




W H E R E F R A G M E N T E D

AML & FCC STACKS

BREAK

Five continuity tests for deciding what to integrate, repair, or replace. An operating framework for anti-money laundering (AML) and financial crime compliance (FCC) leaders and the teams that run their financial crime infrastructure.

Contents

Executive summary	02
Continuity is the test	03
How a stack becomes fragmented	03
Five tests for continuity	04
The fragmentation tax: reconstruction and handoffs	05
The fragmentation tax: change and oversight	06
Follow the complete control lifecycle	07
Where operating pressure exposes gaps	08
When modularity works	09
What the operating layer must connect	10
Modernize without losing coverage	11
Twelve questions for an architecture review	12
Sources and scope	13
Build a lifecycle you can explain	14
Flagright	14



Executive summary

Fragmented compliance stacks break at the handoffs between systems. Each tool can work as designed while investigators rebuild context by hand, ownership blurs between teams, and reviewers lose sight of what was known when a decision was made. This report calls that rework and loss of visibility the fragmentation tax.

The tax shows up in four places: reconstruction, handoffs, change, and oversight. Tool count doesn't predict it. A modular stack can avoid it, and a consolidated platform can still pay it.

The report sets out five tests for continuity.

/1 CONTEXT

Context asks whether an investigator can assemble the full customer picture without reconciling identities by hand.

/2 CONTROL

Control asks whether the team can explain what each control covers.

/3 WORKFLOW

Workflow asks whether work keeps an accountable owner through every handoff.

/4 GOVERNANCE

Governance asks whether leaders can see gaps using consistent definitions.

/5 EVIDENCE

Evidence asks whether a reviewer can recover the inputs and rationale behind a decision.

It then covers when a modular architecture works, what a shared operating layer must connect, a six-step sequence for modernizing without losing coverage, and twelve questions for an architecture review.

WHERE TO START

Trace one recent case through every system it touched, including at least one exception. The breaks in that trace show where to repair first.

Continuity is the test

Consider a single handoff. An alert passes cleanly from one system to another, and the investigator still has to rebuild the customer's story, because the information needed for a sound decision stayed behind. The fragmentation tax names this rework. We make no estimate of its industry-wide cost.

The measure that matters is whether the institution can follow a control from the activity it covers, through investigation and decision, to later review.

The five tests turn that measure into something procurement and transformation teams can check before committing to an architecture. They separate problems that need better integration from problems that need new processes or new technology.

The framework is our own. It draws on the FCA's guidance on holistic monitoring and feeder-data failures, the FFIEC's description of suspicious activity reporting as a set of interdependent components, and the Wolfsberg Group's work on the transition to newer monitoring approaches. Each carries different authority in a different jurisdiction, and none of them treats platform choice as a substitute for effective controls. [\[1\]](#) [\[2\]](#) [\[3\]](#)

How a stack becomes fragmented

Take an institution expanding its payments business. It adds a screening service for onboarding, introduces a monitoring engine for a new product, and acquires a business with its own customer records. A case management tool pulls alerts into a common queue. Each addition solves an immediate problem.

The trouble sits at the boundaries. The acquired business keeps its own customer identifier. A risk score changes in one system and stays stale everywhere else. Screening and monitoring hold different versions of the same address. The investigator can open all four tools and still has to work out which record is current, and what was known at the time of the activity.

TRANSACTION MONITORING

 Alert #2345

RULE TRIGGERED	High-velocity transfers
TRANSACTIONS	14 in 48 hours
COUNTERPARTIES	3 new beneficiaries
CHANNEL	Faster payments



CASE MANAGEMENT

 Alert #2345

CASE	#C-2345 · Open
ASSIGNED TO	Analyst queue
RULE TRIGGERED	
TRANSACTIONS	
COUNTERPARTIES	



Modularity doesn't cause this. Unsettled boundaries do. Fragmentation begins when a capability arrives before the institution has decided how identity, updates, exceptions, and responsibility will move between components. An API can acknowledge a request before the receiving workflow has processed it. That acknowledgment confirms delivery. It says nothing about whether the message carries the right customer, the complete activity, or an actionable next step.

