Aveni Detect · QA teams · Banking

Your QA programme should do two things: improve performance and evidence compliance. Most only do one.

Most tools tell you how well your team performed. Detect takes this further: evidencing whether each customer received a compliant outcome and giving you the audit trail to prove it. Built for QA teams in banking who need more than performance data.

1.38M
Conversations assessed by Detect

77:1
AI assessments completed for every 1 manual review

15 min
Average QA assessment time, down from 90 minutes

182K+
High-risk interactions missed at 2% manual coverage

01 WHERE DETECT IS DIFFERENT

Most tools tell you how well your people performed. Detect tells you whether the customer outcome was compliant, and provides the evidence to prove it.

Whether you're working with a QA platform, an add-on tool, or a manual review process, the same gap tends to exist: performance scoring tells you about team conduct, not customer outcomes. Detect is built specifically to close that gap. It assesses interactions across voice and every channel you connect to it, evidencing whether each customer received a fair and compliant outcome at the case level, across voice recordings, documents and written evidence. That's what Consumer Duty compliance actually requires.

02 WHAT CHANGES FOR YOUR TEAM

The same data. A fundamentally different question being asked of it.

Outcome evidence vs performance scoring

WHAT MOST TEAMS ARE WORKING WITH

> You get team performance scores across a sample of calls. You know which team members need coaching. But the tool isn't asking whether each customer got a fair outcome, and it isn't building the evidence trail to show the FCA they did.

WITH DETECT

> Every interaction is assessed for customer outcome, not just team conduct. The prioritisation algorithm surfaces fraud calls, restriction disputes and vulnerability signals where the outcome was not evidenced. Your analysts review the cases that represent genuine compliance exposure, with the case evidence already assembled.

Case reconstruction time

WHAT MOST TEAMS ARE WORKING WITH

> A scam dispute, an account freeze complaint, a financial difficulty call: each one requires you to pull a voice recording, find the case notes, locate any written communication, and reconstruct the full picture manually before the assessment can start.

WITH DETECT

> Voice and documents are assessed together as a single case. The full journey is already assembled when it reaches your queue. A review that previously took 60 to 90 minutes now takes around 15. You cover more of your highest-risk journeys in the same working day.

Scorecard consistency and coverage

WHAT MOST TEAMS ARE WORKING WITH

> Your Consumer Duty scorecard is applied by individual assessors across a sample of calls. Coverage is inconsistent between reviewers, and journeys that sit outside voice, including written case evidence and document communications, aren't assessed at all.

WITH DETECT

> Your scorecard is built directly into Detect and applied to 100% of interactions the same way, every time, across voice and documents. Assessor variation disappears. Written evidence is assessed as part of every case, not left outside the process.

Audit trail and evidence retrieval

WHAT MOST TEAMS ARE WORKING WITH

> When your CCO needs Consumer Duty evidence for the board pack, or an FOS case lands and someone needs the full interaction trail, your team manually reconstructs records across systems. It takes days. The output depends on who builds it.

WITH DETECT

> Every assessment is audit-ready as it's completed. Evidence is structured and retrievable in minutes. You hand your CCO a consistent output they can take to the board or the FCA, and you're not the one spending three days pulling it together before the deadline.

03 HOW DETECT COMPARES

Performance scoring and outcome evidence are different disciplines

Whatever your current QA setup looks like, the gap is usually the same: it tells you about team conduct, not customer outcomes. Detect is built to answer the question the FCA is actually asking. Here is how it compares on what matters most to a compliance-focused QA team in a regulated bank.



	Aveni Detect	Standard QA tool or process	Manual QA only
Primary question answered	✓ Did the customer get a compliant outcome?	Did the team perform well?	Did the team follow the script?
Team performance scoring	✓ Yes, with outcome context layered on top	Yes, typically the core purpose	Yes, but inconsistent across reviewers
Interaction coverage	✓ 100% of interactions	Often 100% for scoring, sample-based for deep review	Sample-based: coverage gap widens as volumes grow
Voice assessment	✓ Core strength: every call assessed, risk-prioritised for human review	Yes, typically the primary input channel	Yes, but limited to sampled calls
Multiple channel types	✓ Voice and documents assessed as one case	Often voice-focused, or separate configuration per channel	Manually reconstructed across systems per case
Consumer Duty outcome evidence	✓ Built in: every assessment produces FCA-ready output	Possible with configuration, not the core design purpose	Depends on individual assessor documentation
FS regulatory language model	✓ FinLLM, trained specifically on FS regulation and banking language	General NLP or rules-based, adapted for FS	Assessor knowledge
Audit trail for FCA or FOS	✓ Structured, explainable and retrievable by default	Often requires manual extraction and compilation	Manual reconstruction, inconsistent across

EVERY CHECK RETURNS AN OUTCOME, AN EXPLANATION AND THE EVIDENCE

"Were the customer's needs and objectives established?"

