such as London, New York or Hong Kong. The data therefore reinforces the idea that trade surveillance retains a strong expert-driven, judgement-heavy component, making it less suitable for large-scale relocation to lower-cost jurisdictions.

By contrast, e-comms surveillance shows the highest degree of offshoring, with around one third of headcount located in offshore centres. This likely reflects the significant operational nature associated with e-comms monitoring, which often involves large volumes of alerts generated by lexicon-based or behavioural analytics systems. Much of the work consists of first-level triage and review of alerts rather than deep market analysis. As a result, banks appear more comfortable locating a significant proportion of these roles in offshore service centres where the work can be processed at scale and at lower cost. Hub-based teams still remain dominant, however, suggesting that escalation, investigations and policy oversight are still concentrated in core financial centres.

Voice surveillance sits somewhere between these two models but shows a slightly higher nearshore presence than other surveillance types. While just over half of voice monitoring staff remain in hubs, nearly one fifth are located in nearshore locations—higher than either trade or e-communications surveillance. This distribution likely reflects the complexity of voice analysis, which often requires language capability, contextual interpretation and familiarity with regional trading practices. Nearshore centres can offer a compromise between cost efficiency and skill availability, providing strong language coverage and regulatory familiarity while remaining closer to major financial markets than fully offshore centres.

> *"There isn't a great shift in offshore/onshore teams in e-comms, trade and voice… AI may be starting to have an impact, but you're not seeing big shifts in bodies yet."*

Copyright 2026, 1LoD. All rights reserved

CHART 29

What percentage of your surveillance headcount is currently located in:

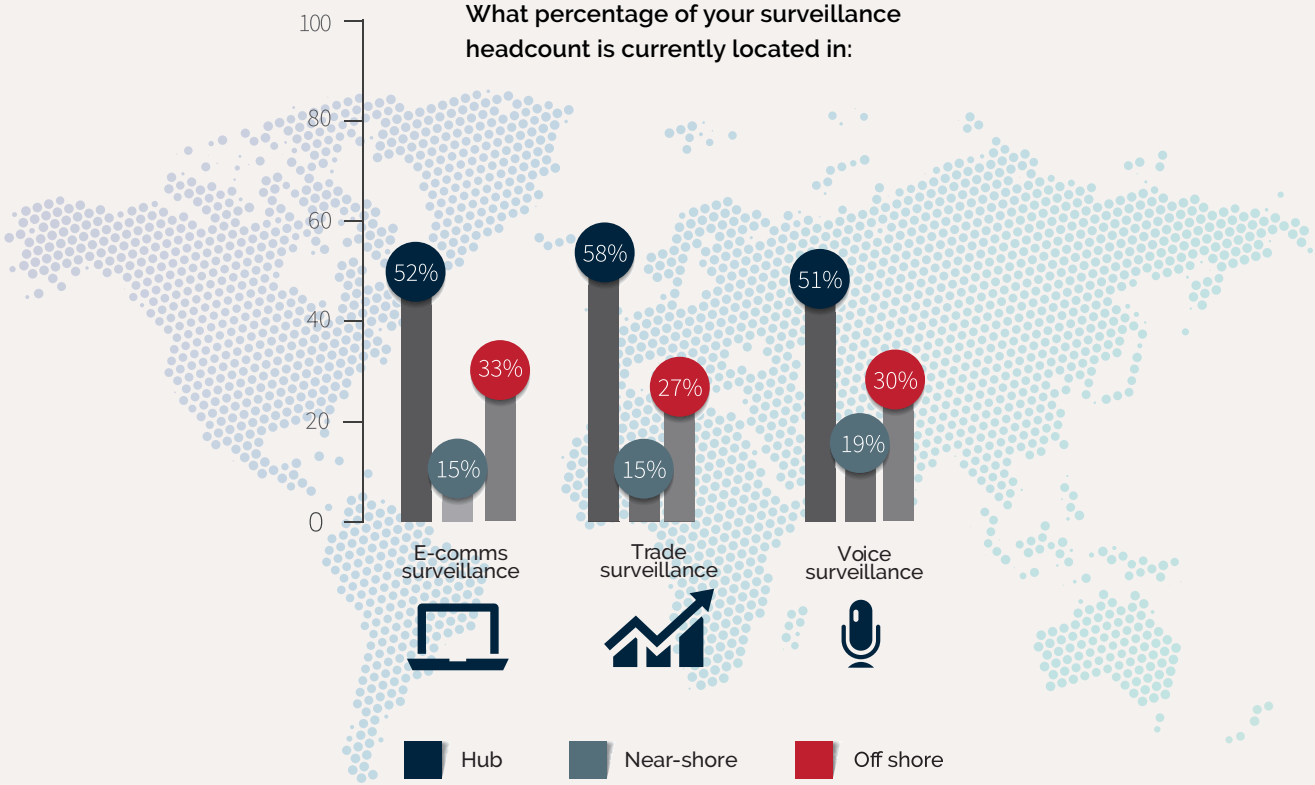

A broader implication of these results is that surveillance functions appear to be evolving into a stratified workforce model. Core expertise, governance and complex investigative work remain concentrated in financial hubs. A second layer of analytical or language-dependent review is increasingly placed in nearshore centres. Meanwhile, high-volume operational tasks—particularly initial alert review in communications surveillance—are more commonly handled in offshore locations. This three-tier structure allows banks to balance cost pressures with the need to maintain regulatory credibility and subject-matter expertise in core markets.

Perhaps the most striking takeaway from comparing the two charts is not the differences between surveillance types but the overall stability of the model itself. Despite significant technological change in the surveillance landscape—including the introduction of advanced NLP, ML-driven behavioural analytics and large-scale data processing platforms—the geographical footprint of surveillance teams has not shifted dramatically. This suggests that while technology may improve efficiency and analytical capability, the human oversight component of surveillance remains critical. Regulatory expectations, investigative judgement and the need for close interaction with trading businesses continue to anchor a substantial proportion of surveillance staff in major financial hubs.

This is confirmed by many respondents with one head of surveillance saying, "There isn't a great shift in offshore/onshore teams in e-comms, trade and voice… AI may be starting to have an impact, but you're not seeing big shifts in bodies yet."

In effect, the data suggests that surveillance technology is not replacing the traditional operating model so much as augmenting it. Banks appear to be using automation to support and scale existing teams rather than radically redesign the geographic distribution of their workforce. The result is an operating model that combines centralised expertise with distributed operational capacity—a structure that, for now at least, appears to have reached a relatively stable equilibrium across the industry.

Section ii: **Coverage**

Copyright 2026, 1LoD. All rights reserved

The results of the survey, taken together with interviews with respondents, reveal that coverage remains the defining operational challenge for surveillance teams. While technology innovation—particularly AI and voice transcription—has improved monitoring capabilities, the core problem remains the same: ensuring that every relevant person, every communication in every material language, and every aspect of trading activity including venue, falls within the surveillance perimeter.

Three themes stand out:

1

Venue and data coverage remain complex, particularly for institutions operating across fragmented trading platforms

2

Language coverage is expanding rapidly, driven by regulatory pressure and advances in AI-enabled monitoring

3

Communications channels are still proliferating, creating continuous challenges for capture and monitoring infrastructure

In practice, surveillance leaders describe coverage as a permanent operational programme rather than a problem that can be fully solved. As new language requirements, channels, and venues emerge, surveillance teams must continually expand their monitoring capabilities to ensure that no material activity occurs outside the surveillance perimeter. Ultimately, technology is seen as the answer to this challenge.

Surveillance scope creep

The broad challenge of surveilling large populations, and the logistics of doing so, have already been discussed. A more granular dive into the details shows a number of significant developments.

When survey participants were asked which business lines they are surveilling **[Chart 15]**, the comparison with the previous (2024) survey is revealing. It's no surprise that in both 100% surveil the markets business, but in most other categories the direction is one way: increased surveillance of non-markets businesses. So, in the latest survey 85% surveil banking versus 77% in 2024; custody is 30% this year versus 23% previously; wealth is 44% versus 36%; retail is 22% versus 14%. A little surprisingly, corporate hardly changes (56% in this survey versus 55% in the last).

The overall picture is clear: a combination of improved data availability, platform trading, and regulatory pressure is encouraging banks to widen the perimeter of surveillance into areas that were previously less scrutinised, including client activity and business lines outside core markets trading.

Banks are expanding market abuse surveillance to custody and wealth management to address gaps in regulatory compliance, rising cross-product manipulation risks, and the need for a holistic view of client activity.

CHART 15
Which business lines are you surveilling?

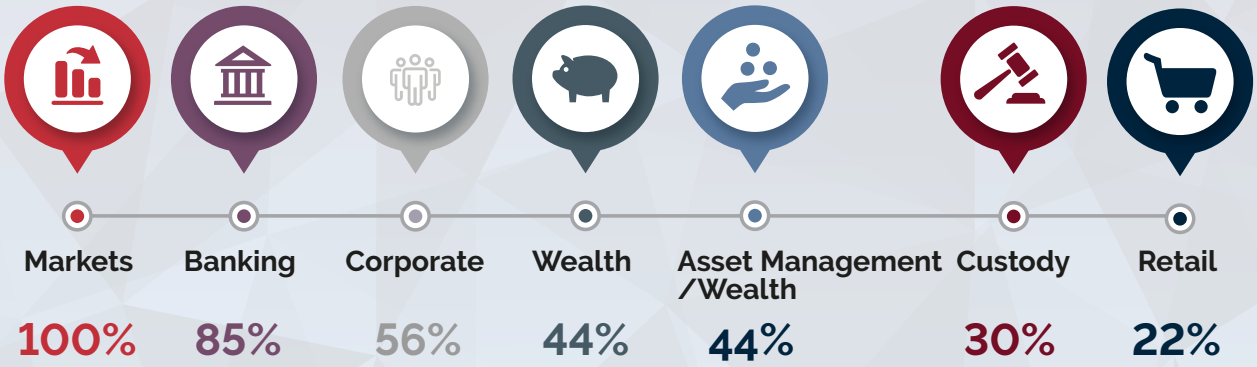

Regulators want more

Perhaps the most direct driver of all of this is sustained regulatory scrutiny of surveillance frameworks. Interviewees repeatedly emphasise that supervisory attention remains intense and is pushing firms to ensure broader coverage of both activity and communications. One UK-based surveillance head described a recent review as "the most intensive supervisory audit I've had in my career… with over a hundred requests for information just on surveillance."

Increasingly, part of these reviews is investigation of coverage across a wider range of activities and business units, including those—such as custody or wealth management—that historically sat outside the core market surveillance perimeter. Regulators, including the FCA, are tightening scrutiny on these areas, viewing them as significant, often overlooked sources of potential market misconduct.

Regulators also want banks to understand more about what their clients are doing. Banks are responding to supervisory actions that highlighted gaps in firms' ability to monitor client activity. As one US surveillance head noted, enforcement cases have shown situations where firms "didn't have the controls to see what their clients were doing." The result is that surveillance teams are increasingly being asked to extend coverage to areas such as DMA and other forms of client-driven trading activity. Reflecting this, another comment emphasised the growing importance of monitoring these channels, noting that "expanding client surveillance around prime clients and DMA clients is becoming more important."

Regulatory demands are also moving in tandem with technological developments. As one global head of surveillance at a Global Systemically Important Bank (G-SIB) explains, technological improvements and richer venue data are now making this possible in ways that were difficult in the past: "technology and the data coming from venues are allowing us to surveil more of our client activity, especially on the listed side."

In practice this means that surveillance is no longer confined to the behaviour of internal traders. Banks are increasingly able—and in some cases required—to examine how clients interact with their trading platforms and execution channels. This has obvious implications for businesses such as prime brokerage, custody trading services, and wealth management platforms, where clients execute orders through the bank's systems rather than through traditional broker-mediated trading.

Regulatory scrutiny of communication capture across multiple platforms, and the adequacy of language coverage across globally distributed workforces, also inevitably draws more businesses into their crosshairs. Wealth advisers, relationship managers, and other client-facing staff generate large volumes of communications and transactional interactions that regulators increasingly expect firms to monitor.

Copyright 2026, 1LoD. All rights reserved

More convergence, more scope creep

Another driver of this expansion is the growing convergence between surveillance and other risk-detection disciplines within the bank. Surveillance teams increasingly interact with functions such as financial crime, transaction monitoring (TM) and trading controls, all of which are analysing overlapping behavioural signals.

As one EMEA head of surveillance explained, financial crime teams may observe suspicious transactional patterns while trading controls or surveillance teams analyse market activity, yet "we're all looking to do the same thing… detect suspicious activity." In practical terms, this convergence encourages the pooling of analytical techniques and investigative tooling across functions. Once the underlying capability is understood as behavioural anomaly detection rather than purely market-abuse monitoring, it becomes logical to apply similar techniques to other activities across the bank.

Data integration and contextual analytics reinforce this trend. Modern surveillance platforms are no longer limited to analysing isolated trade data or single communications channels. As firms build ever-more sophisticated investigative environments that combine communications data, voice transcripts, trade lifecycle data, and external contextual information, then the analytical framework naturally applies to a broader range of activity — including client communications in wealth management, custody transaction flows, or cross-business conduct risks — rather than being confined to trading desks alone.

Governance developments are also expanding the conceptual scope of surveillance. Surveillance models are increasingly being incorporated into enterprise model risk management frameworks, alongside pricing models, credit risk analytics, and other formalised analytical systems. This change reflects a broader regulatory expectation that surveillance systems be documented, validated, and governed with the same discipline as other analytical models. Once surveillance tools are treated as enterprise analytical infrastructure rather than isolated compliance utilities, the natural question becomes where else those analytical capabilities might be applied within the organisation.

Tech, tech, tech

And more generally, the more technology allows things to be surveilled, the more they will be. So, the accelerating adoption of artificial intelligence is another factor widening surveillance's potential scope. AI tools are increasingly capable of performing analytical tasks that historically required significant manual review effort, including thematic analysis of communications, root-cause investigation, and large-scale control testing. While firms remain cautious about relying on AI outputs without human validation, these technologies significantly reduce the marginal cost of analysing large data volumes. As a result, the economic barrier that once limited surveillance coverage to high-risk trading populations is beginning to weaken. Institutions can now realistically consider applying similar monitoring techniques to much larger employee populations and business activities.

Taken together, these developments point to a gradual but significant transformation in the role of surveillance within financial institutions. What began primarily as a market abuse detection capability for trading desks is evolving into a broader enterprise framework for behavioural risk monitoring, drawing on integrated data sources, advanced analytics, and cross-functional collaboration. As surveillance technology becomes more scalable and governance frameworks more formalised, banks are likely to continue expanding the application of these tools across business lines where behavioural risk, conduct oversight, and regulatory scrutiny intersect.

Venues – an ongoing challenge

The difficulty of achieving complete venue coverage has emerged as one of the most persistent operational challenges in modern surveillance programmes. Participants repeatedly describe it as a problem that is both technical and structural: banks must identify every venue on which trading occurs, obtain the relevant order and trade data from those venues, and then ensure that the data is reconciled, ingested, and analysed within surveillance systems. The reality, however, is that the trading ecosystem has become so fragmented that achieving perfect coverage is widely viewed as aspirational rather than realistic.

More than half of respondents to the survey trade on more than 50 venues **[Chart 17]** and a large global bank can easily be using several hundred. Maintaining full coverage requires firms to keep track of every venue they trade and ensure that they have the required data from each of those venues and that reconciliations exist for each data feed.

Many surveillance leaders privately acknowledge that they will be struggling with this for years because their trading activity is spread across a constantly changing network of exchanges, broker platforms, internal systems, and electronic liquidity venues. As new execution venues emerge or existing venues change their data structures, surveillance teams must continually revisit their coverage frameworks.

> *In general, the large banks (who have also been under the most regulatory scrutiny) are the most confident that they have achieved, or are very close to achieving, full coverage.*

Copyright 2026, 1LoD. All rights reserved

Or, as one head of trade surveillance at a very large US bank puts it, "I think the painful part is that, given all the work that has gone into trying to get this right, there should be no more surprises and yet they still keep coming."

There is an increasing divide here between the very largest and best-resourced institutions and the rest. In general, the large banks (who have also been under the most regulatory scrutiny) are the most confident that they have achieved, or are very close to achieving, full coverage.

The only hiccup for the large banks – and it's an issue for everyone – is still acquiring and ingesting venue data. Several survey respondents emphasised that obtaining complete pre-trade and trade data from execution platforms is far from straightforward. Some execution venues simply struggle to provide the necessary data feeds or provide them in formats that are difficult to integrate into surveillance infrastructure. This problem is particularly acute for institutions that operate with multiple fragmented systems rather than a single global order management system (OMS).

Another complication is that venue coverage increasingly extends beyond traditional executed trades to include pre-trade signals and client activity. Some participants highlighted growing regulatory and industry interest in monitoring datasets such as request-for-quote (RFQ) activity, which can reveal potential behavioural patterns before orders are actually executed. Bringing these datasets into surveillance systems requires additional venue integrations and raises questions about data quality and consistency across platforms. As surveillance scope expands into pre-trade activity and client order flow, the data universe that firms must capture becomes even larger.

