CONSORTIX

CONSORTIX · WHITEPAPER · 2026

Augmenting Your Investigation with AI

Where Detection Maturity Actually Sits

Machine learning has been part of compliance and financial crime for years, and teams have invested heavily in cutting false positives. Results vary widely. Some teams have done well; many have not.
Across the sector, false positive rates above 75% still affect close to a third of institutions, around 8% do not track their false positive rate at all, and roughly 30% have no settled KPI framework for measuring AI outcomes.
So detection maturity is uneven, and for a large part of the market the detection investment is still open.

>75%
false positive rate still affects close to a third of institutions

~8%
do not track their false positive rate at all

~30%
have no settled KPI framework for measuring AI outcomes

The picture changes for teams that have already put the work in. If you have run your tuning cycles and your behaviour-based segmentation, your remaining gains on detection are smaller, and the marginal return on further detection work is low.
The headroom narrows once the work has gone in. That is the segment this paper speaks to most directly.

This does not mean detection optimisation work should stop, for anyone. Our companion whitepaper, Improving Detection with AI, sets out the options by maturity level, starting with the easiest quick wins: alert triage, alert prioritisation, and transaction-based segmentation.
That is the subject of the other paper. The point here is narrower:

> ***Where the detection work has already been done, the next gain is more likely to come from investigation than from another round of tuning.***

Why False Positive Reduction Is Overrated for Mature Teams in 2026

Consider a mid-sized bank's AML team

BASELINE · MID-SIZED BANK · AML TEAM

- 2 hours per alert, from start to closure or SAR submission
- 4 alerts per investigator per day
- a 20-person team, which is 1,600 alerts a month

Cut investigation time to 1 hour and the same team can do one of three things.

01

It can clear 1,600 extra alerts from the backlog every month.

02

It can hold the same output with 10 people and move the other 10 onto improvement work rather than fire-fighting, whether that is tuning, scenario validation, or deep dives on the riskiest cases.

03

Or it can simply ease the load: half the pressure, and more time on the cases that deserve real attention.

Now try getting the same gain through detection tuning. At an 80% false positive rate, 1,600 alerts contain 320 true positives.
To match the gain, you would need to halve volume to 800 alerts while keeping every true positive. That means moving from an 80% false positive rate to 60%, and anyone who has tuned scenarios knows how hard a 20-point move is.

For teams that have already run their tuning and behaviour-based segmentation exercises, there is no easy win left on the table.
Regulators are also strict on the true positives you risk losing when you optimise detection. Investigation support with a human in the loop is a far easier conversation to have with them.

Halving investigation time delivers what a 20-point false positive reduction would, without the regulatory fight.

Why the Focus Is Moving to Generative and Agentic AI

The focus is shifting toward generative and agentic AI. The reason is accessibility.
These tools are now available to everyone. You no longer need to be an engineer, data engineer, or computer scientist to work with AI, and that changes what is possible for compliance investigation teams.

Detection optimisation is specialist work. It needs data science effort, and the marginal return is now low.
Investigation is where your team actually spends its time, and generative AI can take on a large share of that work without a specialist build.
The return on investment is therefore stronger here. The cost to start is lower and the time freed up is greater, which is exactly what the example above shows.

So the question becomes: are you investing in detection or investigation optimisation with AI in 2026?

How to Start

Once you accept the return is there, the question is how and where to begin.
This becomes an IT strategy question, a choice between build and buy, resting on two foundations: data integration and your compliance workflow.

Data Integration

Generative AI models start from the same base for everyone. The model or tool you choose matters less than you might expect.
Your advantage comes from the information and context specific to your institution and your knowledge, and from the way you make that available to the tool.

This is why data integration is essential. It determines which data sources and tools you connect, and how you teach your AI.
A useful way to think about it: treat the AI as a new colleague, and ask what you would give them to learn from and work with.

