CONSORTIX

Aurora

AML Unified Rule Orchestration and Risk Agent

Speed up your scenario development with agentic AI

How Consortix and SAS Viya compress the AML scenario lifecycle from weeks of manual work into a supervised, auditable workflow built around AI agents.

A modular agentic AI suite for supervised AML scenario creation, validation, recalibration, and investigation support.

CONSORTIX

SUMMARY

AML detection has an execution bottleneck

For over two decades, anti-money laundering teamsAurora attacks the manual translation layer between compliance expertise and executable AML scenarios.have faced the same operational problem. Billions of transactions flow through financial institutions every day, and only a small fraction carry genuine criminal activity. Finding that fraction is a detection challenge. Keeping detection current as criminal patterns evolve is an execution challenge, and this is where most compliance programs lose the most time.

Scenario development, the engine of rule-based AML detection, has remained a manual, iterative process. A new rule routinely takes weeks to specify, implement, test, and deploy. Every iteration consumes scarce analyst and IT capacity. Machine learning shrank the haystack. The handwork of building and maintaining rules stayed the same.

Aurora is the answer Consortix developed for this bottleneck. Built on SAS Viya and operating alongside SAS Anti-Money Laundering and SAS Financial Crimes Analytics, Aurora is a modular agentic AI suite that turns plain-language rule descriptions into production-ready scenarios, validates them against synthetic and historical data, and supports investigators with enriched cases and narrative drafts. Deployment timelines shrink from weeks to days. Every decision Aurora makes is logged, timestamped, and reviewable by the compliance expert who signs off on the rule.

Traditional scenario creation time.

WEEKS

Aurora target deployment window.

DAYS

Business-to-IT handoffs required per rule.

NONE

Audit trail generated by default.

AUTO

CONSORTIX

AML EFFICIENCY

Scenario-based AML is strong, but slow

Every new risk pattern has to become a precise, testable, governed rule before it can protect the institution.Modern AML programs run on scenario-based detection. Each scenario is a codified rule that scans transaction data for suspicious behaviour. Designing, tuning, and maintaining these rules is where most compliance teams spend the bulk of their development capacity. The process has a predictable shape in almost every institution.

Detection is only half the challenge

Keeping detection current is where compliance programs lose time.

A business expert identifies a new risk pattern. They draft a written description and hand it to IT for technical specification. IT produces an implementation. Quality assurance runs test data through the rule. The business reviews results, flags edge cases, and requests adjustments. The cycle repeats until approval, and only then does the rule reach production. Six handoffs between business and IT is a typical count, and eight is common. In a mature AML program with dozens of scenarios under active maintenance, the gap between identifying a threat and having a rule in production becomes the single largest drag on compliance effectiveness.

CONSORTIX

MANUAL LIFECYCLE

A bunch of handoffs before production

1. **Risk pattern**
   Business expert identifies a new typology.
2. **Specification**
   Description is translated for IT.
3. **Implementation**
   Rule logic is coded manually.
4. **Quality assurance**
   Test data is prepared and executed.
5. **Business review**
   Edge cases are reviewed and adjusted.
6. **Approval**
   Rule reaches production after iterations.

Where the drag appears

Machine learning has helped at the margins. False positive rates have improved with risk scoring models, and analyst attention is directed at higher-value alerts. The manual scenario construction work has stayed almost untouched. Rules are still written in natural language by business experts, translated into logic by hand by developers, tested on sample data, and interpreted through review meetings. The bottleneck sits in the translation from business knowledge to executable code, and that bottleneck is where Aurora does its work.

CONSORTIX

HOW AURORA WORKS

From plain language to deployable AML logic

A compliance expert stays in control while specialised agents turn intent into tested, governed scenario logic.

EXAMPLE NATURAL-LANGUAGE INPUT

"flag offshore transactions above 10,000 EUR when they follow more than three inbound transfers within 72 hours"

Aurora is an agentic AI suite that compresses the entire scenario lifecycle into a single supervised workflow. A compliance expert writes what the rule should do in plain English. A request such as "flag offshore transactions above 10,000 EUR when they follow more than three inbound transfers within 72 hours" is enough to start the process. Aurora takes the sentence and runs it through a chain of specialised agents, each with a narrowly defined job.