Finally, the venue coverage challenge intersects with broader issues of data governance and infrastructure. Surveillance teams increasingly manage hundreds of data connectors and must oversee reconciliation, feed monitoring, and vendor integration across multiple systems. These operational tasks have become so significant that surveillance functions often find themselves performing data management roles traditionally handled by technology teams. The complexity of maintaining these pipelines means that venue coverage is not simply a one-time implementation exercise but an ongoing operational discipline requiring constant monitoring and maintenance.

CHART 17

How many trading venues are you monitoring?

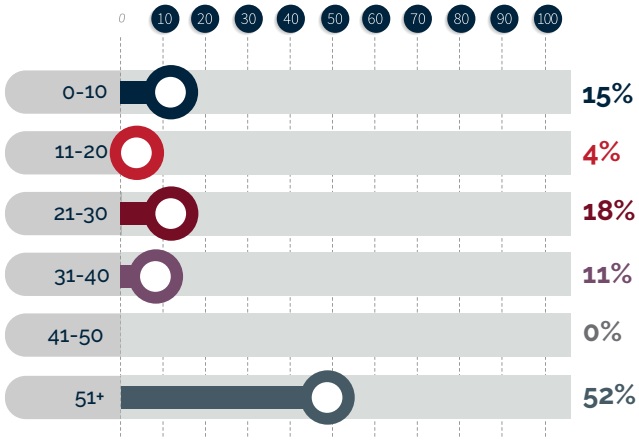

"The painful part is that, given all the work that has gone into trying to get this right, there should be no more surprises and yet they still keep coming."

Keeping up with e-comms

Another major coverage challenge is the continued expansion of electronic communications channels used by traders, sales teams and others. While banks typically monitor core channels—such as email, Instant Bloomberg, and Microsoft Teams—they repeatedly emphasise the difficulty of identifying and capturing newer or niche communication platforms.

The survey results show that 22% of firms are monitoring more than 30 channels and 33% more than 20 **[Chart 10]**. In 2024, just 14% covered more than 30 channels.

One reason for the rise is straightforward: regulatory enforcement actions have spurred banks to audit e-comms channels and are finding more of them. One surveillance leader at a large US bank described an audit that identified approximately 70 additional channels that might require monitoring or restriction. The other key driver is, once again, the development of technologies that enable more channels to be monitored without hiring larger teams or that allow previously hard to capture channels to be integrated into existing processes.

One interviewee explained how regulatory pressure had forced their organisation to expand surveillance coverage: "We had a National Futures Association (NFA) action because we didn't have Teams meetings recorded and surveilled… even though it's not the primary place anybody's committing the bank to execution, there are still potential sales conversations and transactional discussions taking place."

Banks do not, however, seem to have been able to meaningfully speed up the addition of new channels to surveillance. Surveillance teams emphasise that each new channel requires a complicated sequence of building capture infrastructure, and implementing archiving and retention controls, even before integration into surveillance analytics and then testing to confirm data completeness.

CHART 10
How many E-comms channels are you monitoring?

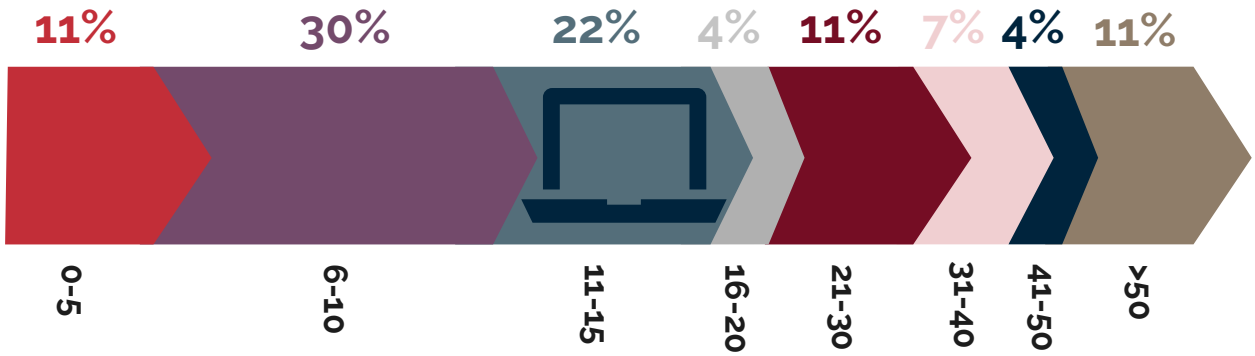

Copyright 2026, 1LoD. All rights reserved

Getting languages licked

It's the same story with language coverage. In 2024, 82% of respondents were covering 10 or fewer languages in their communications surveillance function. Only 9% were monitoring more than 15. And the largest cohort (46%) covered five or fewer.

In this year's survey, the number only covering five or fewer has almost halved to 26% **[Chart 12]**. And 26% (versus 18%) are covering between 11 and 15, with 15% (versus 9%) covering more than 15.

CHART 12
How many languages does your communications surveillance function cover?

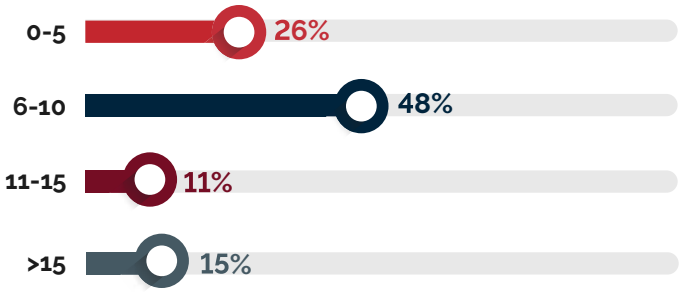

The expansion of language coverage in communications surveillance is no longer a marginal or experimental trend; it has become a central pillar of how banks are redefining surveillance completeness. Across multiple institutions, there is a clear recognition that historical models—largely built around English-language monitoring—no longer reflect the linguistic reality of global trading environments. As one head of e-comms surveillance explains, "We are now undertaking a significant increase in language coverage. We do nine now, and we'll be 14 by the end of 2026." This is not simply incremental expansion; it represents a structural shift towards multilingual surveillance as a baseline expectation.

What is particularly striking is the breadth and diversity of languages now being brought into scope. Beyond the traditional major financial languages—English, French, German, Japanese—firms are increasingly targeting what might previously have been considered peripheral languages within this context. One bank described a formalised "global coverage standard" under which "if we've got something in English… let's make sure we've got each lexicon model also enabled for our eight core languages," with further expansion underway into "Turkish, Dutch, Swedish, Ukrainian, and Korean."

The inclusion of languages such as Ukrainian and Hindi—highlighted by another participant as areas where "a fair few people [are] suddenly starting to want Hindi… that's new"—illustrates how surveillance priorities are increasingly being shaped by actual communication patterns rather than legacy assumptions.

Critically, this expansion is not being driven by innovation for its own sake, but by a growing acknowledgement that language represents a material coverage gap. Several participants explicitly framed multilingual capability in these terms. One described how their "language range was seen as a coverage gap, and… [they are] filling that coverage gap," while also acknowledging that "we are not yet mature… we will add more languages."

Regulators on the case

Regulatory pressure has been a decisive catalyst in accelerating this shift. Firms report that supervisors are no longer satisfied with generic assurances of communications monitoring but are instead probing the alignment between monitored languages and actual business usage. As one bank noted, regulators have "been in and… put pressure on us for that," while another survey respondent highlighted direct intervention, with the US NFA highlighted as "pushing banks to improve language monitoring". The implication is clear: insufficient language coverage is a potential regulatory failing.

At the same time, the expansion of multilingual surveillance has been enabled—perhaps for the first time at scale—by advances in AI and language technologies. Practitioners consistently point to the role of NLP, transcription, and translation capabilities in unlocking this shift. One survey respondent observed that the "biggest leap" has come from "the ability of these AI tools to take language properly, transcribe it accurately, and then… process it," enabling firms to extend coverage into both communications and voice. Another highlighted that modern platforms can support "15 languages out of the box … and they're continuing to roll out more," making large-scale deployment increasingly feasible.

An additional—and relatively new—dimension to this trend is the emergence of data-driven approaches to identifying language gaps. Rather than relying on static assumptions, firms are increasingly using management information and analytics to detect unmonitored linguistic activity. As one practitioner explained, new systems can now identify "your most frequently spoken languages outside of your already detected ones… Korean… Ukrainian," turning what was once a manual, hypothesis-driven exercise into a dynamic, evidence-based process. This represents a significant maturation: language coverage is no longer a one-off design decision, but an evolving control calibrated against observed behaviour.

Despite this progress, participants are clear that multilingual surveillance remains operationally complex and imperfect. Technical challenges persist, particularly around linguistic nuance and real-world usage patterns. Firms report difficulties with "code switching," where individuals move between languages within a single conversation, and with translation accuracy, where "straight translation did not always capture intent."

Even seemingly straightforward implementations can prove problematic; one bank noted that Hindi communications written using English keyboards required entirely separate modelling approaches, as initial implementations "just didn't work." Moreover, not all languages are economically or operationally viable to automate. As one participant acknowledged, some languages remain too infrequent to justify full deployment, requiring reliance on "manual controls… [which are] not… sustainable" in the long term.

Several institutions described deliberate programmes to scale language coverage through lexicon expansion or AI-driven transcription. One surveillance leader noted that their communications monitoring now covers eight core languages, including English, Chinese, Japanese, French, German, Spanish, Portuguese, Hindi, and Italian, with plans to expand to approximately fourteen languages by 2026, adding Turkish, Dutch, Swedish, Ukrainian and Korean.

Regulatory pressure is a major driver of this expansion. Several interviewees referenced supervisory scrutiny—particularly from US regulators—around whether banks have adequate coverage of languages used by their trading populations. One surveillance head explained that their organisation faces particular pressure because of a large Chinese-speaking population, yet their communications surveillance tools historically had weak capabilities for monitoring Mandarin or other Asian languages.

Copyright 2026, 1LoD. All rights reserved

CHART 13
Are you moving towards a risk-based approach to comms surveillance and therefore not surveilling entire in-scope populations?

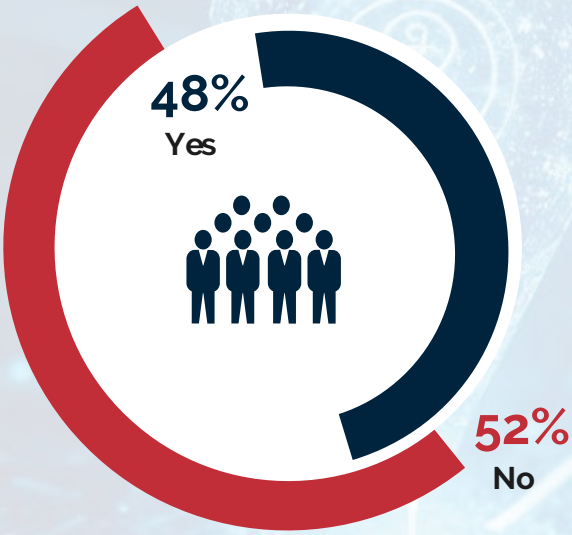

CHART 24
Does the risk of regulatory sanction limit your banks ability to move at pace to a more effective, risk-based surveillance?

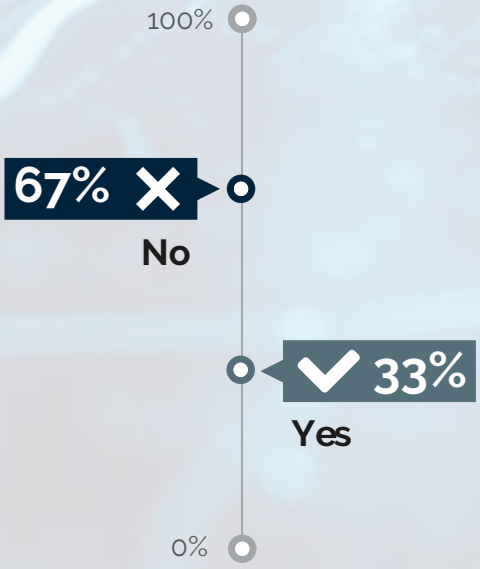

Moving towards a risk-based approach in comms?

Another way to measure progress in surveillance is the extent to which banks are moving towards more risk-based approaches towards coverage, alert prioritisation and the application of technology to areas such as automatic alert dispositioning.

Perhaps the most significant of these is the approach to in-scope populations and here the survey shows that more than half of respondents are not prepared to exclude some members of in-scope populations based on a risk assessment, while just under half are **[Chart 13]**.

Surprisingly there is no underlying pattern in terms of the types of institutions willing / unwilling to take a risk-based approach. The 'yesses' and 'nos' are spread randomly across the large and mature, the large playing catch-up, the mid-tier and the small. The only common feature is the degree to which fears of regulatory opprobrium – and 2nd line culture – dominate thinking.

However, that viewpoint is getting harder to justify. Asked whether the risk of regulatory sanction limits their ability to move to risk-based approaches, 67% of banks say that it does not – up from 59% two years ago. Yes, banks still under regulatory scrutiny feel that their latitude for reducing coverage or automating surveillance is limited by the granularity of the demands made by regulators' on-the-ground teams. But in general, banks accept that regulators simply want to see appropriate levels of monitoring and control with transparency of audit and explainability. They are not standing in the way of risk-based innovation.

A bigger brother?

A surprising area in which surveillance teams seem to have dramatically shifted their opinions in favour of scope expansion is in cultural monitoring. In 2024, many surveillance leaders spoken to by 1LoD were not in favour of using market abuse surveillance infrastructure for more general culture and conduct surveillance for a number of reasons. It is more work. It overlaps with other functions such as human resources (HR). And it is another potential large source of false positives and alerts that need to be processed.

So, while 59% were using their Market Abuse Regulation (MAR) surveillance processes to look for negative cultural indicators, 41% were not. Two years later there has been a dramatic shift. Eighty nine percent now say that they will use communications surveillance tools to proactively monitor culture, with just 11% saying they would not **[Chart 16]**.

So, what has changed? As ever, the hand of the regulator is present. Take the FCA's Policy Statement PS25/23 published in December 2025. The FCA singles out bullying, harassment and violence as being not simply HR issues, but explicitly linked to the "Fit and Proper Test" for employees and senior personnel.

The FCA is intent on preventing the development of workplace cultures that facilitate further wrongdoing and regulatory breaches that harm consumers and damage market integrity. This links conduct rules, culture and behaviour to market integrity and consumer outcomes. So, culture is now a regulatory risk vector, not just an internal people issue.

The report goes into great detail on the types of behaviour the regulator believes are leading indicators of current or future misconduct. Misconduct is not simply harmful in itself, it can "facilitate further wrongdoing and regulatory breaches". A key indicator, it believes, is boundary-pushing language and inappropriate tone in communications.

It makes sense then for banks to use e-comms surveillance tools to get ahead of an evolving regulatory regime. And what is making this expansion into behavioural and cultural monitoring feasible—without a commensurate explosion in headcount—is the rapid maturation of AI-driven surveillance tooling, particularly in natural language processing, transcription, and contextual analytics.

Participants consistently point to the ability of modern systems to ingest vast volumes of communications data, across multiple languages and channels, and extract meaning rather than simply flag keywords. They are also capable of detecting signals that extend well beyond classic market abuse indicators — meaning that, as banks build ever more sophisticated models to meet regulatory surveillance requirements, the same systems naturally became tools for detecting broader conduct and behavioural risks. And the regulators know it.

CHART 16

Are you using your communications surveillance tools to proactively monitor culture? That is, not if it is discovered when surveilling for MAR but looking for cultural indicators separately to MAR?

