CONSORTIX

CONSORTIX · WHITEPAPER · 2026

Aurora

AML Unified Rule Orchestration and Risk Agent

A modular agentic AI suite for supervised AML scenario creation, validation, recalibration, and investigation support. Built on SAS® Viya, operating alongside SAS Anti-Money Laundering and SAS Financial Crimes Analytics.

Speed up your scenario development with agentic AI

The Bottleneck

The execution problem behind AML detection

For over two decades, AML teams have faced the same operational problem. Billions of transactions flow through financial institutions every day, and only a small fraction carry actual criminal activity. Keeping detection current as criminal patterns evolve is an **execution challenge** — and execution is where most compliance programs lose the most time.

Scenario development has stayed a manual, iterative process. A new scenario routinely takes weeks to specify, implement, test, and deploy, consuming scarce analyst and IT capacity at every step. Machine learning shrank the haystack: false-positive rates improved, and analyst attention shifted to higher-value alerts. The handwork of building and maintaining scenarios stayed almost untouched.

The Manual Lifecycle

Every new risk pattern must travel the same path before it can protect the institution.

01
Risk Pattern
A business expert identifies a new typology.

02
Specification
The description is translated for IT.

03
Implementation
Scenario logic is coded manually.

04
Quality Assurance
Test data is prepared and run.

05
Business Review
Edge cases are reviewed and adjusted.

06
Approval
The scenario reaches production after several iterations.

Six to eight handoffs between business and IT is typical. The bottleneck sits in the translation from business knowledge to executable logic — that is where Aurora does its work.

	Before Aurora	With Aurora
Scenario creation	Weeks per scenario	Hours to a few days
Handoffs	Six to eight business-to-IT handoffs	Compliance expert and Aurora only
Implementation	Manual coding by IT, through a ticket queue	Scenario created through the platform API

What Aurora Is

Aurora is Consortix's own AI stack: agentic AI components built for compliance use cases. It sits on SAS Viya, alongside SAS Anti-Money Laundering and SAS Financial Crimes Analytics, and turns plain-language scenario descriptions into production-ready configurations. It validates those scenarios 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 scenario.

How Aurora Changes the Way You Work

Five use cases — from a scenario idea to a filed SAR

The five use cases below trace a scenario, and then an alert, through the full lifecycle. Each compares the manual state today with the supervised, Aurora-assisted workflow.

1 Scenario Creation

BEFORE

- Six to ten handoffs between compliance and IT
- Expert writes in Word; IT interprets and codes
- Interpretation risk at each translation step
- Two to six weeks from typology to production

WITH AURORA

- → Expert describes the pattern in plain language
- → Aurora generates a draft and returns its interpretation
- → Expert corrects or approves — no IT intermediary
- → API deployment; hours to a few days to production

2 Testing and Validation

BEFORE

- Narrow test sets — edge and boundary cases missed
- Manual regression; days between iterations
- No practical overlap mapping across hundreds of scenarios

WITH AURORA

- → Synthetic test data: positive, negative, boundary, and edge cases
- → Automatic regression across all logically related scenarios
- → Structured harmonisation proposal for overlapping or redundant logic

3 Documentation

BEFORE

- Configuration and documentation maintained separately
- Regularly falls out of sync; gap visible at audit
- Version history rare — lives in memory or a spreadsheet

WITH AURORA

- → Auto-generated from live configuration, updated on schedule
- → Links typology rationale to parameters; always in sync
- → Structured change log — full history retrievable at audit

4 Investigation

BEFORE

- Four to six systems; joins made manually by the analyst
- Related transactions assembled by hand
- Two to four hours per medium-complexity case

WITH AURORA

- → All data pre-assembled at alert open — no system switching
- → Pattern interpretation with structured explanation
- → Draft narrative ready; analyst reviews, completes, and approves

5 SAR Reporting

BEFORE

- Manual template filling; data re-gathered after investigation
- Quality varies by analyst; FIU returns and supplements common
- Statutory deadline pressure from manual assembly

WITH AURORA

- → Draft assembled from case data at investigation close
- → Consistent structure and terminology regardless of analyst
- → Direct submission to the FIU system or case platform

How Aurora Works

From plain language to deployable AML logic

A compliance expert writes what the scenario should do in plain English. A sentence like this is enough to start the process:

> *"Flag offshore transactions above EUR 10,000 when they follow more than three inbound transfers within 72 hours."*

Aurora takes the input and runs it through a chain of five specialised agents. Each agent has one job, produces one auditable artifact, and hands it to the next. The compliance expert can inspect every intermediate output.

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

The Five Agents

One job, one auditable artifact, one handoff

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

02 Rule Engineer
Translates the confirmed description into formal, system-readable scenario logic. The output is implementable, documented, and structured for audit.

03 Governance Agent
Checks whether the new scenario overlaps with existing ones and verifies that it contains no discriminatory or legally impermissible filtering conditions. Regulatory and ethical compliance is enforced at creation time.

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

05 Recalibrator
Updates the AI models already in production with fresh data from the new scenario, keeping the broader detection system aligned with the latest logic.

CONSORTIX

Validation and Trust

Designed for audit from the first line of the spec

In AML, every output has to stand up to regulatory scrutiny. Aurora was designed around that constraint from the start. Governance, synthetic tests, and human sign-off are part of the workflow — not a post-production add-on.

Every scenario 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 scenario was built, the answer is already in the trail. If a scenario needs to change, the expert edits the input and the same validation chain runs again.

Readable Artifacts

Clarified description, structured logic, test cases, and results are available at every stage — readable, exportable, and editable by the compliance expert.

Timestamped Decisions

Each agent action is logged with a timestamp and is reviewable. The full audit trail exists before production deployment.

Human Oversight

The compliance expert reviews and publishes. The final disposition decision stays with the analyst, as regulation requires.

Aurora removes the manual translation, testing, documentation, and assembly work that consumed compliance capacity — and returns that capacity to judgment.

CONSORTIX

Get in Touch

Interested in Aurora for your AML programme?

Reach out to schedule a conversation about Aurora and how Consortix can accelerate your AML scenario lifecycle.

Tamás Sváb

CAMS

CHIEF EXECUTIVE OFFICER

tamas.svab@consortix.com

+36 30 412 8575

Csaba Simonyi

ACCOUNT EXECUTIVE

csaba.simonyi@consortix.com

+36 20 919 4959

ADDRESS

Consortix Zrt.
Czuczor u. 2–10
1093 Budapest, Hungary

WEB

consortix.com
info@consortix.com