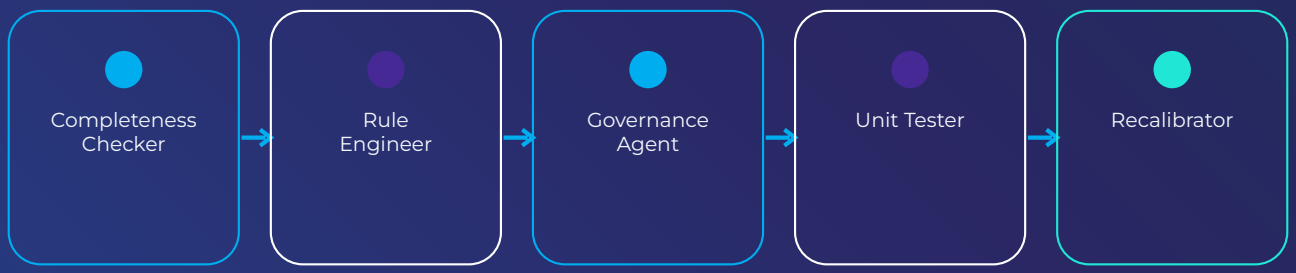

Five agents make up the core of the scenario creation pipeline. Each agent hands its output to the next, and the compliance expert can inspect every intermediate artifact.

At the end of this chain, a deployable AML scenario exists with a complete audit trail attached. The compliance expert reviews the output and publishes it. Work that once took several weeks now takes a few hours of focused review.

CONSORTIX

THE AGENTS

A specific role for every agent

The pipeline is modular: each agent performs one auditable task and hands a structured artifact to the next.

Completeness Checker

Reads the natural-language description and verifies that all the information required to build the rule is present. If something is missing or ambiguous, it asks targeted questions and records the clarified input.

Rule Engineer

Translates the confirmed description into formal, system-readable rule logic. The output is technically implementable, fully documented, and structured for audit.

Governance Agent

Checks whether the new rule overlaps with existing scenarios and verifies that the rule contains no discriminatory or legally impermissible filtering conditions. Regulatory and ethical compliance is enforced at the point of creation.

Unit Tester

Generates test cases automatically, runs them against the proposed rule, and produces a pass/fail report. Where test results suggest refinement, it feeds observations back into the logic.

Recalibrator

Updates the AI models already in production with fresh data produced by the new rule, keeping the broader detection system aligned with the latest logic.

CONSORTIX

VALIDATION AND TRUST

Designed for audit from the first line of the spec

Aurora treats governance, synthetic tests, and human sign-off as part of the workflow - not as a post-production In AML, every output has to stand up to regulatory clean-up step. scrutiny, and Aurora was designed around that constraint from the first line of the spec. Every rule Aurora produces carries synthetic test data, documented guardrails, and a timestamped record of what each agent did and why. If a regulator asks how a rule was built, the answer is already prepared. If a rule needs to change, the expert edits the input, and the same validation chain runs again. Production deployment always runs through the same checks.

Each intermediate artifact is readable, exportable, and editable. The clarified description, the structured logic, the test cases, and the test results are all available to the compliance expert at any point in the workflow. Human oversight is the default state of the system.

Readable artifacts

Clarified description, logic, tests, and results.

Timestamped decisions

Each agent action is logged and reviewable.

Human oversight

The compliance expert reviews and publishes.

CONSORTIX

INVESTIGATION SUPPORT

From generated alerts to better case narratives

Aurora also supports analysts after detection, where triage and SAR narrative drafting consume valuable investigation time. Aurora also operates at the other end of the AML workflow. Once alerts are generated, investigators face the triage problem. Aurora enriches each case with relevant context pulled from across the data estate, suggests a priority order based on configurable risk signals, and drafts narrative summaries that analysts can use as the foundation of a suspicious activity report. SAR narrative drafting is one of the most time-consuming tasks in an investigator's day. With Aurora it becomes a review-and-edit task taking minutes. The final disposition decision remains with the analyst, as regulation requires.

Context enrichment

Relevant data estate signals are brought into the case view.

Priority order

Configurable risk signals help investigators triage.

Narrative drafts

SAR narrative drafting becomes review-and-edit.

CONSORTIX

FROM WEEKS TO DAYS

The shift

A sequence of manual handoffs becomes a single guided session with full automation underneath.
The diagram below contrasts the traditional scenario development cycle with the Aurora workflow. The shift is structural. A sequence of manual handoffs becomes a single guided session with full automation underneath.
The numbers below reflect what Aurora achieves in typical deployments on SAS Viya. They combine observed workflow improvements with recalibration gains on the underlying prioritisation models.

Metric	Before Aurora	With Aurora
Scenario creation time	Weeks to months	Hours to days
Business-to-IT handoffs per rule	Six or more	None required
Iterations before deployment	Four to six	One to two
SAR narrative drafting	Hours per case	Minutes of review
Audit trail	Reconstructed manually	Produced automatically
Lift @ 10% (prioritisation model)	2.8	3.4
AUC-ROC	0.72	0.78
KS Statistic	0.42	0.49

Scenario creation speed is the headline metric. The recalibration numbers tell an equally important story. Aurora compresses development time, and it improves the quality of the downstream prioritisation models by feeding them richer, more consistent data. Institutions gain on efficiency and detection effectiveness in the same programme.