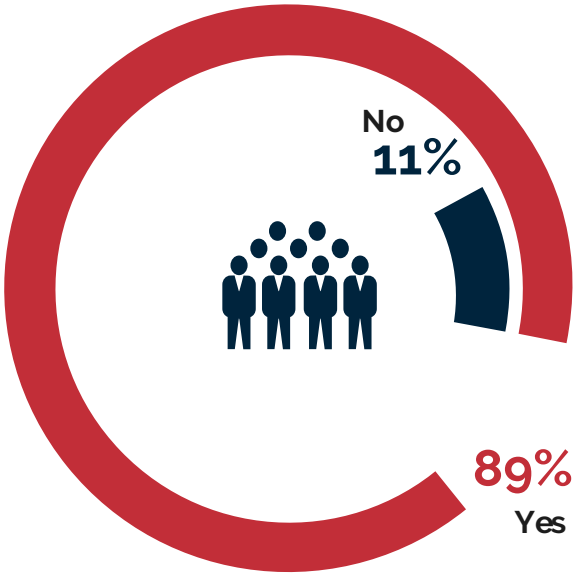

Copyright 2026, 1LoD. All rights reserved

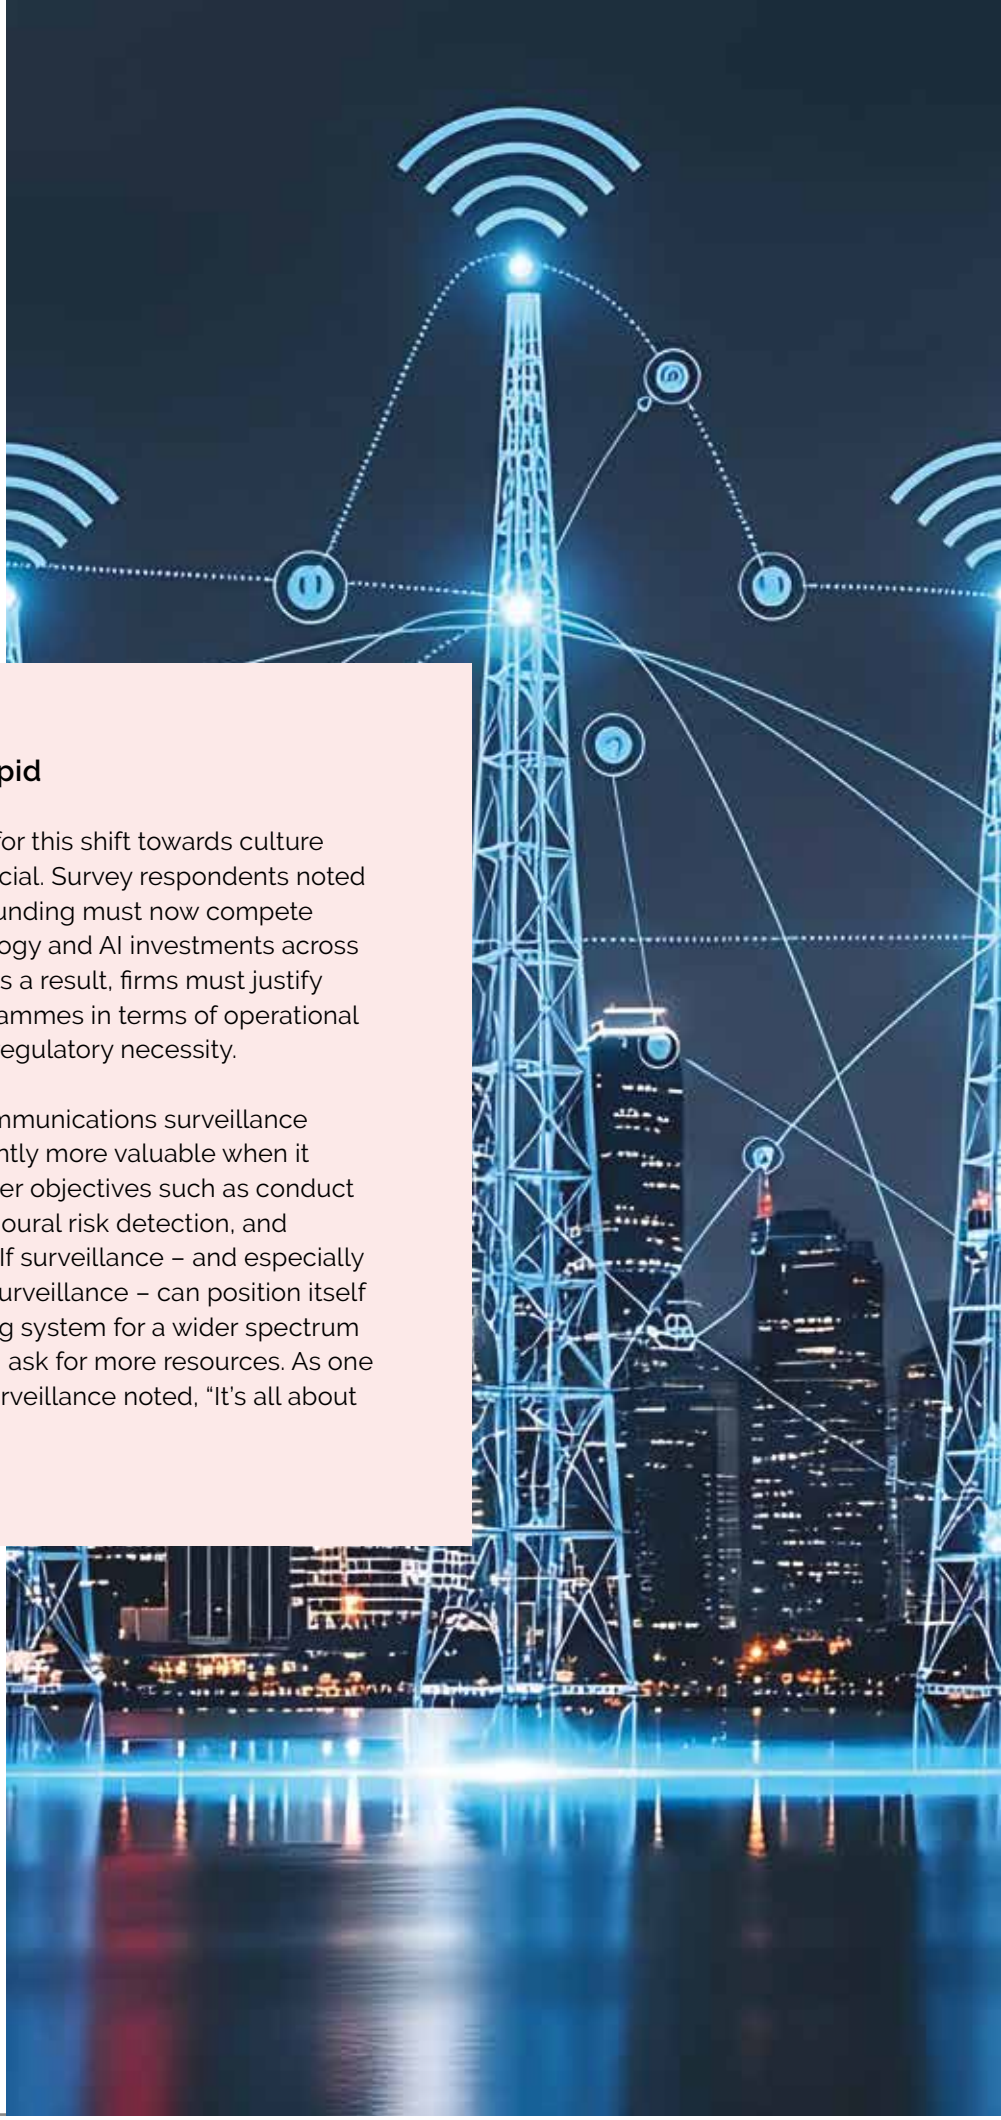

It's the ROI, stupid

One other reason for this shift towards culture monitoring is financial. Survey respondents noted that surveillance funding must now compete with other technology and AI investments across the organisation. As a result, firms must justify surveillance programmes in terms of operational benefit as well as regulatory necessity.

In this context, communications surveillance becomes significantly more valuable when it can support broader objectives such as conduct monitoring, behavioural risk detection, and cultural oversight. If surveillance – and especially communications surveillance – can position itself as an early-warning system for a wider spectrum of risks, then it can ask for more resources. As one head of comms surveillance noted, "It's all about the ROI now."

It's clear that the most significant pain points (the "high challenge" responses), centre on technology performance and data quality, rather than governance, regulation, or organisational design.

CHART 37

How significant are the current challenges in maintaining or advancing surveillance capabilities?

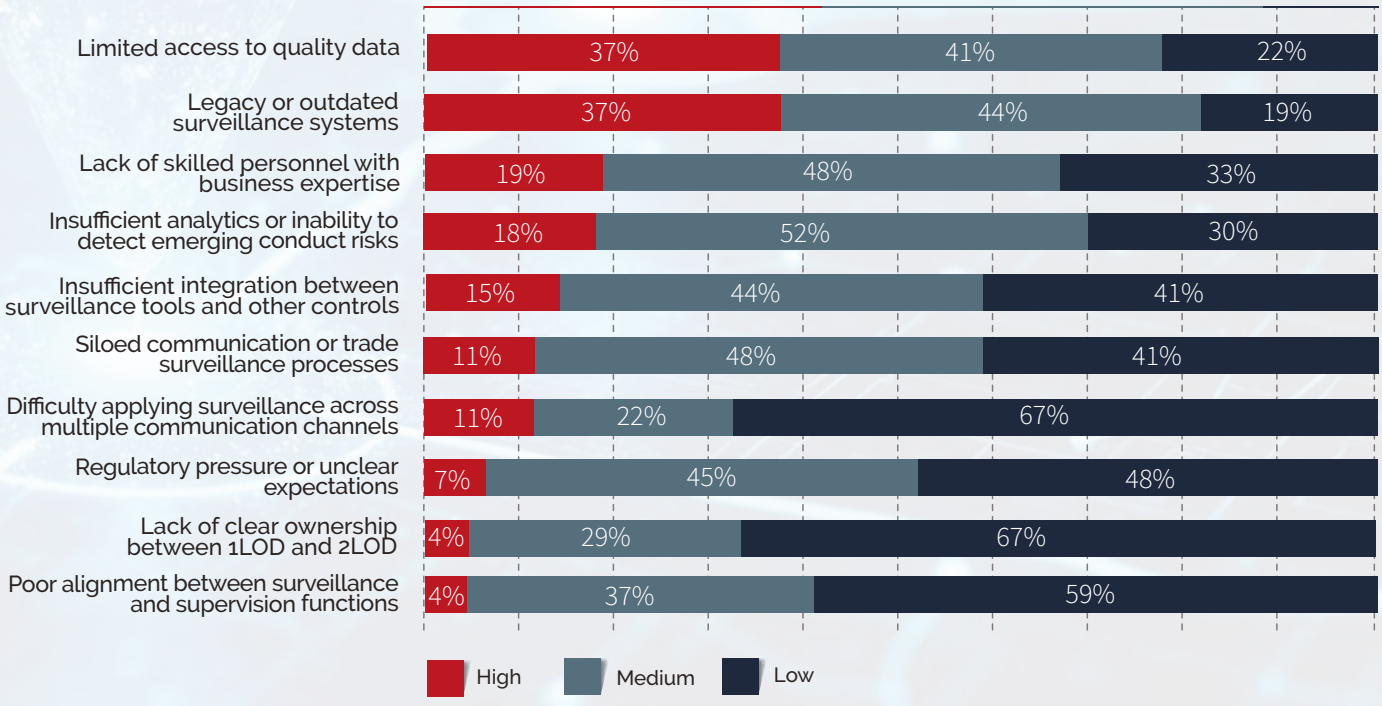

Copyright 2026, 1LoD. All rights reserved

Tech and data still the key pain points

Finally, when asked 'how significant are the current challenges in maintaining or advancing surveillance capabilities?' **[Chart 37]**, then it's clear that the most significant pain points (the "high challenge" responses), centre on technology performance and data quality, rather than governance, regulation, or organisational design.

The standout issue—by a significant margin—is the volume of false positives, with a majority of respondents identifying this as a high challenge. At first glance this suggests that despite years of investment, many surveillance environments remain fundamentally inefficient at the point of detection, with technology still struggling to distinguish meaningful risk from background activity.

However, when the lens **[on Chart 37]** is widened from "high" responses alone to include "high + medium", the picture becomes notably more systemic and less concentrated on a small number of acute pain points. The combined "high + medium" perspective reveals that these issues are not isolated pressure points but part of a broad, interconnected set of constraints affecting the entire surveillance stack. Data quality and legacy infrastructure remain central, but they are now joined by limitations in analytics capability, skills gaps, integration challenges, and budget pressures.

This broader view is, in a way, more problematic because it suggests that the key challenges cannot simply be solved by specific technology solutions. The real obstacles to improvements in surveillance are a wider structural ecosystem of constraint, in which each weakness reinforces the others. Poor data quality does not just drive false positives; it also limits the effectiveness of analytics. Legacy systems do not just create inefficiency; they inhibit integration and the adoption of more advanced techniques. Skills gaps and budget constraints, while not always critical in isolation, further slow the pace at which these foundational issues can be addressed. What emerges is not a single bottleneck but a layered architecture of friction, where incremental improvements in one area are often neutralised by limitations in another.

Perhaps most importantly, this combined perspective highlights a subtle but important disconnect in how firms perceive urgency. Many of these issues are widely acknowledged—indeed, majorities recognise them as challenges—yet relatively few are classified as "high." This suggests an industry operating in a state of "managed imperfection," where deficiencies are understood but tolerated because systems continue to function at an acceptable baseline level. The risk, however, is that this equilibrium is fragile. As surveillance expectations evolve—whether through greater emphasis on behavioural analysis, increased channel complexity, or regulatory expansion into conduct—these medium-level constraints may become critical. In that sense, the "high + medium" analysis paints a more forward-looking and cautionary picture: not of a system failing today, but of one that may struggle to adapt tomorrow without more fundamental investment in data, technology, and capability.

The survey results reveal a striking shift in how surveillance leaders perceive the obstacles to maintaining and advancing their surveillance capabilities. While governance structures and regulatory expectations have historically dominated industry debate, the data now paints a very different picture. The most severe challenges are no longer organisational or regulatory in nature. Instead, they lie firmly in the operational mechanics of surveillance itself: alert quality, data integrity, and technology architecture.

The most dramatic finding in the chart concerns the persistent problem of false positives. A clear majority of respondents (52%) identify the high volume of false alerts as a high-severity challenge, while a further 41% describe it as a medium challenge. This means that 93% of institutions consider false positives to be a meaningful operational issue, making it by far the most widely recognised obstacle to effective surveillance.

This result reinforces what practitioners frequently describe in interviews: traditional rule-based surveillance engines are still generating vast volumes of alerts with relatively low signal-to-noise ratios, forcing banks to devote large review teams to filtering out irrelevant alerts before genuine risks can be identified. In e-comms, where alerts are viewed as lower value, technology is being used to aggressively filter. But in trade, where the risks of missing true positives is perceived as much higher, large human teams are still the norm – with banks still setting up additional offshore hubs in some cases.

"We are still in the business of trying to go from 99% false positives to 95% and I would question the value of that," says one global head of trade surveillance at a large global bank.

Closely following this challenge are issues relating to data quality and technology infrastructure.

Limited access to high-quality data is identified as a high challenge by 37% of respondents, with a further 41% rating it as a medium challenge.

Similarly, legacy or outdated surveillance systems are rated as a high challenge by 37% of respondents and a medium challenge by 44%.

Taken together, these results reveal an obvious reality: surveillance systems are only as effective as the data pipelines and technology environments that feed them. In many institutions, fragmented communications capture, inconsistent trading data, and ageing vendor platforms continue to undermine the reliability and effectiveness of monitoring programmes. These infrastructure limitations often sit upstream of the alert generation problem itself, contributing directly to both false positives and missed risks.

Copyright 2026, 1LoD. All rights reserved

Budgets still an issue

Resource constraints also emerge as a significant operational pressure. Forty-one percent of respondents cite budget or resourcing constraints as a high challenge, with a further 48% rating it as a medium challenge **[Chart 37]**. In practical terms, this means nearly nine in 10 institutions consider funding pressures to be a meaningful barrier to advancing surveillance capabilities.

Interestingly, some of the topics that traditionally dominate industry debates appear to be far less problematic in practice. Governance questions such as the division of responsibility between the 1st and 2nd lines of defence rank among the lowest-severity issues in the survey. Fully 67% of respondents rate a lack of clear ownership between the 1st and 2nd lines as a low challenge, while 59% say poor alignment between surveillance and supervision functions is also a low challenge. These findings suggest that many banks have now reached relatively stable governance arrangements. The long-running debate about where surveillance should sit organisationally may therefore be far less important than the operational effectiveness of the programme itself.

Similarly, regulatory pressure and uncertainty appear less problematic than might be expected. Only 7% of respondents consider regulatory pressure or unclear expectations to be a high challenge, while nearly half regard it as a low challenge. This does not necessarily imply that regulatory expectations have softened; rather, it suggests that many institutions now feel they have a clearer understanding of what regulators expect and how to operationalise those expectations within their surveillance frameworks.

Another surprising result concerns the challenge of monitoring communications across multiple channels. Despite the industry's intense focus on off-channel communications and messaging platform enforcement actions in recent years, two-thirds of respondents describe this issue as a low challenge. This may indicate that most banks now feel technically capable of capturing communications across a wide range of platforms. The greater difficulty may lie not in capturing the data, but in analysing it effectively once it has been collected.

Taken together, the survey results reveal a clear shift in the industry's surveillance priorities. The primary obstacles are no longer structural governance questions or regulatory ambiguity. Instead, they lie in the operational complexity of modern surveillance systems: reducing false positives, improving data quality, modernising technology infrastructure, and delivering these improvements under increasingly constrained budgets.