The best exercise, and the one we run with our customers, is to walk through how you investigate today.
Which systems do your investigators work in? Where do they reach out for specific information?
Those systems and touchpoints are your data sources, or your potential data sources.

One source many teams overlook is how you operate: your playbook, your rule book, your standard operating procedures.
Your current workflow and practices, your risk appetite, and the decisions your investigators make on alerts and cases, including how they turn those into outputs. This is some of the most valuable context you can give the tool.

The final step is to decide which of this information can be made available, and to set the scope of the data integration.

The Business Case

The second part is deciding where to start. A few questions help.

Q1 **Which of your processes can be described most clearly?**

Q2 **Which need the least complex human judgment and carry the least risk?**

Q3 **Which take the most time away from your investigators?**

The processes that score well across all three are where you should start. Typical good candidates are information retrieval for cases, and case summarisation and SAR narrative writing.

Build or Buy

The third question is the one that matters across banking software in the age of AI. It depends on your IT strategy and your resources, in plain terms whether you have more time and internal capacity or more budget.
It also depends on whether you see building these capabilities as a competitive advantage you want to keep developing, or whether your priority is an efficient compliance team as quickly as possible. Do you want to learn from the process, and in what way, or do you want to focus on results?

Those answers point you toward buying a tool or building on the IT infrastructure you already have.

BUILD ROUTE

Your existing AI infrastructure

The build route usually means your existing platforms for creating and running AI agents, which in most organisations sit largely within Microsoft.

BUY ROUTE

Specialised vendor or platform

If you lean toward buying, the next question is whether your current case management or transaction monitoring provider already offers these capabilities, or whether you would rather bring in a specialised vendor.

How Consortix Helps

Consortix is agnostic to these decisions. We are independent advisors.
Our only concern is to give you the best process and results, so you can optimise your AML detection and take your financial crime investigation to the next level.

For both routes we offer options, which we usually settle in a workshop. We start by analysing your current results with you.

Aurora

A Consortix AI methodology for AML investigation and scenario development

Aurora is a Consortix-developed methodology, configuration, and skill set that gives AML teams a concrete starting point for deploying agentic AI in their own environment. It covers two use cases. On the investigation side, Aurora aggregates case data from every connected source, flags the patterns that matter, and drafts both the case narrative and the SAR in the format the regulator expects. The compliance expert stays in control throughout; every agent output is logged, reviewable, and requires human sign-off before anything goes live.
On the scenario development side, Aurora replaces the six to eight manual handoffs between compliance experts and IT that currently slow every new detection scenario to weeks; a supervised pipeline of specialised AI agents turns a plain-English rule description into a tested, documented, deployable scenario configuration in hours, with synthetic test data and regression checks included. For further detail on scenario development with AI, see our companion whitepaper, Improving Detection with AI.

Aurora was built on SAS Viya and Python. The methodology and configuration are transferable to clients' own internal software environments, so teams that do not run SAS can still apply the same approach on their existing infrastructure.

If you want to develop in-house, we form a joint team with you and bring the knowledge you need. That is business consultancy and consultancy on how these use cases are built with AI, with examples and best practice. Where it helps, we also provide technical support, particularly on data integration. The advantage is cost saving through the IT infrastructure you already have, and a lasting lift in your internal AI team's knowledge from the external expertise brought in.

Run Your AI Agent as a Browser Plugin

Some teams want that knowledge delivered differently, as a ready-made solution they can deploy quickly. The most specialised solution we have found in this area is Lucinity, the Luci AI plugin, which runs as a browser plugin on top of your current AML system.
It deploys quickly and easily, it configures well, and we support you through implementation. The advantage is speed and access to one of the best ready-made solutions on the market.

The First Step

Wherever you decide to begin, the first step is the same: an assessment workshop to work out where you should start. That is how you begin to gain an advantage from your AI, and to free up your AML investigation team, starting tomorrow.

CONSORTIX

Get in Touch

Reach out to schedule an assessment workshop — we start by analysing your current results and working out where you should begin.

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