Sygno

Transaction Monitoring under AMLA's Draft Guidelines

Part 2: What the framework must demonstrate, and how it stays effective

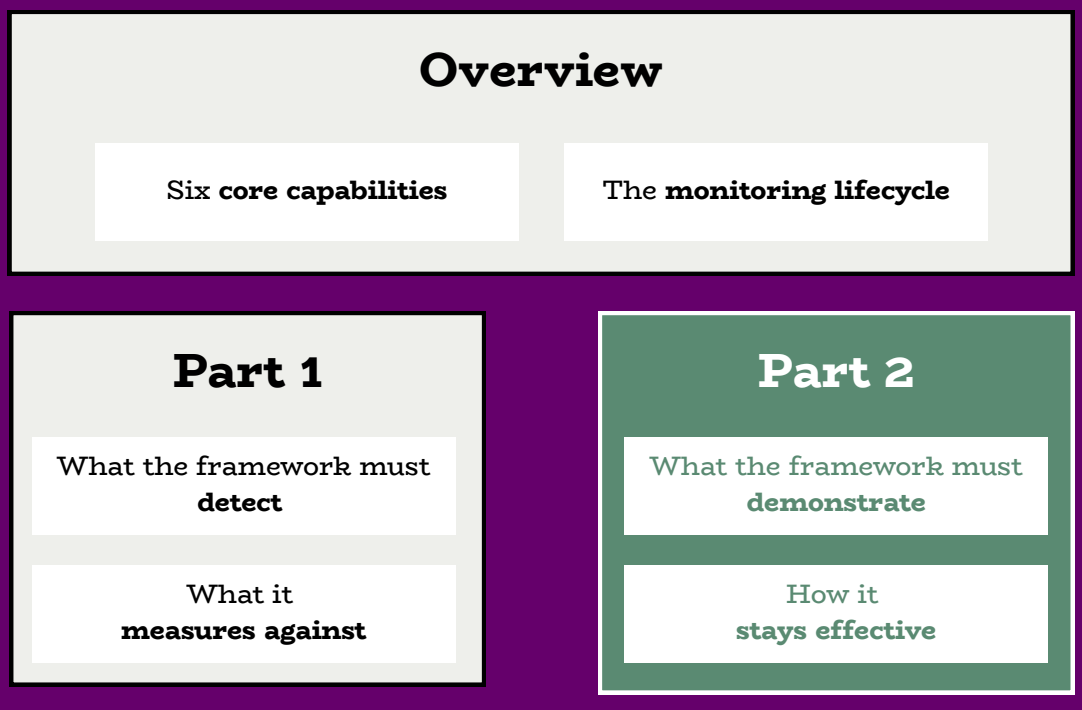
Whitepaper | August 2026

Status of this paper

This paper reflects the regulatory position as at August 2026. It is based on AMLA's consultation paper on draft Guidelines on ongoing monitoring of a business relationship under Article 26(5) of Regulation (EU) 2024/1624, published on 3 June 2026, with the consultation period closing on 3 September 2026 and final guidelines expected in Q4 2026. The guidelines are in draft form and may change following consultation. Where this paper refers to draft regulatory technical standards under Article 28(1) AMLR, those standards have completed consultation but have not yet been adopted by the Commission. References to AMLR articles are to the regulation as adopted and in force.

About this set of papers

This is one of a set of three papers covering AMLA's draft Guidelines on ongoing monitoring. The overview is an extended management summary of what the guidelines ask and what follows from it. Two detailed papers carry the analysis, provision by provision. Part 1 covers what the framework must detect and what it measures against. This part covers what the framework must demonstrate, and how it stays effective over time.



Introduction

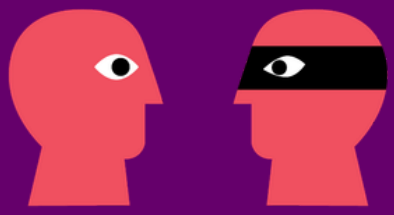
Article 26 AMLR requires obliged entities to conduct **ongoing monitoring of every business relationship**, and Article 26(5) mandated AMLA to issue guidelines on how that should be done. The draft guidelines apply horizontally, across the financial and non-financial sectors, and are written around outcomes rather than prescribed methods.