Section iii: **Technology**

Copyright 2026, 1LoD. All rights reserved

In surveillance, all roads eventually lead back to technology, whether that means the battle to upgrade legacy systems, comply with regulatory requirements around data quality or keep up with developments in AI.

Taken together, the charts in this section suggest a surveillance ecosystem that is technologically approaching a major transition:

Most firms expect increased or sustained investment in surveillance technology over the next 24 months	**Key operational bottlenecks are becoming more problematic**
Data complexity continues to grow	**AI adoption is accelerating**

This combination strongly suggests that the industry may be approaching a technology-driven transformation phase, in which AI and workflow automation begin to reshape how surveillance systems actually detect and investigate market abuse.

When viewed in isolation, the 2026 satisfaction data already suggests that surveillance capability across banks has plateaued below the level at which institutions themselves would consider it analytically effective. However, when placed alongside the 2024 results, what emerges is not a story of steady improvement, but one of uneven technological progress across the surveillance stack and differing approaches to each element.

It's rough in trade

In trade surveillance sentiment is actively deteriorating **[Chart 1A]**. Between 2024 and 2026, full satisfaction declines from 43% to 37%, while dissatisfaction more than doubles from 9% to 19%. This is consistent with a structural mismatch between legacy rules-based detection architectures and the evolution of modern markets.

Surveillance systems designed to detect abuse in relatively transparent, order-book-centric environments are now being asked to infer intent across fragmented execution pathways that span both digital and conversational domains. They are also being tasked to track complex cross-product behaviours which they were not designed to do, and which are inherently challenging.

The survey data therefore points to a widening gap between the behavioural phenomena regulators expect firms to monitor—layering, spoofing, collusive price signalling—and the technical capacity of incumbent trade surveillance tooling to model those behaviours in practice.

Between 2024 and 2026, full satisfaction [in trade surveillance tooling] declines from 43% to 37%, while dissatisfaction more than doubles from 9% to 19%.

CHART 1A
Are you satisfied with your current trade surveillance solutions?

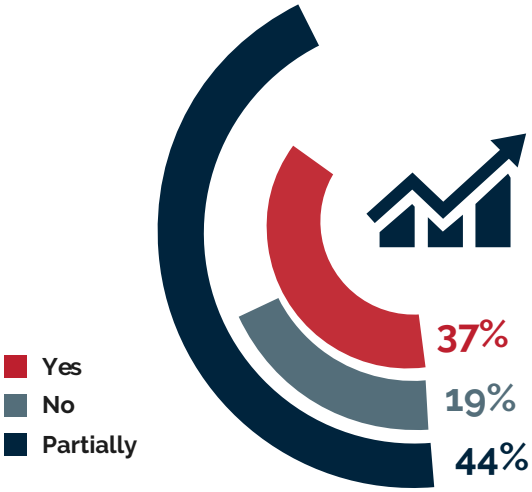

CHART 4
What technology supports your:

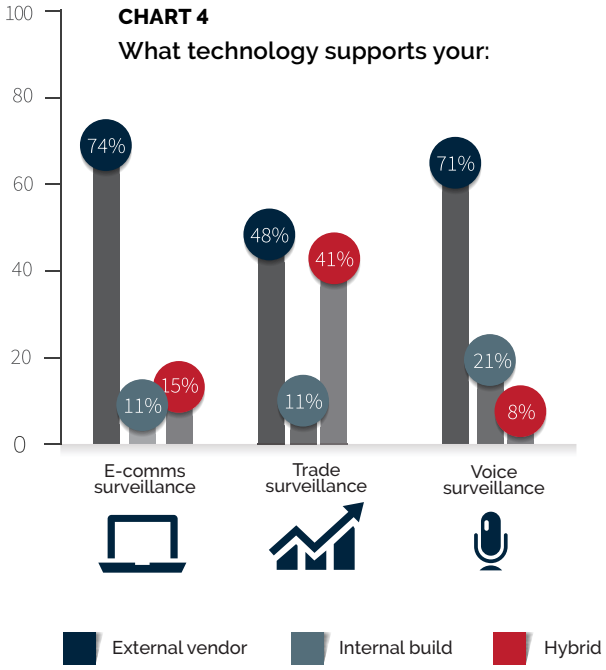

Copyright 2026, 1LoD. All rights reserved

A vendor opportunity?

Trade surveillance is also an outlier in terms of technology sourcing. While almost three quarters of survey respondents now buy voice and e-comms solutions off-the-shelf from third parties, with few taking a hybrid approach, in trade fewer than half of the respondents buy from third parties alone and 41% adopt a hybrid approach.

At the core of this hybrid persistence is the issue of customisation. Multiple respondents point to the fact that vendor solutions rarely deliver complete coverage out of the box. As one head of trade surveillance at a large US bank put it, "there's an awful lot of customisation that has to occur… that's what people are calling hybrid," while another explained more bluntly that vendor platforms "can't provide full coverage… [because of] the lack of customisation… [and] the instruments that they cover." This gap is not marginal but structural: firms consistently find that no single provider meets all their requirements.

One head of surveillance described a typical outcome: "one vendor might give us 80%… another gives 20%… we just didn't find anyone that offered that 100%," leading firms to assemble a composite architecture that blends vendor tooling with internal enhancements. In practice, this results in firms "running both"—using vendor systems for baseline coverage while layering proprietary models on top, particularly for complex behaviours such as spoofing and layering.

For some institutions, this necessity evolves into a strategic preference for in-house development. A small but committed set of banks have taken this route, emphasising the advantages of control, speed, and integration. The head of trade surveillance at one of them explains, "with sufficient investment in data and expertise, the flexibility of what you are able to do… the effectiveness of your models… [means] you are able to quickly pick up things… way, way faster than any vendor," while also embedding surveillance more

CHART 5

What level of investment is your organisation planning in the following areas over the next 24 months?

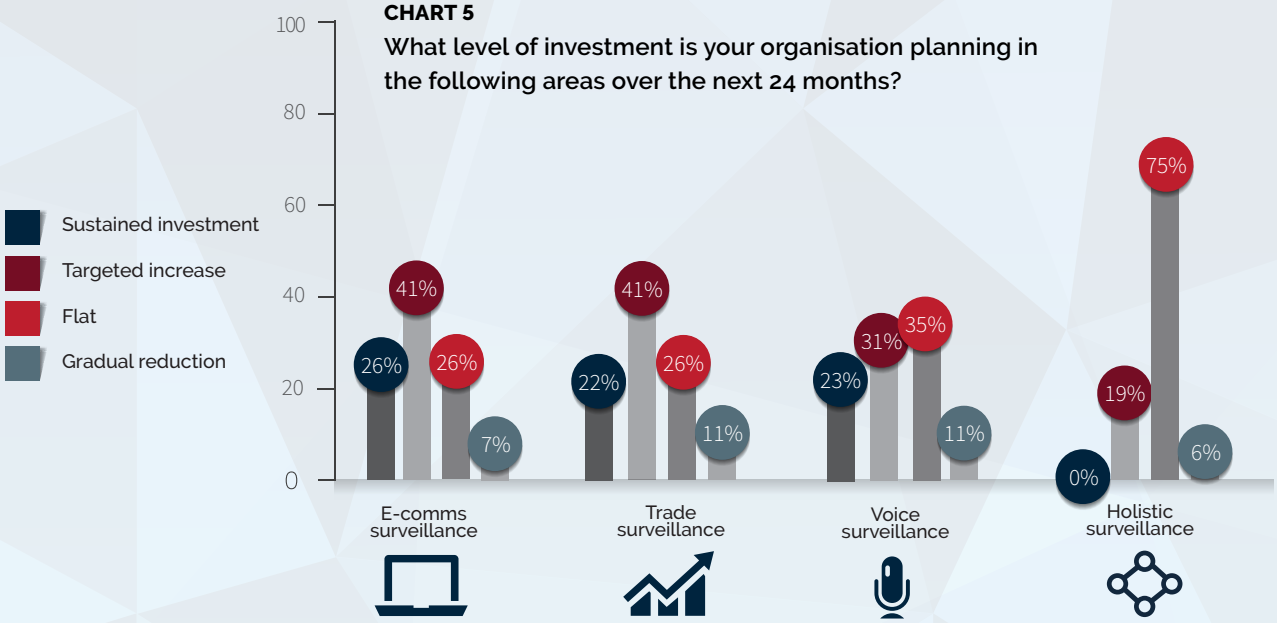

deeply into the broader conduct risk framework. "If you build in-house, it is easier to tell the story internally… in terms of your risk and control coverage and effectiveness."

This perspective highlights an important asymmetry: for firms that have already built the necessary infrastructure, in-house capability is not just viable but potentially superior. However, the same interviews make clear why this model is not widely replicable. The barriers to entry are significant as one survey respondent explains: "the upfront investment… to get the right market data, the reference data, the right models… is going to be huge"—and many firms are constrained by legacy vendor contracts and organisational inertia.

At the same time, there is growing evidence of a countervailing trend: a gradual retreat from fully in-house models as vendor capability improves. One head of trade surveillance making the transition says, "managing and building… [an] in-house trade solution isn't the long-term optimal strategy."

Regardless of which way they choose to go, it seems as though banks will have the money to do it **[Chart 5]**. 63% of respondents predict either sustained or targeted investment in trade surveillance over the next 24 months.

Running to standstill in e-comms

In e-communications surveillance **[Chart 1B]**, the headline picture is one of stasis. Between 2024 and 2026, the proportion of respondents who report being only "partially satisfied" with their existing solutions remains unchanged at 52%, despite a modest two-point increase in full satisfaction and a corresponding decline in outright dissatisfaction.

For a function that has seen sustained vendor investment in NLP-driven lexicon expansion, contextual clustering, and—more recently—LLM-assisted semantic enrichment, this result could be seen as a failure of all the new investment [largely in third-party solutions **[Chart 4]**.

However, respondents make clear that in fact this is not a failure of technology so much as a reflection of a rapidly expanding and increasingly complex problem space. On the one hand, there is clear acknowledgement that vendor capability has materially improved in recent years. As one UK-based surveillance head notes, "there have been considerable enhancements in terms of vendor capability on the comms side," particularly with the incorporation of LLMs and more advanced pattern recognition. Yet these gains are being offset by the fact that firms are now attempting to monitor a far broader and more fragmented communications

Copyright 2026, 1LoD. All rights reserved

CHART 1B
Are you satisfied with your current E-comms surveillance solutions?

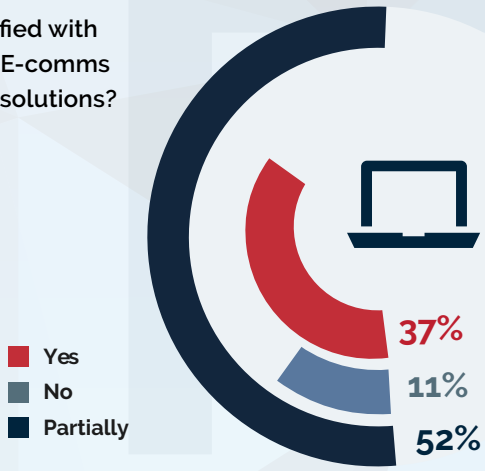

CHART 11
How many lexicons are you using?

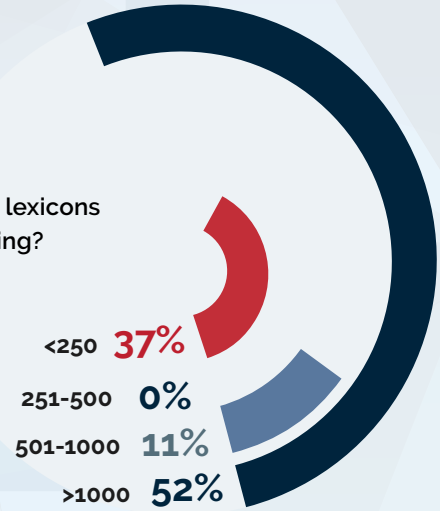

landscape, meaning that improvements in tooling are effectively being absorbed by the increasing scope of surveillance itself.

A central issue underpinning this dissatisfaction is that data completeness and capture remain unresolved, regardless of how sophisticated the analytics layer becomes. Even where firms have invested heavily, the underlying challenge persists: surveillance is only as good as the data it ingests.

As one survey respondent highlights, firms are still "grappling with unsupported communication platforms… encrypted messaging tools… [and] reconciliation processes," with surveillance teams increasingly drawn into complex data engineering tasks simply to ensure basic coverage. Regulatory pressure has reinforced this weakness, with "inadequacies around [data completeness] resulting in some fairly hefty fines," forcing firms to prioritise capture and governance just as much as detection capability. The result is that even as systems improve, confidence in their completeness—and therefore their effectiveness—remains constrained.

It is also the case that while new AI-driven solutions are at the beginning of a new era in e-comms surveillance, they are not mature. More than half of respondents are still using more than 1,000 lexicons **[Chart 11]** which shows that confidence in the new technology is not yet fully established.

As one head of e-comms surveillance puts it, "people just can't yet get comfortable with one or the other… so they're probably using a bit of both," creating duplication, inconsistency, and a lack of clear methodological conviction.

That said there are some signs in the data that lexicon usage has turned a corner. In 2024, 74% of respondents were using between 501 and 1,000 lexicons. That has fallen to 63%. And 23% were using fewer than 250, which has risen to 37%. Overall, therefore the number of people using fewer lexicons than two years ago has risen.

The final reason satisfaction levels may have stalled in e-comms surveillance is that more advanced analytics can uncover a greater number of potential issues, which creates more, not less, work for human review teams. As one survey respondent says, "better technology can generate higher detection volumes… requiring careful resource planning," meaning that human review capacity remains a binding constraint. At the same time, regulators have raised the bar for what constitutes effective surveillance, requiring firms to demonstrate not just the existence of controls but their completeness, accuracy, and coverage across languages and channels.

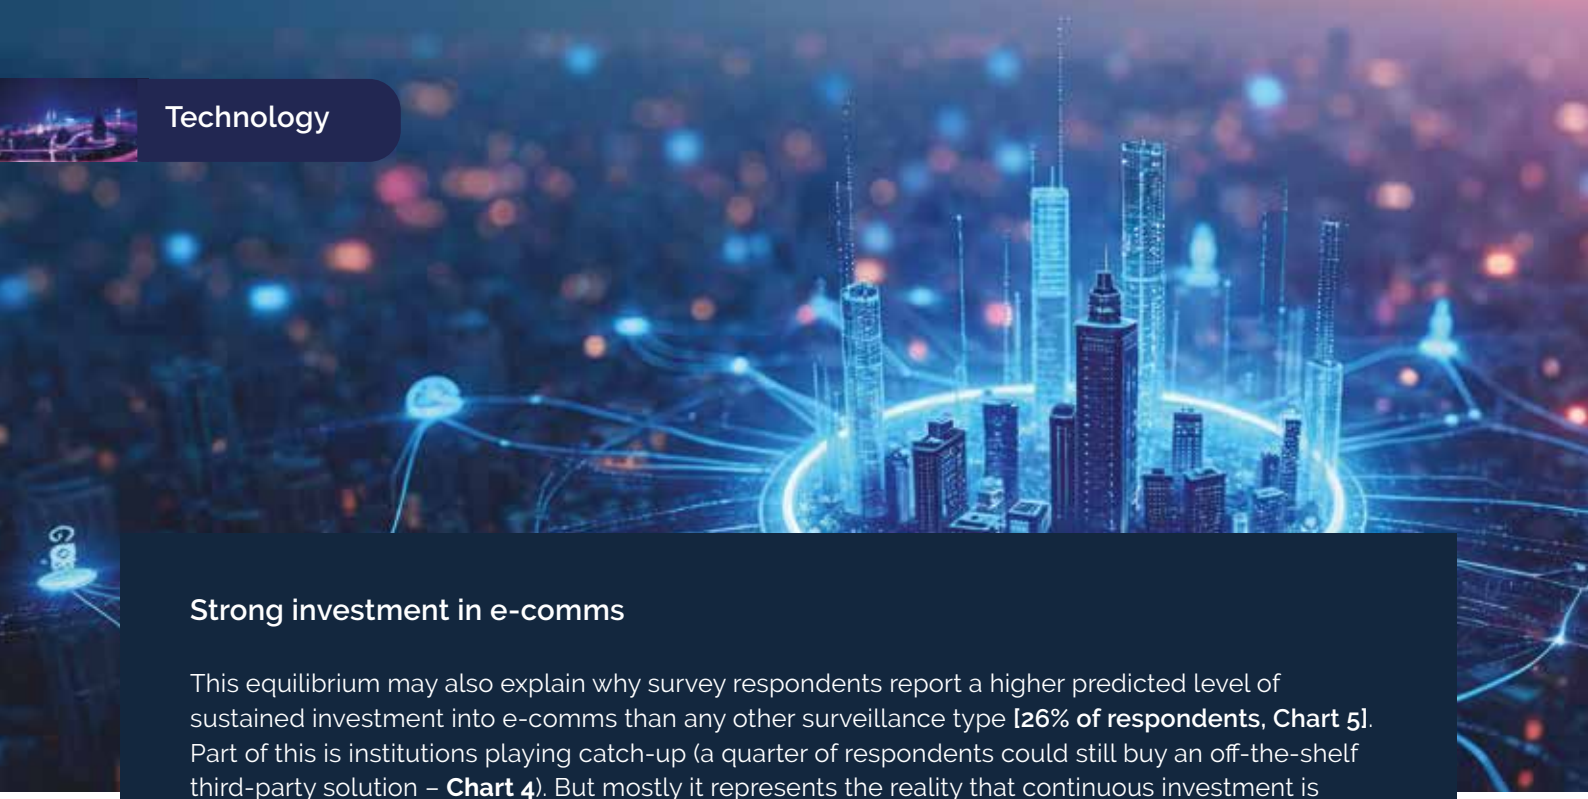

Strong investment in e-comms

This equilibrium may also explain why survey respondents report a higher predicted level of sustained investment into e-comms than any other surveillance type **[26% of respondents, Chart 5]**. Part of this is institutions playing catch-up (a quarter of respondents could still buy an off-the-shelf third-party solution – **Chart 4**). But mostly it represents the reality that continuous investment is required simply to keep pace with increased channels, market changes and regulatory requirements.