✓ Yes | No | Not applicable

WHY

The agent confirmed the reason for the dispute, checked the customer's understanding of the process, and agreed next steps before closing the call.

EVIDENCE

> Agent, 04:12 — "So I've got this right: the payment left your account on Friday and you'd like us to raise a dispute today?"

04 A MORNING IN THE QUEUE

Your caseload, prioritised before you sit down

Prioritised caseload · today · 847 interactions assessed automatically		● Live
VOICE	Conduct risk · #INT-6204 · authorised push payment fraud dispute · case note and voice assessed · outcome not evidenced · FOS risk	Review now
VOICE	Vulnerability signal · #INT-5891 · FG21/1 · account restriction call · financial difficulty indicators · appropriate support pathway not followed	Review today
DOCUMENT	Complaint risk · #INT-5773 · account closure communication · written explanation assessed · closure rationale incomplete · escalation likely	Review now
CASE NOTE	Consumer Duty · #INT-5601 · collections journey · SME arrears · voice and document assessed · fair treatment evidenced · scorecard complete	Complete

05 HOW IT WORKS IN PRACTICE

The detail evaluators ask for

01

Integration and ingestion

Detect connects to your existing telephony and document systems. No requirement to replace your CRM or case management layer. Detect ingests interactions directly from your existing infrastructure and works from there.

02

Configurable scorecards

Your Consumer Duty, vulnerability and conduct risk framework is configured into Detect's scorecard rather than replaced with a generic one. Your QA logic, applied at 100% coverage rather than across a sample.

03

Deployment and data residency

Fully UK-hosted with no customer data used for model training. Designed for regulated firms with strict data governance requirements. Your interaction data stays within your governance boundaries.

04

Human-in-the-loop review

Detect pre-assesses every interaction and routes the highest-risk cases to your QA team for human review. Your analysts make the final call. Detect handles the volume: for every 1 manual review your team completes, Detect completes 77 automatically.

FROM RAW INTERACTION TO AUDIT-READY EVIDENCE



Ingest — Voice, documents and case notes

Assess — Your scorecard, applied by FinLLM

Prioritise — Highest-risk cases surface first

Human review — Your analysts make the final call

Evidence — FCA-ready record, retrievable in minutes

06 CAPABILITIES

Built for how QA teams in banking actually work

Risk-prioritised queues across every journey type

Detect builds your caseload automatically from assessed interactions across fraud disputes, restriction calls, collections and vulnerability signals. Your analysts see the highest-risk cases first, with the evidence already assembled.

Your framework, applied to every interaction

Configure scorecards around your specific journey types: Consumer Duty outcomes, FG21/1 vulnerability, complaint risk, conduct. Applied consistently to 100% of interactions with no assessor variation between reviewers or teams.

Full journey review, not channel-by-channel

Voice recordings and written case evidence, including account restriction letters, collections documents and complaint correspondence, are assessed together as a single case. The full picture is already assembled when it reaches your queue.

FG21/1 across every interaction, including collections

Vulnerability and financial difficulty signals are identified automatically across restriction calls, fraud journeys and collections interactions. Every flag is evidenced and escalation pathways are built into the workflow rather than left to assessor judgement.

FOS and FCA-ready, without the manual build

Every assessment generates a structured evidence record as it's completed. When an FOS referral lands or your CCO needs board reporting, you retrieve what you need in minutes rather than spending days manually reconstructing records across legacy systems.

Pattern data your COO and CCO can act on

Detect surfaces which journeys, teams and scripts are generating repeat risk across fraud, restrictions and complaints. You take a data-backed picture to your leadership rather than a quarterly sample that covered 2% of what actually happened.

aveni

Trained on financial services language, not general text

FinLLM understands the language of banking interactions: authorised push payment disputes, account restriction calls, financial difficulty signals in collections, Consumer Duty outcomes in lending. It recognises the signals that matter in your highest-risk journeys without needing to be configured for each one from scratch. Aveni's scientists are consistently ranked among the top 100 NLP experts globally.

UK-hosted, secure, built for regulated firms

Fully UK-hosted with no customer data used for model training. Your interaction data stays within your governance boundaries. Detect improves continuously from real usage, so the more your team uses it, the more accurately it reflects the risk patterns specific to your journeys and customer base.

See Detect working against your own journeys and caseload

Book a demo and we'll run Detect against your own interactions across fraud disputes, account restrictions, collections or vulnerability calls, so you can see exactly what your team would be reviewing.

Ready when your caseload is

Book a demo