Part 1 covered what shapes a monitoring framework before it goes into use: what enters it, what deviation is measured against, what has to be found across time and connections, and where the detection logic comes from. Para 34(f) marks the point where this part takes over. A framework has to be effective at the point of implementation, and it has to remain effective and fit for purpose over time.

Staying effective asks different things of an institution than getting there did. The guidelines state that three measures institutions commonly report, namely alert volumes, reporting volumes and the breadth of typology coverage, do not on their own settle effectiveness where they do not lead to meaningful detection or escalation. What they ask for instead is qualitative and harder to produce on request. Most of the provisions in this part are of that kind: they ask an institution to establish something about its own framework, and several of them cannot be answered from what the framework produced.

Scope of this paper. The requirements sit mainly in sections 2.10 and 2.11 of Guideline 2, on internal controls and ongoing review and on the use of technology, drawing on provisions elsewhere in the guidelines where these bear on the framework in operation. Where a provision concerns how an institution organises itself rather than what its framework has to be capable of, it falls outside this scope.

Convergence with the PSR. Para 34(a) points obliged entities to systems and processes already required by other Union legislation, naming the proposal for a regulation on payment services in the internal market alongside Regulation (EU) 2023/1113. For payment service providers, the fraud monitoring the PSR requires and the transaction monitoring Guideline 2 requires rest on much of the same capability, and the reversal of the burden of proof under the PSR asks for the same record this part describes. The overview paper sets that convergence out in full.



The four themes

Part 1 covered what shapes a monitoring framework before it goes into use. The requirements in this part sit mainly in sections 2.10 and 2.11 of Guideline 2, on internal controls and ongoing review and on the use of technology, drawing on provisions elsewhere in the guidelines where these bear on the framework in operation.

5

What effectiveness means

The standard the framework is held to. Obliged entities have to define the outcomes their framework is meant to produce and reassess whether those outcomes remain appropriate. Effectiveness is not settled by alert volumes, reporting volumes or breadth of typology coverage alone, where these do not lead to meaningful detection or escalation. It also covers what the framework did not surface at all, and what it surfaced but left unresolved.

6

Staying fit for purpose

What has to keep happening once the framework is live. The framework is reviewed and tested periodically, internally and where relevant externally, and stays adaptable as a customer's risk profile changes. New methods, trends and typologies are taken in without undue delay and earlier decisions revisited, with every change to the detection logic documented, tested and validated on evidence that goes beyond a comparison with the previous version. Its own performance is monitored, including degradation and drift.

7

Explainability, and to whom

Two audiences. Outward, the framework's behaviour, its changes and its decision logic have to be reconstructable on request, and the role, functioning and output of analytical tools explained to a competent authority without requiring full technical interpretability of the model. Inward, the staff who operate the framework have to be able to interpret its output and question the results.

8

Data, tools and responsibility

The inputs and the instruments. The quality of the data has to be controlled, its deficiencies found, and its attribution to customers, counterparties, assets and activities made visible. On the tools, safeguards scale with the role of the tool and model risk is weighed on both sides: the failure modes of the model, and the risk of delaying or limiting improvements in detection. Responsibility for monitoring decisions stays with the obliged entity, which limits what can be taken from a third party.

Each theme sets out the requirements it places on an obliged entity, the provisions those requirements derive from, and the system properties they imply.

5 What effectiveness means

A framework that has been running for a year has produced a population of alerts, escalations and reports. Guideline 2 sets out what has to have been recorded before those were produced, what may be offered as evidence that they are the right ones, and which cases have to be counted that appear in none of them. Three requirements follow from the relevant provisions.

A. Define what the framework is meant to achieve, and reassess it

B. Measure effectiveness on the quality of outcomes

C. Sample beyond what the framework produced

Defining the outcome

Para 83 requires obliged entities to **define the outcomes of their monitoring framework** and periodically reassess whether those outcomes remain appropriate in light of their business-wide risk assessment and, where relevant, publicly available national or Union-level priorities.