Every material connection needs a working contract: what information crosses it, what each field means, when updates must arrive, how rejected records are surfaced, and who resolves failures. The full set of contracts appears under "When modularity works."

Look at what crosses each boundary, and at the work people do when it fails to arrive.

Five tests for continuity

Apply the tests to a defined product, customer population, jurisdiction, and period. Trace ordinary work and the exceptions that matter, including missing data, reassigned work, and reopened cases. The easiest path through a demo proves very little about the full process.

//	TEST	QUESTION TO ANSWER	EVIDENCE TO INSPECT
01	Context	Can an authorized investigator assemble the relevant customer and activity history without manually reconciling conflicting identities?	Source identifiers, relationship mappings, freshness indicators, and the case view.
02	Control	Can the team explain which risks and populations each control covers, including exclusions and changes?	Risk-to-control mapping, input reconciliations, versions, and test results.
03	Workflow	Does work keep an accountable owner and visible status through transfers, approvals, and exceptions?	Queue history, handoff acknowledgments, escalation records, and deadlines.
04	Governance	Can leaders identify gaps and bottlenecks across components using consistent definitions?	Reconciled management information, exception logs, and change ownership.
05	Evidence	Can a reviewer recover the material inputs, rationale, and authority behind a decision?	Point-in-time evidence, decision history, approvals, and governed access.



Record what each trace demonstrates and where it stops, using the ratings at the end of this report. Where a workaround exists, test its reliability and make its operating burden visible.

These tests are built for architecture review. They measure whether the arrangement serves the institution's actual risks and obligations, and they sit alongside, rather than inside, any regulatory rating.

The twelve review questions turn these tests into an agenda. Use the findings to identify specific repairs before deciding which systems to keep, connect, or replace.

The fragmentation tax: reconstruction and handoffs

A monitoring alert flags unusual activity. The analyst pulls onboarding records, searches a separate screening history, checks related accounts, and asks another team why a risk score changed. Each lookup pulls attention away from the actual work of assessing the activity.

That is reconstruction work. Some of it is legitimate, since some information belongs outside the case tool. The avoidable part is the repetition: the same lookups run again and again, inconsistent identifiers reconciled by hand, and information copied manually when a governed connection could supply it with its source and timestamp intact.

Handoff cost appears when the investigation moves to another team. The queue transfer succeeds, yet the receiving analyst can't see what has been checked, who owns the next action, or which deadline applies. The second reviewer repeats the first review, or the work waits without an owner.

A good transfer carries a linked evidence record, the question still open, the last action and its rationale, and a named owner for the next step. Sensitive information can stay in a restricted source system as long as authorized reviewers can reliably retrieve it. The goal is reliable access to the right record, wherever it lives.

To measure this burden in your own institution, observe a defined sample of investigations. Separate the time spent gathering context, waiting for information, analyzing activity, and completing required review.

Log repeated lookups and failed transfers. Compare similar work before and after a change, and include quality and coverage checks. Faster case closure on its own proves nothing about investigation quality.

The fragmentation tax: change and oversight

A change to a payment feed reaches well beyond the database that receives it. Suppose a product changes how it records transaction status. If one control includes pending activity and another counts only settled transactions, a single mapping change can shift coverage while the alert queue looks exactly the same.

Change cost is the work of finding every affected dependency, updating it correctly, testing it, and recording why. It rises when teams maintain separate definitions or can't tell which reports and decisions depend on a given field.

The FCA gives a concrete supervisory example: feeder-system changes and migrations can cause transactions to drop out of monitoring. Its guidance also describes aggregating activity across different levels. Both points argue for checking data coverage along the entire control chain, whatever the underlying architecture. [\[1\]](#)


/	FOLLOW THE CHANGE	CHECK BEFORE ACCEPTING IT
	Source and mapping	Did the expected records arrive, with the intended meaning and timing?
	Detection and risk	Which populations, rules, or models changed, and were the affected scenarios tested?
	Investigation	Do case views, prioritization, and analyst procedures reflect the change?
	Oversight and evidence	Can reports and reviewers distinguish the old configuration from the new one?