Leading the way in third-party solutions

As has already been noted, new technology is easiest to justify when its ROI is easy to demonstrate. Because language is where AI has excelled to start with, AI-augmented solutions have allowed banks to keep up with the explosion in the number of communications channels they are required to monitor and the number of languages regulators now want included in surveillance without having to significantly increase the size of human teams.

As one head of EMEA surveillance at a very large Asian bank says, "AI in the comms space is the best use case to extract both effectiveness and efficiency and in fact do things that you really couldn't do at scale without it."

No surprise then that 74% of respondents rely on third-party vendor solutions in this space, and the banks who have built in-house solutions look increasingly eccentric. In fact, when asked why they have taken that route, they tend to point to an engineering or IT-dominated culture, rather than any benefits in terms of customisation or flexibility. "We are a bank run by engineers who like to build stuff in-house," is a typical comment from one European surveillance head.

They do point to the cost of third-party solutions, implying that in-house builds are cheaper. But since almost no-one at any large organisation in any sector can actually calculate the true cost of ownership of any piece of IT, third-party or in-house, it's not obvious that going in-house does actually save money.

Copyright 2026, 1LoD. All rights reserved

Big rise in voice satisfaction

What is true for e-comms is also largely true for voice. Again, the core problem is based in language analysis – transcription, translation and analytics. So here too, vendor solutions now dominate the space, with 71% of banks using them **[Chart 4]**. Banks agree that they are getting much better at providing consistent and scalable capabilities, while advances in transcription, multilingual processing, and search have reached a point where improvements are both tangible and operationally meaningful. As one participant noted, firms are "beginning to see the fruits of those ambitions coming through."

These advances explain the huge rise in respondent satisfaction with their voice surveillance activities. Two years ago, just 26% were satisfied and 44% were partially satisfied. In this year's survey 48% are satisfied, just 28% are partially satisfied and the share of those who are not satisfied has remained static **[Chart 1C]**.

A key driver of this improvement is the extent to which voice has become a standardised, vendor-led domain. As one respondent notes, "in both e-comms and voice you buy an external solution… [and] you are wildly an outlier if you don't," highlighting the degree of convergence around third-party platforms.

There are still issues, of course. While AI has undoubtedly improved capability, participants are clear that it has not yet solved the problem. As one respondent puts it, "as you go across multiple languages… transcription is getting better, but it's not there yet."

Many banks also still report poor recording quality as a major stumbling block to improving voice surveillance results. Noisy audio, overlapping speech and generally poor line quality fundamentally limit detection accuracy.

Others struggle with a broader issue – that of assigning a unique identifier to an individual that allows them to be tracked across the different surveillance processes. Just within voice, as one surveillance head at a large US bank says, "We still haven't really solved the problem of implementing a unique identifier across all voice channels that definitively prevents duplication or misidentification." But the issue goes beyond voice as another US surveillance chief explains, "How you are uniquely identifying [a person] across every single one of your platforms to make sure you're assigning voice and e-comms to the right person all the time – that is a challenge."

The answer is a mix of old-fashioned (surely the line identifies the voice?) and new (AI is pretty good at voice identification – though not perfect). But there is still a lot of work to do. Fortunately, for 54% of organisations, budgets for voice surveillance will rise over the next two years and only 11% see any chance of even a gradual reduction **[Chart 5]**.

CHART 1C
Are you satisfied with your current voice surveillance solutions?

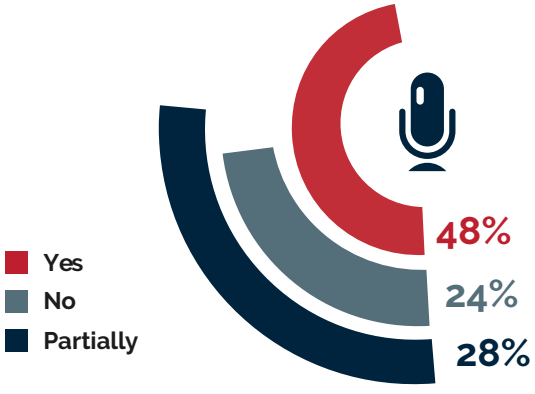

CHART 2
Are you satisfied overall with your case management / workflow tools?

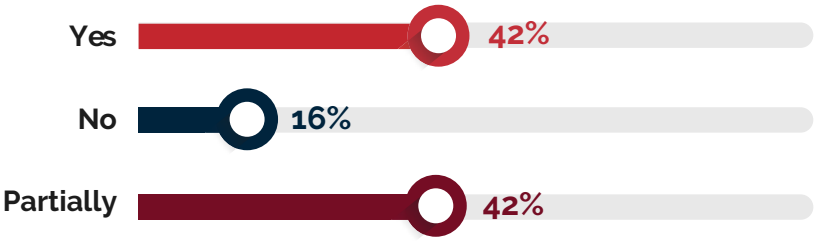

CHART 3
Have you bought or developed your own case management / workflow tools?

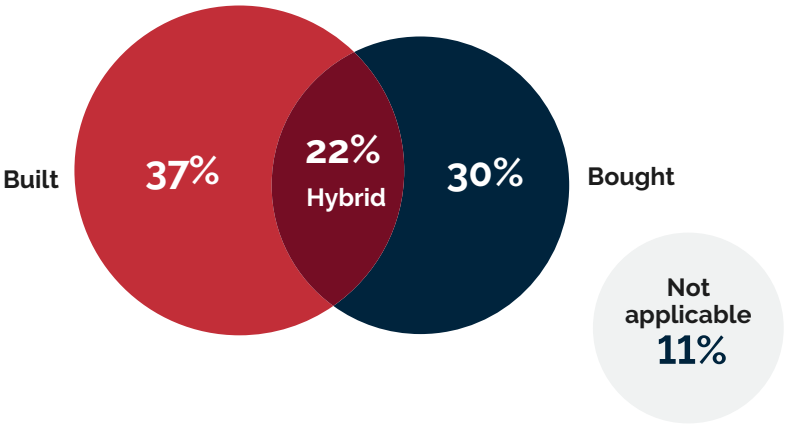

CHART 8
Do you anticipate moving your technical architecture from on-premises to the cloud in the next 24 months?

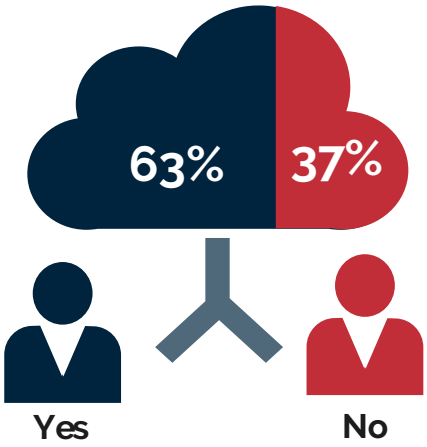

In-house workflow solutions at end-of-life?

Satisfaction with workflow and case management systems has deteriorated over the past two years and it's not hard to see why. In 2024, 52% of respondents were satisfied overall with their case management / workflow tools, with just 4% not satisfied. Two years later, just 42% are satisfied and 16% are not **[Chart 2]**.

Copyright 2026, 1LoD. All rights reserved

Satisfaction appears to have declined not because those platforms have deteriorated, but because the nature of surveillance itself has changed around them. Historically, these systems were designed to support a world of limited channel monitoring and relatively contained alert volumes. However, the current surveillance environment has become more complex, and firms demand more from case management systems in terms of better aggregation of data across voice, e-comms and trade, and in terms of smart analytics in the investigative process. Also, better alert tooling has increased the number of alerts worthy of investigation, and those volumes put pressure on case management systems.

It's also true that as surveillance has expanded, and as regulators have focused more heavily on data quality and completeness, gaps in data have become a much larger issue over the past two years. Those gaps often emerge during investigations and place additional strain on case management systems. It may also cause teams to blame the workflow tools for inadequacies whose root cause is far upstream of those systems.

Finally, increasing regulatory scrutiny has raised the bar for documentation, auditability, and defensibility. Firms are expected to evidence not only what decisions were made, but why they were made, what data was considered, and how complete that data was. This has expanded the documentation burden within case management systems, often without a corresponding improvement in usability or automation. At the same time, AI has begun to introduce capabilities such as automated triage and alert prioritisation, further highlighting the limitations of traditional workflow tools. In effect, expectations have moved ahead of delivery.

Interestingly, it is in case management and workflow tooling that reliance on in-house solutions is highest **[Chart 3]**. It is tempting to suggest that those falling satisfaction levels are linked to the use of tools that have been cobbled together in-house over the years and which now cannot cope with the modern surveillance environment. It is certainly true that banks are most likely to defend their in-house workflow tooling on the basis that they understand best the fragmented datasets and required connectors and the investigative processes they want followed. It is increasingly hard to agree with this point of view in a world of third-party data aggregation experts, AI, Application Programming Interfaces (APIs) and high-quality workflow and process management tooling.

The use of in-house tools certainly runs counter to the trend in cloud adoption evident from **Chart 8**. Asked whether they anticipate moving technical architecture from on-premises to the cloud in the next 24 months, 63% say 'yes', 37% 'no'. Given that two years ago it was 91% yes and 9% 'no', it might appear that banks are shifting dramatically away from the cloud just as vendors drive even further towards Software-as-a-Service (SaaS) and Managed Service offerings. But that is not the story of this data.

Instead, the story is one in which an increasing number of banks are moving more and more of their surveillance processes to third parties whose software and processing is itself increasingly cloud-based. Just the use of LLMs and AI-driven solutions usually implies the use of cloud-based solutions.

So, the rising number of 'nos' and declining number of 'yesses' simply reflects the fact that many people have already moved what they want to the cloud and do not plan any further moves.

Here again, vendors seem to be leading the market – and not just with SaaS solutions. The larger providers are introducing fuller surveillance-as-a-service solutions that include alert triage using a combination of AI and human SMEs. This is seen as too big a leap for some of the largest and most sophisticated institutions and will likely appeal first to smaller, mid-market firms and the buy-side.

The head of trade surveillance at one large US investment bank is typical of their peers when they say, "We would absolutely never do that [use a managed service provider for surveillance that triaged alerts]. We would regard that absolutely as a loss of control of a core process. And there is a regulatory angle: just using a vendor at all is considered outsourcing and getting auditors and regulators comfortable with it a challenge. How you would do that with a managed service I don't know."

CHART 38

Which of the following data-related issues most hinders the effectiveness of your surveillance systems?

24%	24%	23%	12%	9%	4%	4%
Fragmented data across silos or systems	Lack of standardised data formats or identifiers	Inconsistent or poor-quality data	Insufficient data governance or lineage visibility	Difficulty linking structured and unstructured data (e.g. chat, voice, trade logs)	Limited real-time data processing capability	Other

Surveillance? Or data engineering?

The single biggest issue in surveillance raised in 1LoD's Leaders' Networks and in conversations with surveillance heads across trade, voice and e-comms is still data. And while banks have invested heavily in both enterprise data architectures and surveillance data capture, archiving and cleaning, the problem is still as pressing as ever. This partly is because the challenges are inherently difficult, and partly because of the explosion of surveillance data volumes caused by increased market activity, the expansion of voice and e-comms into new channels and languages, the ability of technology to capture previously unmonitorable activities or datasets and regulatory pressure.

So, when asked 'which of the following data-related issues most hinders the effectiveness of your surveillance systems', **[Chart 38]**, the surveillance leaders paint a sobering picture of the current state of surveillance data maturity.

Rather than pointing to a single dominant challenge, the results reveal a cluster of closely related issues at the top: 24% of respondents cite fragmented data across silos or systems, a further 24% highlight a lack of standardised data formats or identifiers, and 23% point to inconsistent or poor-quality data. Taken together, these findings suggest that the core problem facing surveillance teams is not one discrete weakness, but a systemic breakdown in the underlying data architecture. Fragmentation, inconsistency, and lack of standardisation are not separate challenges—they are interdependent symptoms of a deeper structural issue.

The combination of fragmentation and lack of standardisation is why fixing the data supply chain is so difficult. Fragmentation in itself is not insurmountable; many large institutions operate across multiple systems and datasets. However, when those systems lack common identifiers, consistent schemas, or aligned taxonomies, the ability to integrate and reconcile data breaks down.

The prominence of poor data quality as a reported issue is as likely to be a downstream effect of the more fundamental problems of fragmentation and standardisation as it is to be a cause of problems in itself. Data quality issues—such as missing fields, duplication, or incomplete records—are often the result of disparate systems feeding into surveillance environments without consistent controls or alignment. In other words, poor data quality is less a

Copyright 2026, 1LoD. All rights reserved

CHART 28

What percentage of employees in your surveillance function are working exclusively on Change the Bank (CTB) projects?

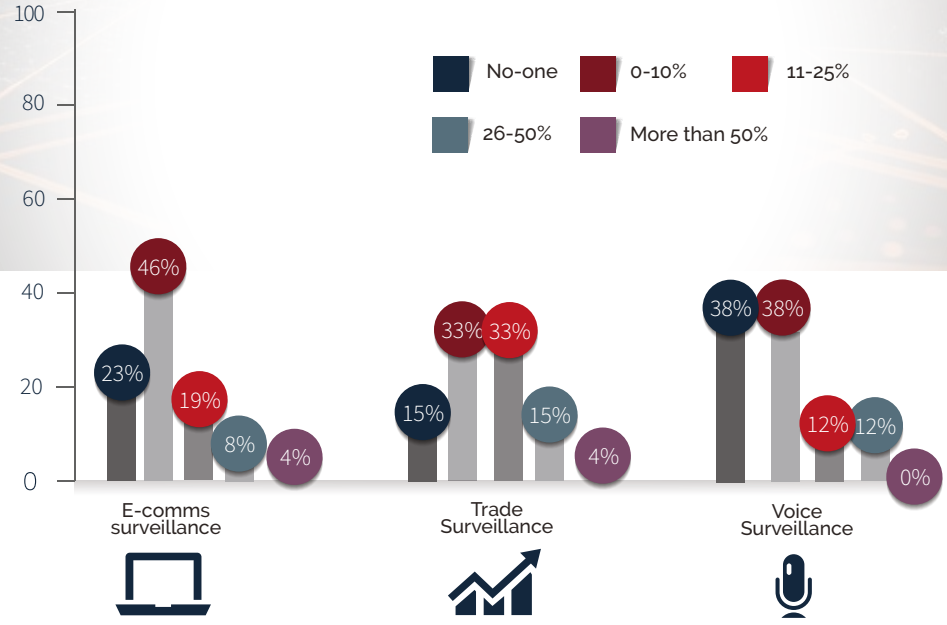

root cause than a visible symptom of a fragmented and poorly standardised data landscape.

Interestingly, only 9% of respondents explicitly identify the difficulty of linking structured and unstructured data—such as chat, voice, and trade logs—as a primary issue. However, this figure arguably understates the scale of the problem. The ability to link these datasets is inherently dependent on resolving the higher-ranked challenges. Without consistent identifiers, clean data, and aligned formats, meaningful integration is simply not possible. As such, the difficulty of linking datasets is embedded within the top three issues, rather than standing apart from them.

Overall, the chart reinforces a broader conclusion that emerges strongly from interviews and discussions with surveillance heads: the primary constraint on surveillance effectiveness is not a lack of analytical capability or technological innovation, but the weakness of the underlying data foundation. Until firms can address fragmentation, establish consistent data standards, and improve data quality at source, more advanced ambitions—whether holistic surveillance, AI-driven detection, or real-time monitoring—will remain difficult to realise in practice.