The obligation runs before any measurement. An institution that has not stated what its framework is meant to produce has nothing to assess its performance against, and no basis on which to judge whether a change to the detection logic was an improvement.

Para 5 ties those outcomes to the institution's own risk picture. The framework has to be **aligned with the risks identified in the business-wide risk assessment**, which places the same document at both ends of the exercise: it determines what the framework is built to find, as Part 1 sets out, and it is the reference against which the framework's outcomes are judged.

Para 34(f) carries that forward in time. A framework has to be effective at the point of implementation and **remain effective and fit for purpose over time**. The second half of that requirement is what makes the reassessment in para 83 recurring rather than a one-off exercise at go-live.

What effectiveness is measured on

Para 84 states what effectiveness rests on and what it does not rest on. Obliged entities should consider the extent to which their framework produces outcomes that are **relevant, risk-based and capable** of supporting the identification and reporting of ML/TF risks. It then limits the weight of three measures: effectiveness should not be assessed solely by reference to alert volumes, reporting volumes or the breadth of typology coverage where this does not result in meaningful detection or escalation outcomes.

What it points to instead is qualitative: the timeliness of escalation, the appropriateness of outcomes, whether deficiencies recur, and whether the framework remains capable of addressing the risks.

All three of those measures are countable, and all three are commonly reported. Two of them, alert and reporting volumes, describe how much the framework produced. Typology coverage describes how many known patterns it has logic for. Neither group answers whether the output was relevant or whether escalation happened in time, and para 84 makes the qualitative measures the ones that decide.

Para 87 applies the same reasoning to technology. The deployment of **automated or advanced analytical tools should not, in itself, be considered an indicator of effectiveness**, which is instead assessed by reference to whether the tools enable the institution to detect and escalate relevant risks without undue delay. Having a model is no more evidence of effectiveness than having a high alert count.

Beyond what the framework produced

Para 84 points obliged entities to the timeliness of escalation when assessing effectiveness. Output that is raised and never worked leaves the risk it identified unaddressed.

Paras 78 and 79 turn to output the framework raised and then closed on its own. Para 78 **restricts automated closure** to cases which, following automated analysis, do not indicate suspicious or unusual transactions and activities or other material risk indicators, and rules it out where there are indicators of heightened risk, including higher-risk customers and situations requiring enhanced scrutiny. Para 79 requires the effectiveness of those mechanisms to be **subject to effective human oversight, including periodic validation through sampling**, the review of potential cases where suspicious transactions or activities were not identified, or other appropriate checks.

An institution using a model to suppress alerts from an existing rule set is operating an automated closure mechanism, whatever it is called internally. Every suppression has to hold to the limit in para 78, and para 79 requires the underlying decision logic to be understood and documented, which means the output has to **show which risks contributed to the score** it was closed on.

Neither provision reaches the transactions that never produced output at all, but two others do. Para 84 points to recurring deficiencies as a signal that effectiveness is lacking, and a deficiency in the detection logic recurs in the transactions the framework passes over, which are absent from every case file and every review of closed cases. Para 49, covered in Part 1, requires design choices to be weighed against both detection capability and the **risk of undetected activity**. Establishing that risk means testing against activity the framework did not flag, which no measure taken from its output can do.

Para 84 puts it a third way. Effectiveness rests on outcomes that are capable of supporting the identification and reporting of ML/TF risks, and a framework that has no logic for an entire category of risk is not capable of supporting the identification of it. Establishing whether that is the case calls for **a method that can surface risk the framework has no logic for**, since its own output holds nothing from those categories.

Testing against activity the framework did not flag is known as **below-the-line testing**. The guidelines do not use the term and neither provision prescribes a method, but without something along those lines an institution cannot establish its own effectiveness.