Oversight becomes expensive when this chain has to be rebuilt for every review. A dashboard can add up totals while hiding incompatible definitions of an alert, a case, a backlog, or a completed investigation. Agree on the definitions first, then combine the numbers.

Give one person or function responsibility for assessing the full effect of a change, and keep the specialists who own each component. A shared change record should link the affected controls, test results, open exceptions, and the decision to proceed. Dependencies should be visible before a gap reaches production.

Follow the complete control lifecycle

Monitoring is only as good as the work that follows a signal. The institution has to establish what happened, assess it in context, make the relevant decision, take any required action, and keep an appropriate record for later review. The FFIEC's BSA/AML Examination Manual describes five interdependent components: identifying unusual activity, managing alerts, deciding whether to file a suspicious activity report (SAR), filing where appropriate, and monitoring continuing activity. Its discussion of investigations ties analysis directly to customer due diligence and other relevant information. [2]



Alerts lead to a range of outcomes. An investigation may end in closure, further inquiry, escalation, or a reporting decision, and screening and other controls follow their own rules. The diagram shows the shape of a reviewable lifecycle. Each control will trace its own path through it.

Continuity means each stage can explain its inputs and outcome to the next authorized participant. A decision should point to the evidence it relied on and the authority under which it was made.

Later changes should remain clearly separate from what was originally considered.

Feedback closes the loop. Findings from investigations and quality assurance expose missing data, weak prioritization, and controls that need tuning. Those findings belong in governed change review, where they can improve detection deliberately. Left ungoverned, they can quietly rewrite detection logic or turn one analyst's judgment into a universal training label.

A NOTE ON DOCUMENTATION

The FFIEC chapter carries a notice that it is awaiting updates, and the October 2025 interagency FAQs state that the Bank Secrecy Act and its implementing regulations do not require or expect institutions to document a decision not to file a SAR. Evidence design should follow applicable requirements and institution policy. This framework adds no recordkeeping obligations of its own. [2] [4]

Where operating pressure exposes gaps

Business change reveals whether interfaces preserve the meaning of what they carry. Use a planned launch or system change to test the stack before relying on its existing integrations.

/	CHANGE	QUESTION FOR THE CONTROL REVIEW
	New product	Do the new customer and transaction types reach the intended controls?
	Faster payment rail	Are data freshness, decision timing, and exception handling suitable for the relevant control?
	New jurisdiction	Can the workflow preserve local responsibilities, access restrictions, and reporting requirements?
	New typology	Can teams identify affected populations, test detection changes, and explain any remaining gaps?
	Feeder-system migration	Can expected activity be reconciled across extraction, transformation, and ingestion?
	AI in investigations	Can reviewers trace generated material to its sources and identify unsupported conclusions?

AI deserves particular attention, because fluent output can hide missing inputs. An assistant that only sees one case tool can produce a readable summary that leaves out relevant evidence held elsewhere. The risk comes from the architecture, so that is where to test it.

The Wolfsberg Group’s guidance on transitioning to innovative monitoring approaches covers integration with the wider control environment, data availability, investigator readiness, governance, and explainability. It is industry guidance rather than binding regulation. [\[3\]](#)

For any AI-supported workflow, test incomplete inputs, conflicting identities, stale data, and failures to reach a source. Make uncertainty visible to the reviewer. Keep the evidence and system details needed to examine an output, within policy and retention requirements. Define which actions need human approval and which the system is authorized to take.

Automating a handoff helps only when the receiving process can interpret and act on what it receives. Otherwise, automation moves the same gap through the workflow faster.

When modularity works

A specialist component may offer capabilities, local coverage, or deployment options that a broader platform can't match. Replacing it can weaken the control, or cost more in migration than the problem is worth. Consolidation brings dependencies of its own: on a vendor's roadmap, its commercial terms, and its support for exit.

A coherent modular architecture is a legitimate target. It requires explicit agreement on how information and responsibility move between components, backed by tests that keep running after implementation.