CONSORTIX

WHY AURORA WORKS ON SAS VIYA

Open platform, supervised automation

Most AML vendors ship closed solutions that can be configured inside their own boundaries and stop there. SASViya exposes APIs, supports custom models, and allows third-party components to plug into the core AML engine. Aurora takes direct advantage of this design. It runs alongside SAS Anti-Money Laundering and SAS Financial Crimes Analytics, adding a layer of automation and intelligence on top of what institutions already have in production. Aurora is designed to add an automation layer on top of existing SAS AML and Financial Crimes Analytics SAS Viya is an open, extensible analytics platform. This architectural choice matters for two practical reasons. First, institutions keep the investments they have already made. There is no rip-and-replace programme to justify, and existing scenarios, models, and integrations continue to operate as before. Second, Aurora grows with the organisation. New agents can be added for new use cases without re-architecting the detection stack. The same extensibility that makes Viya a foundation for Aurora makes Aurora a foundation for the next wave of compliance tooling.

Consortix domain expertise

Consortix brings the domain expertise that turns platform capability into working solutions. The firm's AML practice has spent years building on SAS foundations, and Aurora is the distilled output of that experience.

CONSORTIX

FOUR PRINCIPLES

How to carry Aurora forward

The architecture of Aurora rests on four principles that generalise beyond this single product.

01

Viya is built for customisation and extension. Vendor-locked solutions cannot keep pace with the speed of regulatory and criminal change.

02

Aurora is modular by design. Organisations can adopt the agents they need first, and add others as their AI readiness grows.

03

Bring your own AI, and use domain experts. The combination of platform flexibility and compliance expertise is what makes agentic AI deployable in regulated environments.

04

Start with the use case that matches your current AI readiness. Scenario creation, validation, and investigation support are each standalone entry points into the full workflow.

LOOKING AHEAD

Waiting is the expensive option

Aurora makes audit evidence a by-product of normal operation instead of a separate sprint after the rule is built. Agentic AI is already in production at financial institutions that are serious about the efficiency of their compliance function. Aurora is one concrete path to that state with governance kept intact from the first day. It reduces the time from regulatory signal to operational detection. It reduces the manual load on compliance teams. It produces complete audit evidence as a by-product of normal operation, so documentation work that used to sit in its own sprint now arrives with the rule itself.

The standing question

The standing question for every AML programme today is whether the cost of waiting weeks for a new rule still fits the current risk landscape. For programmes facing faster evolving criminal patterns, tighter regulatory deadlines, and constrained headcount, waiting is already the most expensive option on the table.

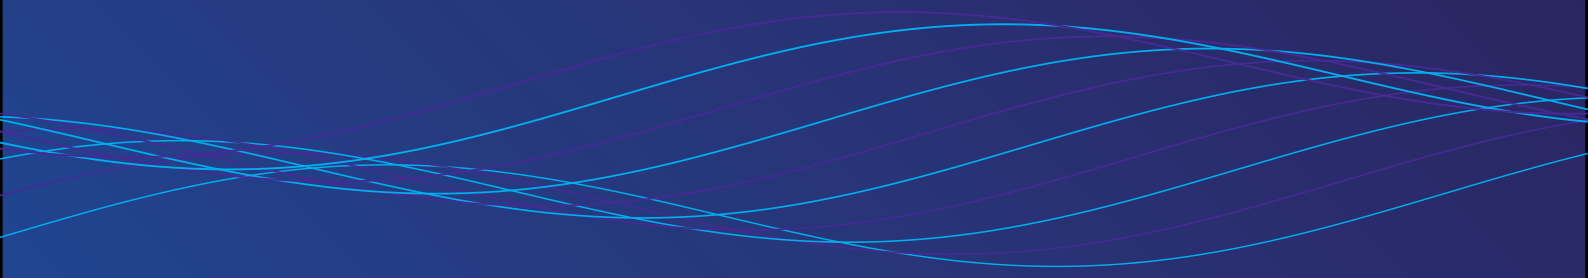

CONSORTIX

ABOUT CONSORTIX

Financial crime compliance expertise, delivered on SAS foundations

Consortix is an advisory, implementation, and support firm specialising in anti-money laundering, financial crime analytics, risk management, and customer intelligence. Founded in 2015 by specialists from banking, IT, and AML, the firm combines domain expertise with analytics technology delivery across Europe.

Consortix works with banks and financial institutions to implement, extend, operate, and fine-tune transaction monitoring, customer screening, customer due diligence, KYC onboarding, regulatory reporting, and related AML/CTF environments. Its practice builds on strategic technology partnerships including SAS and Lucinity, with Aurora as the flagship component of its agentic AI portfolio for compliance.

- Transaction monitoring
- Customer screening
- Customer due diligence
- KYC onboarding
- Regulatory reporting
- AML/CTF support

https://linkedin.com/company/consortix

https://www.consortix.com/aurora