KEY PROVISIONS

- **Para 5**: Framework effective and aligned with the risks identified in the business-wide risk assessment
- **Para 34(f)**: Framework effective at the point of implementation and remaining effective and fit for purpose over time
- **Para 49**: Design choices weighed against both detection capability and the risk of undetected activity
- **Para 78**: Automated closure limited to cases showing no indication of suspicious or unusual activity or other material risk indicators, and excluded where indicators of heightened risk are present, including higher-risk customers and situations requiring enhanced scrutiny
- **Para 79**: Effectiveness of automated mechanisms subject to human oversight, including sampling, the review of potential cases where suspicious transactions or activities were not identified, and other appropriate checks
- **Para 83**: Outcomes of the framework defined and periodically reassessed against the business-wide risk assessment and, where relevant, national or Union-level priorities
- **Para 84**: Effectiveness rests on outcomes capable of supporting the identification and reporting of risks, considering the timeliness of escalation, the appropriateness of outcomes and whether deficiencies recur, rather than on alert volumes, reporting volumes or breadth of typology coverage alone
- **Para 87**: Deployment of automated or advanced analytical tools is not in itself an indicator of effectiveness

System implications	
Requirement	**Monitoring capability**
A. Define what the framework is meant to achieve, and reassess it	Detection logic readable and traceable to the risks in the business-wide risk assessment Detection logic that can be updated or regenerated without rebuilding the framework
B. Measure effectiveness on the quality of outcomes	Measurable performance on known historic cases, with high detection and low false positives Detection that reaches criminal activity the earlier logic kept missing, so recurring deficiencies close rather than persist
C. Sample beyond what the framework produced	An explainable basis for every score, showing which risks contributed to it Testing against activity that produced no alert, sufficient to estimate what the framework misses (below-the-line testing) Sampling results feeding directly into an updated model Anomaly detection built on what legitimate behaviour looks like, so anomalies surface without a pre-defined risk typology

6 Staying fit for purpose

A framework that was effective at go-live does not stay that way on its own. Guideline 2 asks two questions of it over time, and sets conditions on the record an institution answers them from. Three requirements follow from the relevant provisions.

A. Know whether the framework still performs as validated

B. Know whether it still covers the risk, and adapt where it does not

C. Keep the framework auditable and its problems visible early

Whether the framework still performs

Para 34(f) requires the framework to remain effective and fit for purpose over time through **ongoing testing and adaptation to emerging risks**. Testing and adaptation are named as the mechanisms, which makes both of them obligations.

Para 95 covers the continuous side. Analytical tools have to maintain the expected performance over time, including through appropriate monitoring and the identification of material **performance degradation or drift**, in a manner commensurate with the role of the tool and the associated risk. Detecting drift means comparing how the tool behaves now against how it behaved when it was validated, which requires earlier behaviour to have been recorded in a form that supports the comparison.

Para 81 covers the periodic side, and draws on the first. Obliged entities review and test the framework periodically to establish that it remains effective and aligned with the risks in the business-wide risk assessment. The provision lists what feeds that review: **monitoring performance, the outcomes of escalations and suspicious activity reports, emerging risks and other relevant feedback**, including, where relevant, the performance and limitations of analytical tools used within the monitoring framework.

Whether it still covers the risk

Performance and coverage are different questions. A model can behave exactly as it did at validation while the risk it was built for has moved.

Para 43 addresses the outside view. The framework has to be capable of identifying and assessing, without undue delay, new ML/TF methods, trends and typologies. Obliged entities assess and document whether the framework needs amendment, and should **revisit earlier decisions regarding their framework** where needed, including in response to internal cases, supervisory feedback or other credible external information.

Two things follow. The trigger sits outside the institution, so it has no control over when the obligation applies, and "**without undue delay**" attaches to a change in the framework rather than to a case. Revisiting earlier decisions means the design is never settled: a scenario justified two years ago has to be defensible again when the picture changes.

Para 5 points the same way from the customer side. The framework has to be **adaptable to changes in the customer's risk profile** and aligned with the risks in the business-wide risk assessment, including emerging risks related to products, services, customers, geographic exposure and distribution channels.