/	CONTRACT	WHAT IT NEEDS TO ESTABLISH
01	Identity	How entities map across sources, how uncertain matches are handled, and how corrections propagate.
02	Data	Field meaning, lineage, freshness, completeness checks, and visible treatment of rejected records.
03	Workflow	Ownership, state changes, acknowledgments, retries, and protection against duplicate actions.
04	Decisions	How evidence, rationale, versions, and approvals remain available to authorized reviewers.
05	Change	Who assesses cross-system effects, verifies coverage, and resolves exceptions.
06	Exit	Whether data, relationships, case history, and decision records can be exported in usable form.

Evaluate these contracts as seriously as detection features. Ask the team or vendor to demonstrate an identity correction, a failed transfer, a reopened case, and the export of a completed investigation.


The tradeoff is real. Maintaining these contracts takes engineering and operational capacity. A shared platform can absorb some of that work, and it still needs reliable connections to source systems and a clear operating model. Compare the full burden of each option, including retention, resilience, migration, and exit.

Consolidate where it improves the lifecycle you need to support. Keep specialist components where their value justifies the coordination they require.



What the operating layer must connect

Every stack that works has an operating layer: the shared capabilities that keep context, work, and evidence coherent across controls. It can live inside a platform, in coordinated services, or in a mix of both. Its value is measured by what teams can demonstrate in practice.



Start with identity and data meaning. Preserve source records, explain how related entities are linked, and make data quality failures visible. A common customer view should keep the differences that matter, including legal entities, jurisdictions, and when each fact became known.

Connect that context to monitoring, screening, and risk assessment, and respect the different decision timing and purpose of each control. Investigation workflows then need reliable ownership, access to the relevant evidence, and explicit authority for consequential actions.

Governance spans the whole arrangement. Teams need consistent definitions, visible exceptions, and a way to trace a change through every affected component.

Evidence should remain retrievable under the institution's access, retention, and legal requirements, wherever it is stored.

This architecture creates dependencies of its own. Plan for identity resolution going down, a provider sending incomplete data, and the common case service failing. Define degraded operation, recovery, reconciliation, and accountability in advance. Concentrating work in one layer also concentrates the consequences of failure.

Acceptance should rest on representative traces, coverage reconciliations, and tests of material failure paths across the agreed scope. A diagram shows connections. Evidence shows what those connections preserve.

Modernize without losing coverage

Start with the parts of the lifecycle that create the most exposure or avoidable work, and keep each transition bounded and testable. Tailor the exit conditions below to the institution's own risks and obligations.

/1 MAP THE CURRENT PATH

Trace activity through controls, investigation, and outcomes. Include manual steps, open work, and external dependencies. Move on when coverage, ownership, and known exceptions are recorded for the chosen scope.

/3 RECONCILE THE DATA

Test identity mappings, completeness, field meanings, freshness, and access to history. Include rejected, duplicated, late, and corrected records. Move on when material discrepancies are resolved, or explicitly assessed and placed under controlled treatment.

/5 TRANSFER WORK AND EVIDENCE

Reconcile open cases, deadlines, ownership, restrictions, and decision history. Test retrieval of older evidence and the recovery or fallback arrangements. Switch responsibility once the receiving process handles both ordinary work and the in-scope exceptions.

/2 CHOOSE THE FIRST REPAIR

Weigh the consequence of each break against its frequency, its operating burden, and the reliability of the current workaround. Define the expected improvement and the checks that would reveal deterioration. Set acceptance criteria before looking at results.

/4 RUN A CONTROLLED COMPARISON

Compare the current and proposed processes on relevant activity and representative exceptions. Judge differences against approved control objectives and test cases. Matching the old process proves only that the two agree. Look at coverage, quality, and workload. If you run both processes in parallel, designate one authority for live actions so they can't produce duplicate filings or conflicting interventions.

/6 VERIFY BEFORE RETIREMENT

Check live reconciliation, unresolved discrepancies, backlog, and decision quality against the criteria agreed in step 2. Retire the old component once coverage and required records remain available, responsibilities are settled, and recovery arrangements have been tested.

The transition is complete when the control still works and can still be explained.



Twelve questions for an architecture review

Define the scope and review period. Ask for evidence against each question, including difficult cases and relevant exceptions. For each gap, record an evidence reference, an owner, and a next action.