The transition resourcing problem

With surveillance at a key transition point, the question of who funds and runs the transition becomes critical. Here, it appears that resourcing may still be an issue. Asked "what percentage of employees in your surveillance function are working exclusively on Change the Bank (CTB) projects?", respondents paint a picture of an industry in which almost no institution has meaningfully ringfenced capacity to actually deliver that transition toward a more technology-led, AI-enabled operating model.

In 2024, 22% of firms had no dedicated CTB resource at all. While 50% had between 0% and 10% of staff on change initiatives. Only 6% had anything approaching a meaningful (25% to 50%) allocation.

This year's survey reinforces this picture **[Chart 28]**: across e-comms, trade, and voice, the dominant model is still sub-10% allocation and "more than 50%" is effectively non-existent (0% to 4%). So, the overwhelming majority of institutions are still operating with marginal CTB allocations, and almost none have ringfenced transformation teams at scale.

In practical terms, this means that the industry is attempting to deliver a fundamental operating model shift without creating the organisational capacity required to execute it. The transition is not being driven by a deliberate, well-funded transformation effort—it is emerging slowly, unevenly, and largely by necessity, carried on the shoulders of already overstretched teams.

There is an irony here. The relentless pressure to expand surveillance coverage that is driving the transition to a more efficient technology-led model is itself the reason for the short-term rise in BAU burden that then sucks resources away from change-the-bank programmes. As ever more time and headcount are consumed by alert review, monitoring, and coverage obligations, CTB activity is pushed further into a side-of-desk activity.

There are also subtle differences across surveillance domains that reinforce this interpretation. Voice surveillance, which is relatively less mature and still undergoing active build-out, shows slightly higher levels of mid-range CTB allocation, suggesting that where capability is still being established, firms are more willing to invest in change resource. Trade surveillance, by contrast, appears to be in a phase of incremental optimisation, with modest but not transformative investment. E-comms surveillance shows the least capacity for change, reflecting the sheer weight of ongoing coverage requirements. In effect, the more established a surveillance function becomes, the more it is locked into BAU and the harder it becomes to transform.

As this head of surveillance at a large US bank describes the challenge, "We're quite slim as an organisation … it's a juggling act".

Vendors stepping up

So, how is this situation to be resolved? No-one spoken to for this survey appeared to believe that large transformation budgets would be made available in the absence of a large regulatory action or investigation. This leaves an opportunity and a challenge for the vendor community. Vendors are already carrying an increasing share of the innovation burden—developing AI models, analytics capabilities and integration frameworks that banks then adopt more slowly. It seems as though they will also have to take a more active role in driving CTB activities too.

Ultimately, the charts suggest that the industry is trying to execute a generational shift in surveillance capability without stepping out of its existing operating model. That tension—between strategic ambition and resourcing reality—is likely to define the next phase of surveillance evolution.

CHART 34

What is your planned investment spend in Change the Bank (CTB) surveillance projects over the next 12-18 months?

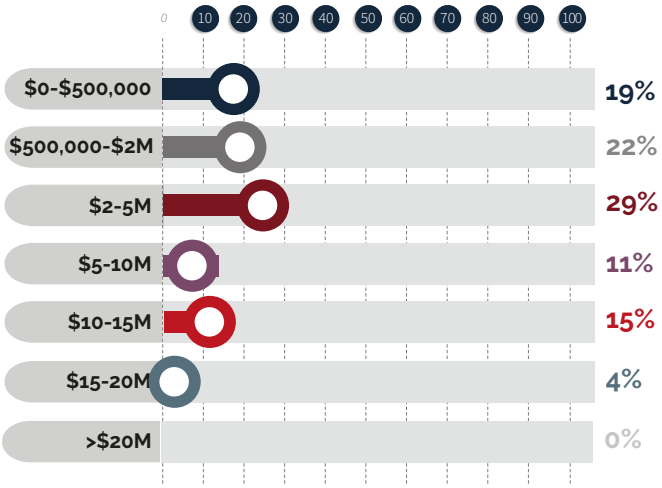

Copyright 2026, 1LoD. All rights reserved

Enough investment in CTB?

This slightly pessimistic picture is somewhat countered by the levels of investment going into CTB surveillance projects. **Chart 34** shows the answers to the question "What is your planned investment spend (US$) in CTB surveillance projects over the next 12-18 months?". Whereas in 2024, the distribution of answers was roughly flat ($0–500k, $500k–$2M, and $2–5M all sat at around 22%) this survey reveals a concentration of spend in the $2M–$5M range (29%). But respondents investing more than $5M is 30%, and those investing $10M or more is 19%.

This reflects a division in the market between larger, high-investment leaders ($10M+ to spend) whose spending levels look realistic if they are to meet the challenges of developing AI-enabled surveillance, data transformation and significant platform upgrades. A second group ($2M to $10 million in spend) comprises larger banks with more constrained budgets and mid-sized banks with budgets that are appropriate to their scale. Here, the larger banks are limited to incremental development, not fundamental re-architecting. The mid-sized banks have budget for what for them are significant upgrades and additions. The final group (<$500 to $2M) are highly constrained and have little transformation capacity beyond tactical fixes.

The result is a transformation that is both real and uneven - advancing rapidly in some parts of the market, while progressing more gradually in others, and shaped as much by institutional scale and complexity as by strategic intent.

Who holds the purse strings?

As for who is making the technology buying decisions, the answer, as vendors will know only too well, is a committee of sorts **[Chart 35]**. It's no surprise that plenty of people get involved, but it is a surprise that so many people without a technology background may have a say. This aligns with conversations at 1LoDs Surveillance Leaders' Networks around AI adoption and experimentation (and also digital assets) which give the impression that while banks frequently describe themselves as 'technology companies', and acknowledge their absolute dependence on technology, many still run technology innovation within silos on a side-of-desk basis (like CTB projects) run by people who simply expressed the most interest in the particular challenge.

The technology transformation required in surveillance (and banking more generally) requires a shift in skillsets, or a shift in the way that technology is managed within businesses and functional areas like compliance.

CHART 35

Who are the enterprise wide stakeholders involved in external technology buying decisions at your banks?

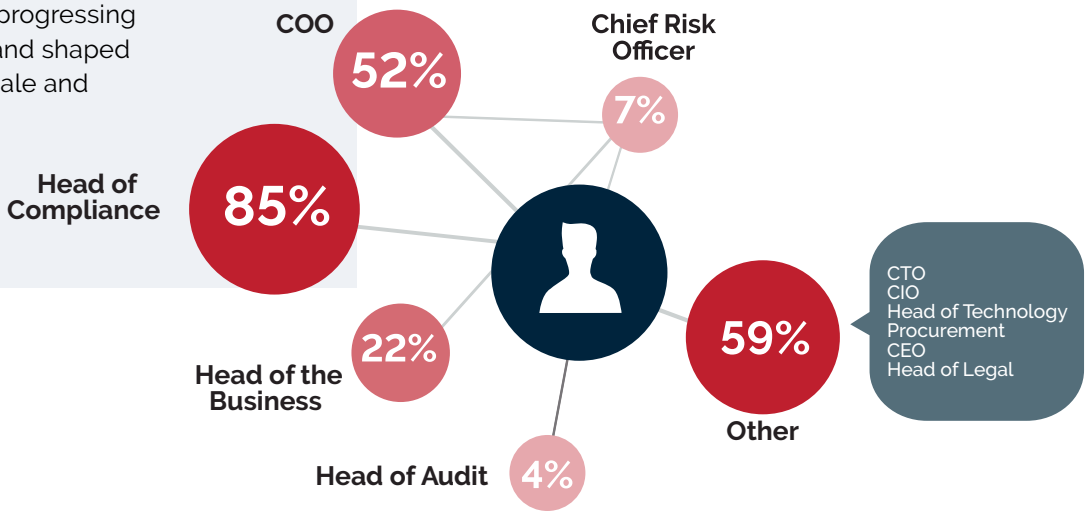

AI adoption on the rise

AI may be the hottest topic in the media, and in the self-interested worlds of venture capital and private equity, but when it comes to real businesses shipping fully formed AI-augmented solutions into real environments, then the picture is different. That is true in healthcare, manufacturing and financial services. And it is certainly true in market abuse surveillance.

As **Chart 6** makes clear, when it comes to having AI solutions 'fully embedded and operating as part of a target operating model', the numbers are very small: none in trade surveillance, none in holistic surveillance, 8% in e-comms surveillance and 9% in voice surveillance – both areas where LLMs should excel.

"We've got multiple AI initiatives running, but very few that you would call truly operational," is a typical response from, in this case, a global head of surveillance at a European bank. This gap between activity and operationalisation defines the current state of the industry. AI is no longer theoretical—but nor is it yet transformative at scale.

However, the chart maps the trajectories banks are taking on the way to that final goal, and here the data is more nuanced. 70% of respondents have e-comms solutions in POC or active deployment. In voice, the number is 68%. In both cases a further quarter of respondents are in the research and development (R&D) phase. So, just where you would expect AI to have made the greatest strides, it is doing that. And a number of respondents have recently onboarded third-party solutions.

"We've implemented an AI-enabled solution in e-comms globally and that has been a big success story for us," says the surveillance head of a large Asian institution. "One immediate benefit has been in being able to meet new regulatory challenges over channels and languages. Foreign languages are a key thing for us and being able to use AI for that has enabled us to really increase our coverage."

In voice it's a very similar picture. More than a third of respondents are in full build, with a third in POC. As with e-comms, the main drivers of adoption are the technology's direct application to language analysis, and the need to expand coverage and prove monitoring of all regulated populations.

CHART 6
What best describes the state of AI adoption in your surveillance function in the following areas?

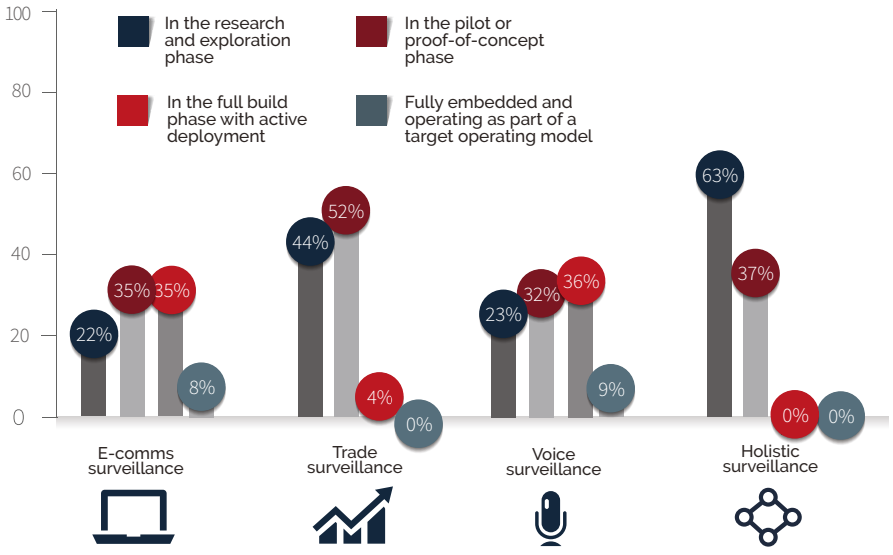

Copyright 2026, 1LoD. All rights reserved

Trade lags behind

The picture in trade is very different: 44% of respondents are still in the R&D phase, with 52% in POC. Just 4% are in full build with no-one fully embedded.

So why the difference? Unlike e-comms or voice, trade surveillance has always been model-driven with scenarios underpinned by complex rules and statistical thresholds. Many of these scenarios are now treated as complex models by banks and regulators and many models already incorporate sophisticated ML. It's not immediately obvious that the LLMs driving most of the AI hype cycle are that useful in the structured data, quantitative analytics world of trade surveillance. And, again unlike with voice or e-comms, introducing it means very significant changes to existing models, which then need to be revalidated, and to the rest of the process too. This can make AI seem like a bet on an unknown and possibly incremental improvement with disproportionate implementation cost and risk.

Governance and explainability issues are also much more critical in trade surveillance as it derives directly from very specific regulatory requirements. So, if AI is seen as introducing opacity or explainability issues, then adoption will be much slower.

Perhaps most important of all though, the main challenges in trade surveillance – false positives, cross-market / cross-product abuse, prevention rather than detection – cannot be solved simply by overlaying AI onto existing datasets or alert logics. The next breakthroughs in trade surveillance will be when it is possible to shift from largely static, rule-driven monitoring to adaptive, 'smart' systems that work on broader datasets to identify patterns and behaviours rather than simply spitting out alerts based on datapoints that trigger scenarios.

But that requires a data revolution in banking that reaches far beyond surveillance, and it requires the return of a focus on holistic surveillance. And as the chart shows, that is not evident in banks' thinking yet.

Aspiration versus ambition

These findings carry through into the availability of AI-augmented capabilities within core surveillance solutions **[Chart 7A]**. AI-enhanced trade surveillance is present in just 11% of firms, yet 89% say they want it. In e-comms and voice surveillance—adoption sits at only 26%, with two-thirds of firms still aspiring to deploy it. Generative AI assistants for analysts are even less mature, with only 7% reporting current capability against 78% expressing demand.

Who needs holistic surveillance?

AI-enabled holistic surveillance is the only area where no respondents are either in build or embedded. More than 60% are still in the R&D phase with the rest in POC. This makes sense **[Chart 6]**. If coverage is still a challenge, and if building out accurateness and completeness in data quality, then the foundations for holistic surveillance do not exist – let alone the technology. In that case banks are right to focus on the easiest use cases in which to see how AI implementation and integration work, before moving on to more ambitious projects.

The lack of holistic surveillance also reflects a more general lack of prioritisation of the whole topic. As **Chart 7B** shows, just 4% of firms report having such functionality, while 59% say they do not have it and do not require it. This is despite the long-standing narrative around holistic surveillance and the clear risks associated with behaviours that span multiple channels and asset classes. Regardless, vendors are pushing ahead with holistic surveillance platforms, so this may be yet another area where third-party solution providers drive the market forward rather than the banks.

CHART 7A
Which features are available in your surveillance platforms?

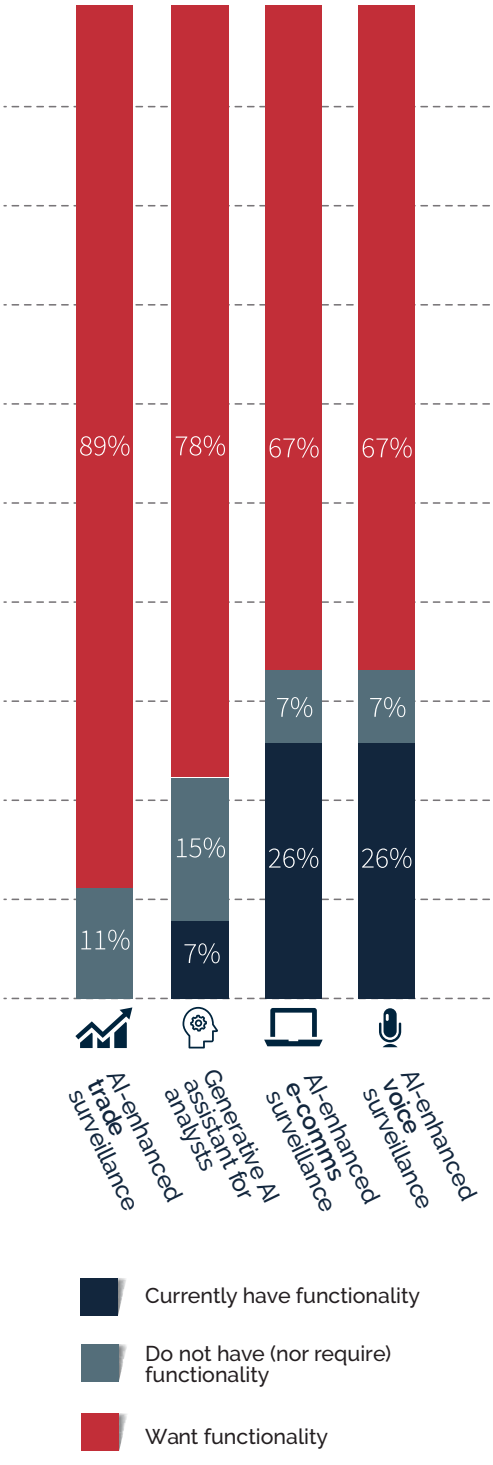

Copyright 2026, 1LoD. All rights reserved

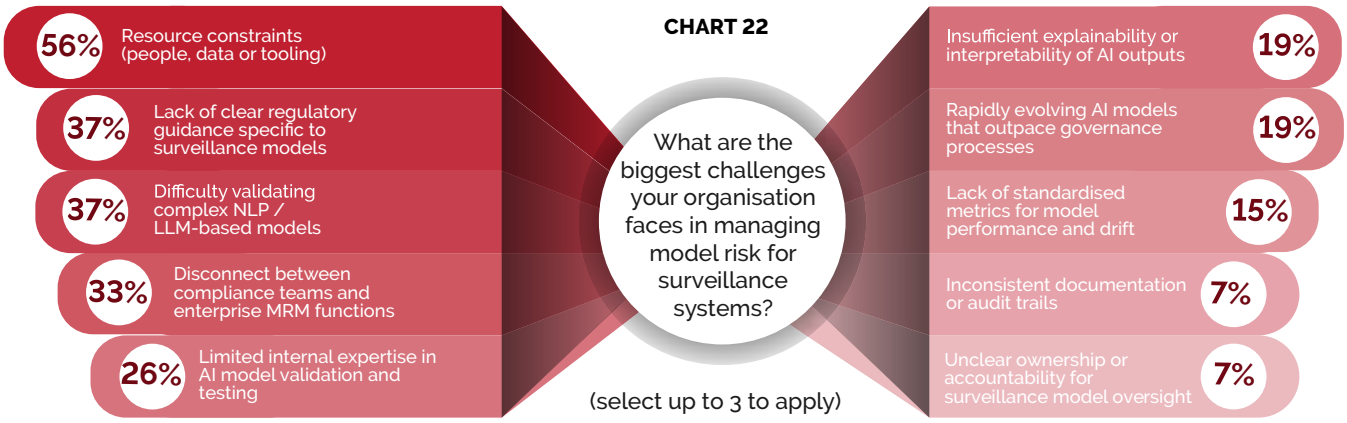

Explainability still an issue

Governance and explainability issues are also still significant obstacles to the broader adoption of AI-augmented surveillance tools.