Para 85 sets the standard for validating a change. Where obliged entities materially change their monitoring logic, analytical methods or approach, they apply a documented transition and validation framework to **assess whether the revised approach remains effective and consistent with the defined outcomes and the business-wide risk assessment**. Validation should include testing that goes beyond comparison with legacy outputs. Comparative analysis with prior approaches may be used where it provides meaningful insight into potential gaps or risks.

Keeping the framework auditable

Para 45 requires **periodic internal and/or external reviews** to maintain the effectiveness of the framework. This requires that the detection logic can be read, the reasoning behind it is written down, and the reviewer can establish what the framework was doing at a given point without asking the people who built it.

Para 49 sets the same condition on every change. Updates to detection logic, underlying configuration and, where relevant, models or behavioural baselines have to be **documented, tested and recorded**. This works both ways: recording a change without testing it leaves an institution unable to say what effect it had, and testing without recording leaves it unable to show that it did.

Para 76 requires obliged entities to define and maintain internal arrangements including mechanisms to identify, monitor and address any accumulation of pending monitoring outputs that may affect the effectiveness of the monitoring framework. The provision addresses the queue, which is where the problem becomes visible. It forms earlier, in the rate at which the framework produces output that leads nowhere. **Monitoring the false positive rate** allows an institution to act before the queue accumulates to the point where it affects effectiveness.

KEY PROVISIONS

- **Para 5**: Framework adaptable to changes in the customer's risk profile and aligned with the risks in the business-wide risk assessment, including emerging risks
- **Para 34(f)**: Framework remains effective and fit for purpose over time through ongoing testing and adaptation to emerging risks
- **Para 43**: New methods, trends and typologies identified and assessed without undue delay; earlier decisions on the framework revisited where needed
- **Para 45**: Periodic internal and, where applicable, external reviews to maintain the effectiveness of the framework
- **Para 49**: Updates to detection logic, configuration, models or behavioural baselines documented, tested and recorded
- **Para 76**: Mechanisms to identify, monitor and address accumulation of pending monitoring outputs that may affect effectiveness
- **Para 81**: Framework reviewed and tested periodically against the business-wide risk assessment, drawing on monitoring performance, escalation outcomes, reports, emerging risks and the performance and limitations of analytical tools
- **Para 85**: Material changes subject to a documented transition and validation framework, with validation that goes beyond comparison with legacy outputs
- **Para 95**: Analytical tools maintain expected performance over time, with identification of material performance degradation or drift

System implications

Requirement	Monitoring capability
A. Know whether the framework still performs as validated	Drift monitoring in place, comparing current model behaviour against its behaviour at validation Investigation outcomes and risk assessment changes fed back into the next model update
B. Know whether it still covers the risk, and adapt where it does not	Detection logic regenerated on current data without rebuilding the framework Validation of a new model measurable against previous versions and the risk assessment
C. Keep the framework auditable and its problems visible early	Detection logic and its rationale readable by an external reviewer Logging of every change to logic, configuration or baseline, as well as test results and effective date False positive rates tracked over time, alongside true positives

7 Explainability, and to whom

Explanation runs in two directions. A supervisor asks what the framework did and why. The people who run it work with its output every day and have to be able to push back where it looks wrong. Three requirements follow from the relevant provisions.

A. Make the framework's outcomes and decision logic reconstructable

B. Be able to show it, on request and within the time allowed

C. Enable operating staff to interpret the output and challenge it

Reconstructing what the framework did

Para 45 requires the monitoring framework to operate in a controlled and sufficiently transparent manner, underpinned by clear governance arrangements and proportionate documentation. It then sets out what that transparency has to deliver: obliged entities should be able to **explain material changes in monitoring behaviour or outcomes**, including those resulting from updates to rules, thresholds, models or analytical methods. The state of the **logic and its configuration need to be recoverable** at each point in time to be able to explain changes in monitoring behaviour and outcomes.

Para 79 makes the same demand of automated mechanisms. Obliged entities should ensure that **the underlying decision logic can be understood and documented**.