//	TEST	QUESTION
01	Context	Can we follow one customer and its relevant relationships across the systems involved?
02	Context	Can an authorized reviewer distinguish current information from what was known at the time of the decision?
03	Context	Can we detect and correct an identity mismatch, a stale input, or a missing record?
04	Control	Can we trace an identified risk to its controls, covered populations, exclusions, and tests?
05	Control	Can we show how a source or configuration change affected coverage?
06	Workflow	Can we identify the owner, status, and relevant deadline at each handoff?
07	Workflow	Can we recover from a failed transfer without losing work or duplicating an action?
08	Governance	Can we reconcile management reporting to the underlying work using agreed definitions?
09	Governance	Can we identify who accepts a cross-system change and how unresolved gaps are handled?
10	Evidence	Can an authorized reviewer retrieve the material evidence, rationale, and approvals for a sampled decision?
11	Evidence	Can we explain the material rule, model, or AI output involved, including its relevant version and limits?
12	Evidence	Can we preserve required records and access restrictions through correction, migration, and vendor exit?



	/Rating	/Meaning
	Demonstrated	The trace and relevant exception tests support the answer within the stated scope.
	Qualified	A bounded limitation remains, with a compensating process that has been assessed and tested.
	Broken	Observed evidence shows a material break in the defined process.
	Not demonstrated	Evidence is missing or the test hasn't been run. Investigate before drawing a conclusion either way.

Keep these ratings separate. Adding them into a maturity score hides what matters most. Prioritize findings by their consequence and by the reliability of current mitigations. One serious coverage gap can matter more than several inconvenient lookups.

Sources and scope

[1] **UK Financial Conduct Authority.**

Financial Crime Guide, FCG 3.2.5A G, “The use of transaction monitoring,” introduced 29 November 2024 (PS24/17). See the examples-of-practice table for holistic monitoring, activity aggregation, and omissions after feeder-system changes. UK supervisory guidance; not a prescribed technology architecture.

[Read the FCA Financial Crime Guide](#)

[2] **Federal Financial Institutions Examination Council.**

BSA/AML Examination Manual, “Suspicious Activity Reporting,” particularly “Systems to Identify, Research, and Report Suspicious Activity” and “Managing Alerts.” US bank examination context. The chapter carries an update notice referring to newer interagency FAQs; use those FAQs for the topics they supersede.

[Read the FFIEC Suspicious Activity Reporting chapter](#)

[3] **The Wolfsberg Group.**

Statement on Effective Monitoring for Suspicious Activity, Part II, “Transitioning to Innovation,” 27 August 2025. See section A, pages 2–5, on transition and control integration, and section C, pages 8–10, on explainability. Non-binding industry guidance.

[Read the Wolfsberg statement](#)



[4] **FinCEN, Federal Reserve, FDIC, NCUA, and OCC.**

Frequently Asked Questions Regarding Suspicious Activity Reporting Requirements, 9 October 2025. See question 4, page 4, on documenting decisions not to file, and question 2, pages 2–3, on continuing activity reviews. Interpretive FAQs that do not create new requirements.

[Read the October 2025 interagency FAQs](#)

[5] **Flagright.**

Company information, used solely for the product description on the back cover. First-party product information, not independent evidence of efficacy.

flagright.com

Sources checked 24 September 2026. An institution's obligations are determined by the rules in its jurisdictions and by its own circumstances, and this report's architecture recommendations should be read alongside that assessment.

Build a lifecycle you can explain

Before adding a component or replacing one, ask the team to trace a real investigation through the proposed arrangement and run the twelve questions against it. Follow the identity, the information received, the decisions made, and the evidence retained. Include at least one exception and its recovery.



The AI operating system for financial crime compliance

Flagright brings transaction monitoring, watchlist screening, customer risk scoring, investigations, case management, and governance into one unified, risk-based platform. It embeds AI across alert investigation, recommendations, and system optimization while preserving human oversight. Customers can adopt the capabilities they need and expand across the platform over time. [5]

Talk to us about the control lifecycle your institution needs to support, the components you plan to keep, and the evidence you'll need to validate a change.

GET A DEMO ↗