This is most critical in trade surveillance as it derives directly from very specific regulatory requirements. But it is true beyond trade too and evidence for this comes from the clash between model risk management (MRM) teams, now all-pervasive in surveillance, and AI models **[Chart 22]**.

The survey asked what the biggest challenges surveillance teams face in managing model risk for surveillance systems, and of the top seven issues, four are AI-related: difficulty validating complex NLP-/LLM-based models, limited internal expertise in AI model validation and testing, insufficient explanability or interpretability of AI outputs and rapidly evolving AI models that outpace governance processes.

Of these the most fundamental is the 37% struggling with validation. In rules-based, deterministic surveillance models validation is well understood: you can backtest, you can define thresholds, you can evidence behaviour. With LLM/NLP models, you lose that determinism. Outputs are probabilistic, context-dependent, and often non-reproducible in a strict sense. This makes validation – not strictly explainability – the real issue: yes, if you cannot reliably explain why a model produced an alert, you cannot defend it to compliance, regulators, or 2nd line.

But the underlying issue is that the model itself does not behave in a way that fits existing control frameworks. MRM teams are effectively trying to retrofit AI into governance structures designed for linear, static models. Those structures assume stable model behaviour, clear feature attribution, repeatable outputs and well-defined performance metrics. But AI models do not behave like this and what is needed is a fundamental reappraisal of model evaluation and validation by both banks and regulators.

That said, vendors increasingly understand the banks' need for explainability and accept that, as one chief revenue officer at a market-leading solution says, "transparency around our models and how they work, full documentation and willingness to help banks explain solutions to regulators are now just table takes for vendors who want to sell to the surveillance community."

However, the combination of challenges related to AI revealed in the survey suggests that it will take a significant drive by banks as enterprises, as well as by vendors and the regulators, to accelerate AI usage across the surveillance ecosystem.

Section iv: **The future of surveillance**

Almost a quarter of banks believe their market abuse surveillance infrastructure does not provide effective risk management of market abuse risk

Copyright 2026, 1LoD. All rights reserved

The future of the surveillance function depends largely on how banks resolve the challenge of ever-expanding BAU obligations while realising an ambitious, but resource-constrained transformation agenda.

On one side, surveillance leaders remain consumed by the fundamentals—extending venue coverage, onboarding new communications channels, scaling language capabilities, and keeping pace with the growth in regulated populations.

On the other, there is the unavoidable need to embed AI and more sophisticated analytics at the core of the function, not simply to optimise alerting but to fundamentally reframe surveillance as a form of business intelligence. This in turn requires investment in enterprise data quality and related technologies beyond the control of surveillance functions.

How individual banks balance these two sets of challenges will determine their future surveillance operating models, in terms of staffing, structure, technology, efficiency and effectiveness.

BAU budgets healthy

Money, as always, will play a big part and banks seem to believe that as far as BAU is concerned, they have enough. Asked 'do you have sufficient budget to build and maintain a surveillance function capable of effectively managing the risk of market abuse within your bank?', 78% say 'yes' **[Chart 36]**.

We have seen **[Chart 5, Section iii]** that many banks report rising budgets across e-comms, voice and trade – much of that increase described as 'targeted'.

This is confirmed by **Chart 33** which reveals that 44% of respondents expect their available budget for investing in market abuse surveillance to rise this year (up from 33% in 2024), with 30% (up from 11%) anticipating a decrease and 26% (56% in 2024) believing budgets will remain constant.

So, although in 2024 82% believed budgets would be flat to up, and that has fallen to 70% in 2026, the numbers seeing an increase are sharply up.

The splits here largely reflect maturity (as is also the case with headcount – see below). The largest, most advanced banks predict falling budgets as their transformation spend is largely finished and they start to see the rewards in terms of lower ongoing budget requirements and lower staffing.

For many banks though, that transformation spend is still to come and they are either still wrestling with regulatory demands or upgrading legacy technology to cope with the expanding surveillance data surface.

Overall though, it's clear that more than three quarters of banks feel that budgets are meeting the challenges of increased coverage and scope that have been discussed throughout the report and the largest single cohort see budgets overall rising in 2026 to meet the BAU challenge.

CHART 33
Do you expect your available budget for investing in market abuse surveillance in 2026 to:

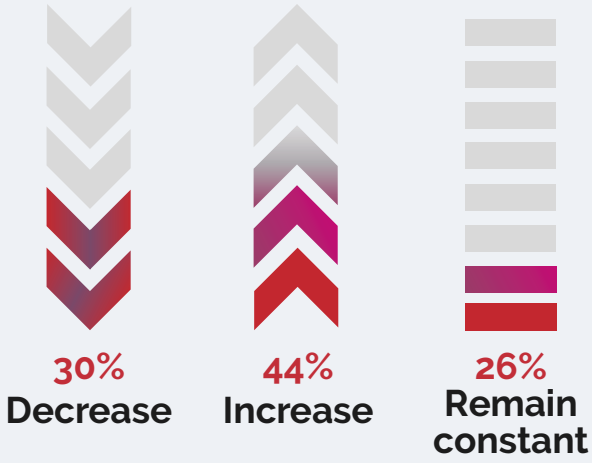

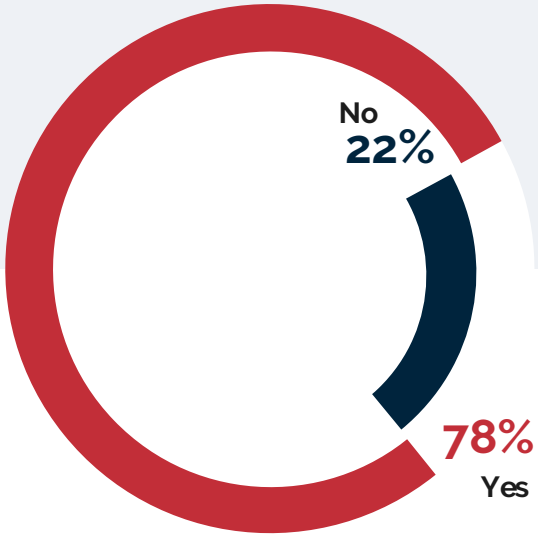

CHART 36
Do you have sufficient budget to build and maintain a surveillance function capable of effectively managing the risk of market abuse within your bank?

Copyright 2026, 1LoD. All rights reserved

CHART 23
How confident are you that your surveillance function is effective enough to detect potential future cases of market abuse in your bank?

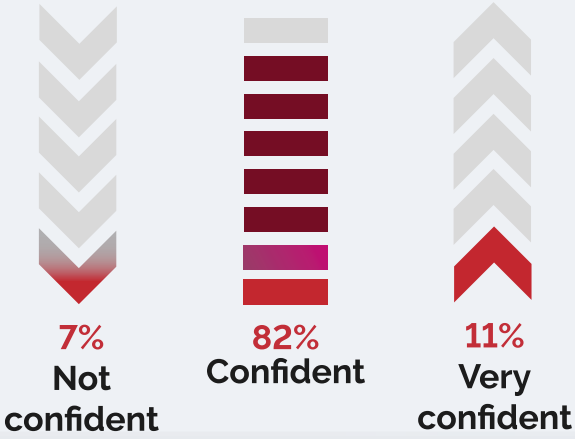

Transformation still a tough pitch

However, as we have also seen, the fact that BAU surveillance is largely seen as covered by todays budgets, the same is not true for CTB projects. But **Chart 36** may provide part of the explanation: if market abuse risk can currently be effectively managed, then why plough huge sums into transformation projects?

The answers – scalability, effectiveness, efficiency – are easy to list, but harder to pitch. Surveillance functions have always had a hard time getting investment today for rewards tomorrow, and little has changed in that regard.

Risks still not effectively managed

It is also striking that (just as in 2024), a sizeable 22% still admit that 'no', current levels of spend meant that market abuse risk was not effectively being managed. These replies did not only come from smaller institutions; they also came from large organisations, some of whom under real regulatory pressure feel that they are some way from meeting the requirements that have arisen from enforcements and investigations, and who need more budget to get there.

As the surveillance head at one of these organisations said, "We're definitely one of those banks that are investing more, as we know that there are certain areas where we need to catch up with our peers as well as satisfy regulators."

Those regulators would likely be alarmed that almost a quarter of banks surveyed believe their market abuse surveillance infrastructure does not provide effective risk management of market abuse risk.

A sizeable 22% still admit that 'no', current levels of spend meant that market abuse risk was not effectively being managed.

CHART 27
How do you anticipate the headcount in your surveillance function changing over the next 12 months?

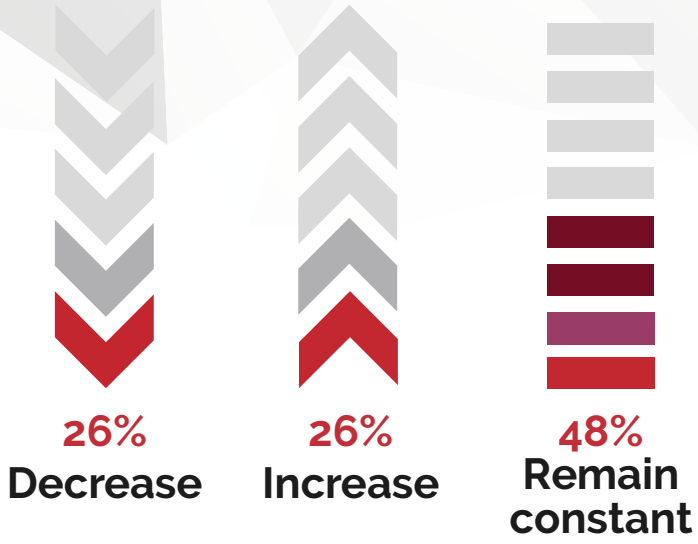

Technology investment drives headcounts

Despite all the hype around technology, the biggest costs in surveillance are still humans and here we see the same pattern as in budgets – with far fewer people than in 2024 believing that things will stay the same **[Chart 27]**

Asked 'how do you anticipate the headcount in your surveillance function changing over the next 12 months?' 48% believe that they will remain constant (versus 67% in 2024), and then there is a 26%/26% split between those who believe headcounts will rise and those who think they will fall.

As with budgets, these different responses represent very distinct groups of institutions: those who believe headcounts will fall are almost all global institutions whose heavy investments in surveillance technology transformation are now bearing fruit.

Those who believe headcounts will rise are largely in organisations still building out their surveillance capabilities to match their peer groups. These are generally larger mid-tier banks or those with a significant regulatory enforcement burden still running.

Those who believe headcounts will remain static are mostly at larger mid-tier banks who are happy with their current surveillance operations and do not believe that market abuse risk is so significant for them that they have to invest today in technology transformation. That said, these banks are still upgrading their capabilities on a case-by-case basis. A number of them report static headcounts but rising budgets – the clear implication being that the additional spend is going on technology.

Copyright 2026, 1LoD. All rights reserved

Holistic surveillance struggles

The future of surveillance depends on budgets, but getting budgets depends on demonstrations of ROI on individual initiatives. One place where that demonstration has clearly not worked yet is in integrated/holistic surveillance.

We saw earlier **[Chart 5, Section ii]** that while levels of investment in the individual surveillance buckets was anticipated to be healthy, banks foresaw no sustained investment in holistic surveillance, little in the way of targeted increases, and mostly flat investment levels for the next 24 months – with the clear implication that that level was not significant.

Those findings are borne out by the responses to the question **[Chart 20]** "how best describes the current state of your holistic surveillance capabilities?" Just over half of the banks replied that 'some' of their surveillance controls are linked; just under half said 'none' were. Given how long holistic surveillance has been a topic of conversation, this is not particularly rapid progress.

As might be expected, the banks most advanced in linking surveillance controls are the largest and best funded; they are also those who show the most interest in broader surveillance transformation. But the real divide is not between "advanced" and "lagging" firms, but between those who can afford to attempt integration and have the data architecture, organisational alignment, and surplus capacity required to make it work, those who are forced into it by scale pressure, and those who simply cannot yet prioritise it.

As for genuinely combining trade and comms data pre-alert, and creating alerting algorithms to generate alerts from those combined datasets, no banks represented in the survey have a process in production even though there are vendors now offering products of that kind.

CHART 20
How best describes the current state of your holistic surveillance capabilities?

Optimising silos takes priority

The main reasons given for this lack of progress fall into two categories: first, few banks believe that they can aggregate the comms and trade data in the way that would be required to create real-time alerting from algorithms that used the combined data signals.

And they point out that the fundamental difference in the two datasets (trade is structured, e-comms is unstructured) creates additional issues linked to the different speeds at which each type of surveillance is incorporating new technology, particularly AI.

As one trade surveillance head at a large US bank explains, "It's hard to marry the AI outputs you get in e-comms with the results of [the remaining] lexicons for a start, and then to marry those two with trade data, which uses different technology again, is challenging to say the least."

Also, more generally right now, the need to tie specific comms to specific trade data pre-alert would hugely slow down the alerting process in both areas.

Second, even fewer banks believe that the combination of data would actually decrease noise or detect more true positives. "If you try to map patterns and trends through the different risk systems you just end up in a sea of noise," explains one global head of surveillance.

All of this means that banks are focusing on optimising each silo first. Holistic surveillance requires high levels of maturity in each, and only then "can we move to merging the two [trade and comms]. There are definite use cases – in front running, for example, you very much need comms to look at the trades, but we are not there yet," according to one UK-based surveillance head.

AML + trade surveillance = nowhere

Slow progress towards holistic surveillance is further reflected in the lack of headway being made in aligning MAR surveillance functions with transaction monitoring functions in anti-money laundering (AML). Just 15% of respondents say that they do align, usually trade surveillance, with TM, with 85% saying they do not **[Chart 21].**

Banks generally accept that transaction monitoring and trade surveillance are inherently linked because market abuse such as insider trading and manipulation are predicate offences to money laundering, and therefore separating the detection of those activities risks missing the full lifecycle of

Copyright 2026, 1LoD. All rights reserved

financial crime. But they point out that trade data and payment data are inherently different and that the two domains operate on fundamentally different data sets and logics—TM focuses on money movement, while trade surveillance focuses on market behaviour—and linking a payment to a trade at a granular level remains technically difficult, reflecting legacy architectures that were never designed for this purpose.

The key though is that heads of surveillance do not see much value in the alignment whereas heads of financial crime do. Put another way, there is an asymmetry in risk: market abuse may imply money laundering, but the reverse is rarely true, limiting the case for full bidirectional convergence. As one surveillance head says, "Market abuse suspicion… could also hold an AML suspicion… the other way around is not really true."

So, if this alignment is to happen, it is more likely to be driven (and funded) out of financial crime departments than surveillance. And there is regulatory impetus for just such a development.

In its January 2025 report on 'Assessing and reducing the risk of Money Laundering Through the Markets (MLTM)', the FCA is critical of TM systems and encourages "collaboration … between trade surveillance and TM teams to identify and review potentially suspicious activity' because 'most firms have found that using their current automated TM systems in isolation provides limited success in identifying suspicions of MLTM."

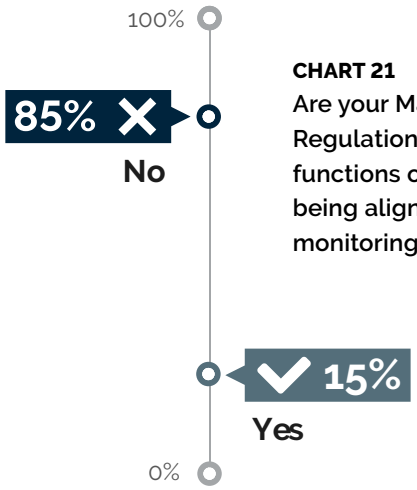

CHART 21
Are your Market Abuse Regulation (MAR) surveillance functions organisationally being aligned with transaction monitoring functions in AML?