Para 8 covers the decision trail around the framework. Governance and the related decision-making processes are documented in a manner that supports traceability, accountability and effective supervision, and obliged entities should be able to **understand what was assessed, what decisions are taken and the rationale behind them**. The provision applies to the framework as a whole, which places the reasoning behind a monitoring decision on the same footing as the decision itself.

Producing it on demand

Para 42 closes with an obligation on the framework as a whole. Obliged entities should be able to **demonstrate the functionality, rationale and effectiveness** of their monitoring framework to the competent authorities upon request. It is the only provision that puts the framework as a whole under an on-request obligation; para 90 does the same for tools, and para 45 for material changes.

Para 90 sets the standard for those tools. Obliged entities should be able to demonstrate and explain, upon request, the role, functioning and outputs of such tools to competent authorities in a manner that enables **effective understanding, assessment and challenge of those outputs, without requiring full technical interpretability of the underlying model**.

The final clause draws the line at understanding rather than at interpretability. A supervisor does not have to be able to follow the mathematics of a model, but does have to be able to challenge what the model produced. That **rules out a black box** whose output cannot be questioned, and it also **rules out an explanation that only a data scientist can act on**.

Article 69(1) AMLR sets a deadline for a related capability. Obliged entities have to reply to requests for information from the FIU **within five working days**, and in justified and urgent cases, FIUs may shorten that deadline, including to **less than 24 hours**. What that establishes is that retrieving and assembling what the framework holds is a matter of days at most, and sometimes hours.

Enabling the people who operate it

Para 7 addresses the people who operate the framework. Staff involved in ongoing monitoring have to have the necessary knowledge, skills and professional judgement to operate the monitoring framework effectively, understand and interpret its outputs, raise questions or challenge results where appropriate. The requirement on the output is therefore that **the score, the logic behind it and the data it drew on are available and presented** in terms staff can follow.

Para 86 places that on the institution rather than on the individual. Governance arrangements, including policies and procedures, should enable relevant staff to effectively understand and challenge monitoring tools and their outputs.

Para 94 adds what the output has to be good for. Outputs generated by analytical tools, including artificial intelligence, **should support monitoring decisions** and be subject to appropriate governance, oversight, challenge and control.

KEY PROVISIONS

- **Art. 69(1), third subparagraph, AMLR**: Requests for information from the FIU answered within five working days, shortened in justified and urgent cases to less than 24 hours
- **Para 7**: Staff able to operate the framework, understand and interpret its outputs, and raise questions or challenge results
- **Para 8**: Governance and decision-making documented to support traceability, accountability and effective supervision, covering what was assessed, what was decided and on what basis
- **Para 42**: Functionality, rationale and effectiveness of the monitoring framework demonstrable to competent authorities on request
- **Para 45**: Framework operating in a controlled and sufficiently transparent manner, with material changes in monitoring behaviour or outcomes explicable on request
- **Para 79**: Underlying decision logic of automated mechanisms understood and documented
- **Para 86**: Governance enabling relevant staff to understand and challenge monitoring tools and their outputs
- **Para 90**: Role, functioning and outputs of analytical tools demonstrable and explicable to competent authorities, without requiring full technical interpretability of the model
- **Para 94**: Outputs of analytical tools support monitoring decisions and are subject to governance, oversight, challenge and control

System implications

Requirement	Monitoring capability
A. Make the framework's outcomes and decision logic reconstructable	Logic and input configuration logged in a form that cannot be disputed and can be explained Every alert stored with the features and values that produced it
B. Be able to show it, on request and within the time allowed	Model documentation covering method, validation and output, available without a preparation exercise Reporting easily generated from the data available in the tooling
C. Enable operating staff to interpret the output and challenge it	Each alert accompanied by a plain-language statement of what drove it Model documentation written in the same plain language, for the people who operate the framework

8 Data, tools and responsibility

What the framework runs on, and what it runs with. Guideline 2 sets conditions on both, and leaves the responsibility for each with the institution. Two requirements follow from the relevant provisions.

A. Control data quality, find its deficiencies and make its limits visible

B. Know what the tools do, and weigh model risk in both directions

The data the framework runs on

Para 92 requires the data used within the monitoring framework to be of sufficient quality and integrity to support the effective identification and escalation of ML/TF risks, with para 82 adding the controls around it: robust data quality controls, data validation processes and regular data cleansing. Para 92 adds a second obligation. Obliged entities should understand and, where relevant, **document material limitations in the data or its use**, including how data is attributed to customers, counterparties, assets or activities.

Attribution is named separately: a field can be complete, accurate and current and still be attached to the wrong entity, and the framework will treat that as fact. Where attribution is unreliable, the aggregation and entity resolution described in Part 1 produce a picture that is internally consistent and wrong.

Para 93 covers what happens where deficiencies are material. Obliged entities should document the deficiencies identified, the basis for their assessment and the mitigating measures applied. **Documenting a deficiency presumes finding** it, and stating the basis for the assessment presumes a method of establishing how far it reaches. Part 1 covered completeness and timeliness of the data entering the framework; these provisions cover what happens when it is neither.

The tools it runs with

Para 87 requires obliged entities to assess whether automated or advanced analytical tools, including algorithms, machine learning, artificial intelligence and other innovative technologies, would **enhance the effective identification and escalation of ML/TF risks**.

Para 88 sets the **safeguards around such a tool** proportionate to size, nature, complexity and risk, and to the impact of the tool on monitoring outcomes or decisions, with outputs that can be reviewed and challenged. Para 89 names what those safeguards have to weigh, and it names two things.

The first is **model risk**, including risks arising from inaccuracies, biases, lack of robustness or limited transparency of analytical tools that may affect monitoring outcomes.

The second is the **risk of delaying or limiting improvements** in detection capabilities, along with unnecessary constraints that do not materially enhance those outcomes.

Para 95 puts the answer with the institution. Obliged **entities remain responsible for monitoring** decisions and should ensure these can be reviewed and, where necessary, adjusted. Para 86 requires the responsibility for overseeing effectiveness to be clearly assigned, with identified deficiencies addressed without undue delay.

Para 91 draws the consequence for what can be taken from a supplier. Where obliged entities rely on third-party providers, they should **ensure sufficient information is available** to meet the requirements in paras 88 to 90. Where this is not the case, they should not rely on such tools for monitoring functions that materially influence monitoring outcomes or decisions. Para 50, covered in Part 1, puts heavy restrictions on using default settings. Para 91 goes further: where the provider does not supply enough to explain and oversee the tool, the tool cannot be used for anything that materially affects the outcome, and documenting the gap does not cure it.

The two halves of para 89 rule out both a model an institution cannot answer for and a governance regime that holds back detection it could have had.

KEY PROVISIONS

- **Para 82**: Robust data quality controls, data validation processes and regular data cleansing
- **Para 86**: Responsibility for overseeing the effectiveness of the framework clearly assigned, with identified deficiencies addressed without undue delay
- **Para 87**: Assessment of whether automated or advanced analytical tools would enhance the identification and escalation of risks, taking into account whether manual processes are sufficient
- **Para 88**: Safeguards proportionate to the size, nature, complexity and risk of the entity and to the impact of the tool, with outputs that can be reviewed and challenged

KEY PROVISIONS (continued)

- **Para 89**: Model risk weighed, including inaccuracy, bias, lack of robustness and limited transparency, alongside the risk of delaying or limiting improvements in detection and of unnecessary constraints
- **Para 91**: Where a third-party provider does not supply sufficient information to meet those requirements, the tool is not to be relied on for functions that materially influence monitoring outcomes or decisions
- **Para 92**: Data of sufficient quality and integrity, with material limitations understood and documented, including how data is attributed to customers, counterparties, assets or activities
- **Para 93**: Material data deficiencies documented, together with the basis for the assessment and the mitigating measures applied
- **Para 95**: Responsibility for monitoring decisions remains with the obliged entity, with decisions reviewable and adjustable

System implications

Requirement	Monitoring capability
A. Control data quality, find its deficiencies and make its limits visible	Data quality controls and validation running as part of the monitoring process Data quality reporting easily generated
B. Know what the tools do, and weigh model risk in both directions	Performance reporting readily available for assessment Model output presented in understandable terms Model information transparent and available Improvements and model regeneration easy to carry out

From theory to practice

This paper has looked at **what the framework must demonstrate** and **how it stays effective**. Guideline 2 approaches that through four themes: what effectiveness means, what has to keep happening once the framework is live, what has to be explainable and to whom, and the data and tools the framework depends on. The requirements those themes place on an institution converge on three core capabilities.