Budgets, staffing sufficient

A key element of future surveillance capabilities is cross-product market abuse detection and prevention. This remains a significant challenge **[Chart 18]**. The majority of banks (56%) are still developing a capability that they would not even define as 'basic'. 26% say that they would describe the current state of their ability to surveil for cross-product market abuse as 'basic' and just 18% claim that it is 'intermediate'. No-one believes that their capability is 'advanced'.

The banks who describe their capability as 'intermediate' are all large, well-resourced global institutions – another example of this divide in the survey data. Even they find the challenge tough as this surveillance head emphasises: "Cross-product manipulation is still the most challenging area both because of the number of relationships there are between different types of securities and because it is extremely difficult to identify which are the leading and lagging markets and then from that see which relationships can be exploited."

That said, all respondents position cross-product as a key strategic priority, driven both by regulatory focus and the limitations of traditional, single-asset surveillance. And they are all actively expanding capability in this area.

However, for most, the practical implementation is highly resource-intensive and largely in-house. Some depend on data sources that require custom development and for cross-product models to be designed and built from scratch.

As one EMEA head of surveillance at a large Asian bank explains, "we can't go to [vendor x] and say… switch on this cross-product alert. And in general, current vendor solutions are not yet mature enough to handle complex, cross-asset behaviours out of the box."

Cross-product hopes pinned on AI

This is because the core constraints for most banks is both data and modelling. "Data is an issue but even with perfect data, differentiating arbitrage across products from illegitimately creating a price move that you then exploit in a different asset is hard and even harder to operationalise," says one US-based head of trade surveillance.

Banks are hoping that the real breakthrough will not come from more rules, but from AI-driven linkage of products and markets and analysis of behaviours. Even then, they are sanguine about the chances of success. As one says, "AI might help identify patterns and surface suspicious scenarios, but still just telling it what to look for would be the complication."

Trader profiling? No way.

Perhaps the clearest indication that there are currently limits to the level of innovation deemed acceptable by surveillance chiefs is the response to the question 'do you use trader behaviour analysis to filter your alerts' **[Chart 19]**. Here a very surprising 93% of banks say that they do not. Even more surprisingly, this is up from the 86% who said 'no' in 2024.

The objections fall into two broad groups. Most surveillance heads feel that categorising traders on the basis of their trading and other (non-trading) behaviours is in some way is unethical. They also fall back on a privacy/General data protection regulation (GDPR) argument that it is an illegal intrusion of the employer into the employee.

Neither of these seems to bear much scrutiny: far more intrusive employee surveillance is carried out both by banks on their entire populations and by companies in other sectors, and the trend is for that intrusion to grow not shrink. GDPR makes specific exceptions for activities that could be illegal and no-one is suggesting that trader profiling is used for HR purposes anyway.

The better arguments against profiling come from those who don't believe it works. As one head of trade surveillance at a large US investment bank says, "I'm pro the theory but dead against the practice right now. I don't think we [the market] have the level of sophistication to build models complex enough to do anything except generate more noise.

Copyright 2026, 1LoD. All rights reserved

CHART 19
Do you use trader behaviour analysis to filter your alerts?

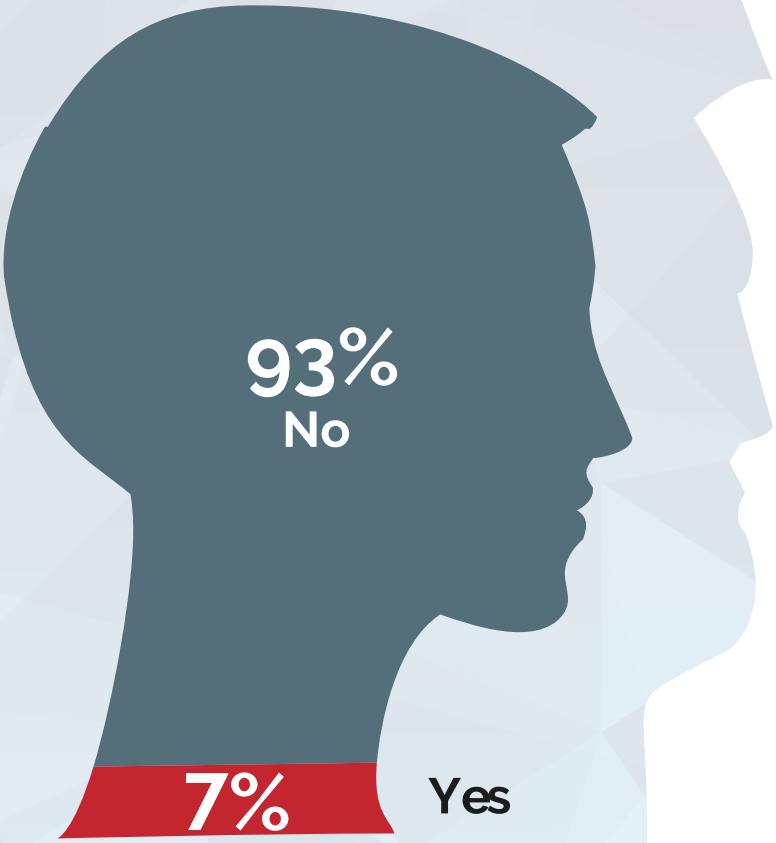

CHART 18
How best describes the current state of your ability to surveil for cross-product market abuse?

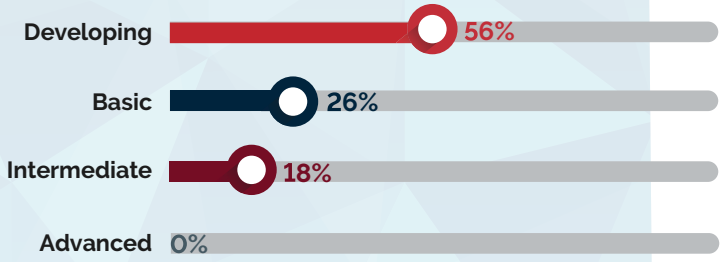

Yes, we are being pushed by regulators and auditors to do this. And it hasn't really yielded any results. I'm yet to see something to support that investment."

Another sceptic says, "I've seen idiosyncratic cases… bespoke cases, nothing scalable… nothing that I could say was a genuinely tangible success."

In this view, the types of behavioural data available are simply not diagnostic enough to make a difference.

But there are banks who have implemented sophisticated trader profiling. They use disparate trade, comms and non-surveillance data to rank traders by perceived riskiness and they claim that those rankings enable them to disposition alerts in ways that have yielded an increase in true positives as well as highlight serious HR-related issues that have been resolved well.

"We have built an in-house system that incorporates data from trading behaviours, communications, previous incidents on the desk, HR and other data and we use it to assign risk scores to individuals," explains one surveillance head. "It can be used to prioritise alerts, which is essentially about efficiency. But it also surfaces new risks and new instances of abuse that we did not detect before. It also acts as an early warning system and so is actually a preventive control too."

It is hard to reconcile such conflicting views. It may be that banks are distinguishing between what they see as mechanistic risk-scoring and a more AI-driven behavioural modelling. But the almost blanket refusal to countenance profiling is unusual.

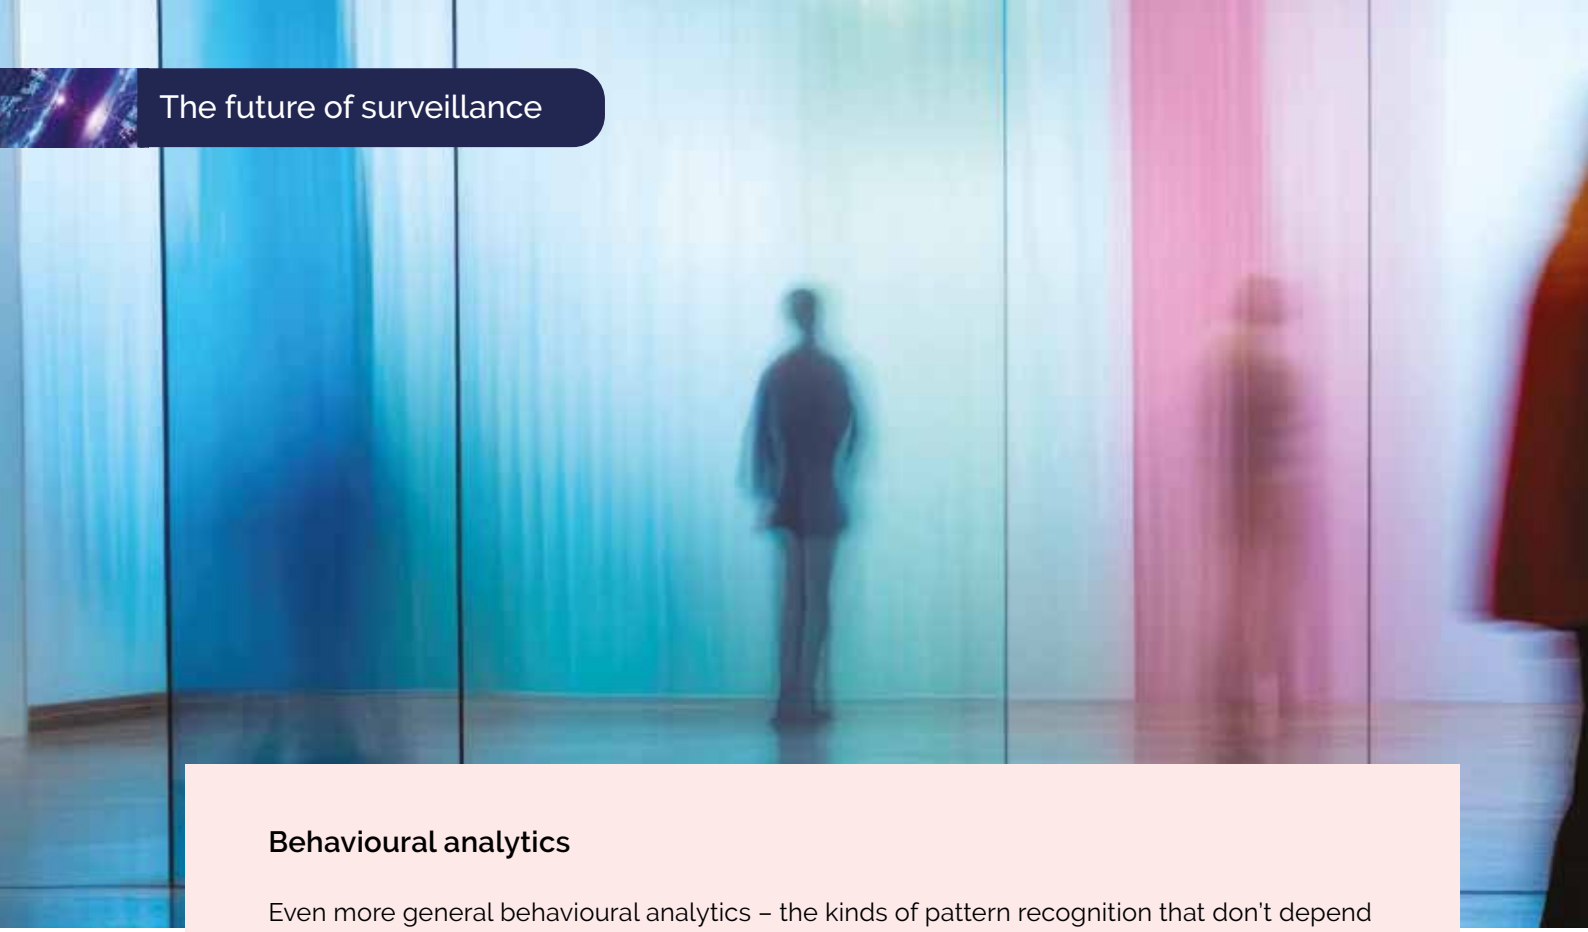

Behavioural analytics

Even more general behavioural analytics – the kinds of pattern recognition that don't depend necessarily on trader profiles – are still uncommon (11% have in surveillance platforms), and somewhat unwanted –37% do not require this functionality **[Chart 7B]**.

CHART 7B

Which features are available in your surveillance platforms?

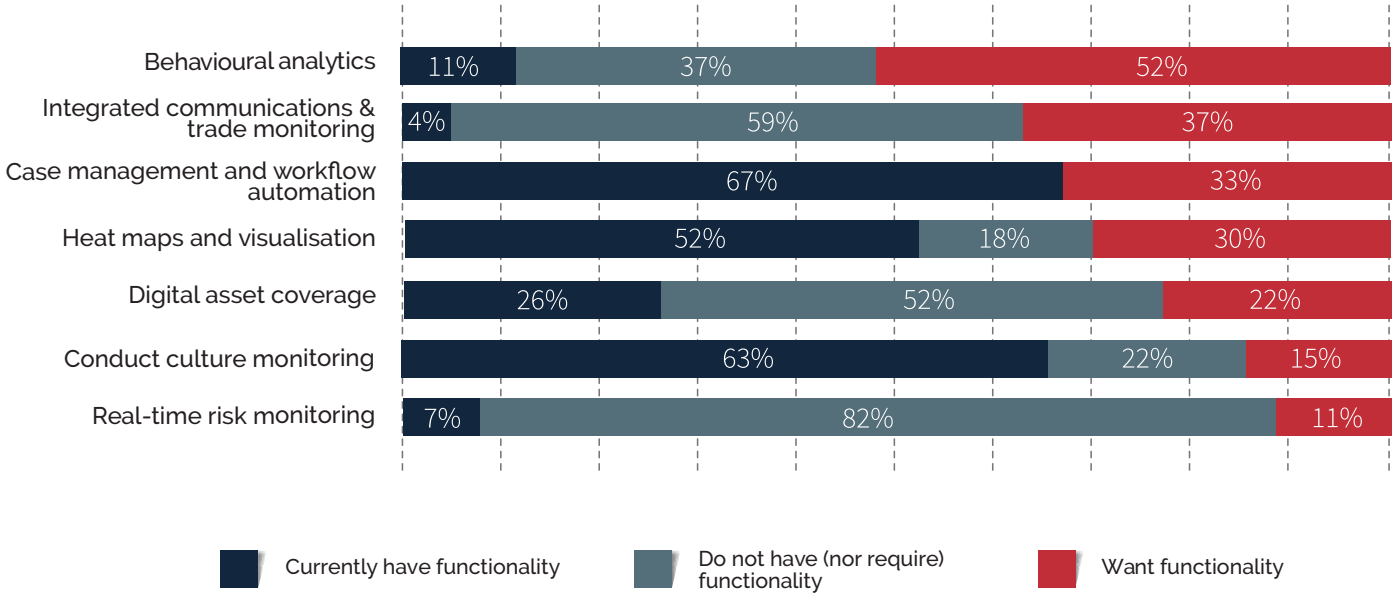

Copyright 2026, 1LoD. All rights reserved

The real tipping point

The survey reveals several key tipping points for the surveillance function, from technology to skillsets to operating models. But all of these are essentially symptoms of a more fundamental required transformation. Is surveillance going to remain a somewhat static, detective compliance function, or is it going to develop into a true risk management function?

The survey suggests that surveillance teams want that transformation but still face very significant hurdles to achieving it – with key issues beyond their control.

A final data point that shows how far surveillance needs to travel to match other risk functions comes in **Chart 7B**: just 7% of respondents report having real-time capabilities, and an overwhelming 82% say they neither have nor require them. This is perhaps the clearest indication that the dominant model of surveillance is still detective rather than preventive. Alerts are generated after the fact, investigations follow, and controls remain largely reactive. This stands in contrast to regulatory pressure—particularly from prudential authorities—to move toward more preventive and automated control environments. The industry, at least for now, has not made that transition.

Taken together, the chart reveals a surveillance function that is operationally well-developed but analytically constrained. Banks have built effective machinery for managing alerts and conducting investigations, but they have not yet transformed the underlying logic of surveillance itself. The aspiration is clearly to move toward more intelligent, integrated, and real-time systems, but the reality remains one of fragmented data, legacy architectures, and cautious governance frameworks.

The result is a widening gap between ambition and delivery. Surveillance functions are still largely designed to detect isolated events within specific silos, using static or semi-static logic, and doing so after the fact. Yet the risks they are meant to address are increasingly behavioural, cross-channel, and dynamic. Until that mismatch is resolved, the industry will continue to operate in a state of partial effectiveness—highly active, increasingly sophisticated in process, but still lacking the holistic understanding required to fully manage modern market abuse and conduct risk.

A 1LoD Survey

The 2026 Surveillance Benchmarking SURVEY & REPORT

Please contact

Andrew Lennon
Head

andrew.lennon@1lod.com
+44 (0) 7914 215 980

www.1lod.com