4

A complete record of what the framework did and why

Every alert is stored with the features and values that produced it. Every change to the logic is logged with its rationale, its test results and the date it took effect. The detection logic itself can be read, and so can the documentation behind it. A periodic review, a supervisor's request and an analyst questioning a result all draw on the same material, so an institution that keeps it properly only has to keep it once.

5

Performance measured on evidence

Detection is tested against known cases. False and true positive rates are tracked over time. Model behaviour is compared against how it behaved at validation, so drift shows up before it reaches outcomes. Data quality is checked as the data comes in, not once a problem surfaces. These figures come from measurement built into the framework, and they cannot be reconstructed afterwards from case files.

6

The framework learns from what it finds

A model built on what legitimate behaviour looks like flags anomalies that no existing rule describes, which is how activity outside the known typologies gets found at all. What the investigation makes of those anomalies goes back in, together with sampling results and any change in the risk assessment. Regenerating the logic on current data is a shorter route than rebuilding the framework, which is what lets an institution keep up as the risk moves.

Together these three determine whether an institution can show that its monitoring framework works and keeps working. They belong at the centre of the decisions that follow: how to **evidence effectiveness**, how to **organise validation and review**, and how to **answer for the framework** when a supervisor asks.

How Sygno addresses this

The Sygno platform generates transaction monitoring models from an institution's own data. Because the logic is generated rather than configured, the record of how it was built, what it does and why it changed comes out of the same process that produces the model.

Complete, accurate and indisputable audit logs

Each alert carries the features and values that triggered it, so the reason travels with the alert instead of having to be reconstructed. Model documentation covers the method, the validation and the output, and it is available without a preparation exercise. Every regeneration and every change to the logic is versioned, with its rationale and effective date, which is what a periodic review or an external validator works from.

Performance measurement built into the platform

Sygno measures detection against known historic cases and reports false and true positive rates over time. Model behaviour is compared against its behaviour at validation, so drift is visible before it reaches outcomes, and data quality is checked as part of the process rather than as a project that precedes it. All of it is available in dashboards that can be configured to what an institution wants to follow, and the same data generates reports for internal and external stakeholders without a separate extraction step.

Easy model regeneration that takes in what was learned

Sygno models what legitimate behaviour looks like for each customer, so anomalies surface without a pre-defined risk typology, including behaviour that no existing rule was written for. Investigation outcomes, sampling results and changes to the risk assessment feed the next regeneration, which runs on current data without rebuilding the framework around it. Regeneration is a routine step rather than a project, so the logic can follow the risk at the pace it moves.

Compliance with operational value

AMLA asks a great deal of the monitoring domain, and meeting it takes investment. Generating models from an institution's own data answers what Article 26 and Guideline 2 require, and at the same time changes what the framework finds: more confirmed criminal activity detected, and substantially fewer false positives. Monitoring and reporting take less effort, because the figures an auditor or supervisor asks for are the figures the platform already produces.

About Sygno

Sygno provides automated model generation for transaction monitoring, helping financial institutions and other obliged entities detect financial crime more effectively and efficiently. By modelling the good, expected behaviour of customers, Sygno makes suspicious activity stand out, drastically reducing false positives while increasing the detection of unusual and suspicious transactions. Models are generated from your own data, are transparent and explainable to supervisors, and deploy into your existing transaction monitoring system without migration.
Know good, catch bad.

Contact



Sjoerd Slot

Chief Executive Officer

Sygno
www.sygno.com
getincontact@sygno.com
linkedin.com/in/sjoerdslot/

Zonnehof 40, 3811 ND Amersfoort
The Netherlands