hubbis

WEBINAR REPORT SEPTEMBER 2026

AI in Banking: From Front-Office Friction to Practical Delivery

IN PARTNERSHIP WITH finalix

Artificial intelligence (AI) is already helping wealth teams draft, search and analyse information. Turning those capabilities into dependable improvements across a bank is harder. Client records sit in different systems, onboarding passes through several teams, and a faster answer still needs a clear source and an accountable reviewer. Know your customer (KYC) work brings those pressures together, making it a useful test of whether AI can improve the business process around an adviser as well as the adviser's own workload.

That was the setting for a Hubbis webinar with Finalix on practical AI in banking. The accompanying presentation deck maps the conversation from front-office pain points and use-case selection through KYC, data integration, human oversight and the move into production. Ronnie Chan and Andrew Rufener opened with Finalix's preliminary adviser-productivity research and an incremental approach to AI-enabled client lifecycle management. Michael Stanhope then moderated a Q&A with wealth practitioners from banks and a multi-family office. Their experiences gave the research a practical counterpoint: the same KYC challenge looks different when viewed from a large institution, a smaller office or the client facing repeated requests for information.

View part 1 of this webinar On Demand

View part 2 of this webinar On Demand

MODERATOR

Michael Stanhope
Founder and Chief Executive Officer, Hubbis

SPEAKERS

Andrew Rufener
Senior Manager, Data and AI, Finalix

Ronnie Chan
Senior Manager, Finalix

Chiara Bartoletti
Managing Partner and Chief Operating Officer, Eightstone

Ajit Vaidya
Head of International Wealth Management Operations, Nomura

Gourab Kundu
Head of Digital Growth and Wealth Continuum, Asia South, Citi

Key Takeaways

- Finalix's preliminary research identifies onboarding, KYC data capture, periodic reviews and information gathering as substantial demands on adviser time.
- The most promising initial uses combine a heavy, repeatable workload with a level of user trust that permits AI to assist without displacing judgement.
- A workable KYC application can sit alongside existing systems, structuring source information before generative AI prepares a summary for review, allowing rapid benefits without massive investment and substantial system changes.
- Institutional KYC policies, source-of-wealth complexity and data from multiple custodians prevent a single automated approach from fitting every firm.
- Adoption must be measured in changed work, better outcomes as well as improved fit to regulatory demands; access to a tool or a successful demonstration is insufficient.
- A human reviewer remains responsible for decisions, but needs enough expertise to challenge assembled evidence and resolve exceptions.

The Finalix Presentation: From Research to KYC Delivery

Ronnie Chan described Finalix's preliminary research as a productivity study involving more than 200 wealth management respondents across Europe and Asia. It followed the advisory journey from understanding the client through advice, execution and review. Instead of asking whether people liked AI, the research asked where work consumed the most effort and where assistance might relieve it.

Ronnie set out the starting principle: "The key principle is simple, we start with identifying business problems first, then discussing the technology solutions."

As presented, the findings concentrated on KYC data capture, onboarding, periodic review and the gathering of information. Documentation, manual entry and regulatory requirements drew effort away from client-facing work. The issue was larger than one adviser's productivity: fragmented records had to move among the front office, operations and control teams. Investment selection was not the main administrative burden the research revealed.

52%
identified KYC data capture as the biggest effort

45%
cited onboarding

38%
cited periodic reviews

76%
felt administration took too much of their time

48%
were ready to trust AI to support periodic KYC review

FINALIX PRELIMINARY RESEARCH, AS PRESENTED IN THE WEBINAR

Ronnie reported that 52% identified KYC data capture as the biggest effort, while 45% cited onboarding and 38% periodic reviews. He said 76% felt administration took too much of their time, and 48% were ready to trust AI to support periodic KYC review. These preliminary figures gave the argument a sharper edge: the burden sits in gathering and maintaining client information, precisely where many advisers are willing to accept assistance.

Ronnie linked that result to the design of the wider process: "The problem is not advisor capability, the problem is the complexity of the entire operating model."

> "The problem is not advisor capability, the problem is the complexity of the entire operating model."
>
> RONNIE CHAN, SENIOR MANAGER, FINALIX

Choose the Work Users Will Trust AI to Support

Finalix then compared two dimensions: how much effort a task takes and how comfortable users are with AI assistance. KYC capture, periodic reviews, research and information gathering appeared to combine a substantial workload with a willingness to use a tool. Ronnie also identified portfolio risk monitoring, data consolidation and regulatory document handling as activities where respondents showed some trust in AI support. Repeatable, information-heavy tasks with clear rules offer a more practical first step than asking a system to make an investment decision.

Ronnie expressed the distinction directly: “Advisors are not asking AI to make investment decisions and proposals for them. They actually want AI to help them remove administrative burden.”

> **“Advisors are not asking AI to make investment decisions and proposals for them. They actually want AI to help them remove administrative burden.”**
>
> RONNIE CHAN, SENIOR MANAGER, FINALIX

The proposed opportunity is therefore to make client information easier to collect, check and reuse. The return depends on the task chosen, the quality of its inputs and whether the output helps the next person in the chain. A high-effort activity is not automatically a good AI project if its evidence is too unreliable or its users cannot act on the result.

Build a KYC Use Case Around Existing Systems

Andrew Rufener moved from the research to implementation. Onboarding, periodic reviews and source-of-wealth analysis draw on large volumes of structured and unstructured material. Client relationship management, workflow and monitoring systems often hold different portions of the record. Andrew’s longer-term vision was a common engine able to support several views of the same client information, but an institution need not wait for that full architecture to begin.

03

KYC and CLM in Practice

Where document extraction, translation, rule-based checks and generative AI can improve onboarding and periodic reviews while preserving human approval.

FROM THE WEBINAR DECK · WHAT THE PANEL WILL COVER

A first use case can gather the material needed for one onboarding task, analyse it alongside established systems and return a usable result to the existing workflow. The same approach can then expand to another stage of the client lifecycle. A common client-data engine could ultimately support different views for name screening and KYC, but Andrew stressed that "it doesn't have to be a big bang". A wholesale platform replacement need not be the condition for progress.

Andrew cautioned against treating every step as generative AI: "And a lot of this stuff is not AI. Many of these technologies have been around for a long time."

> **"And a lot of this stuff is not AI. Many of these technologies have been around for a long time."**
>
> ANDREW RUFENER, SENIOR MANAGER, DATA AND AI, FINALIX

Andrew continued: "It's structuring data, extracting data the right way and establishing the relevant relationships before ultimately using Gen AI to generate a defensible and regulatory compliant output." The structured evidence and repeatable checks matter to the reviewer and to regulatory scrutiny. Generative AI can help prepare the summary, while the path from source document to conclusion remains visible.

That visibility is part of the efficiency case. A faster delivery of a client profile is of limited use if no one can explain where a facts originated, what decisions were taken or why an exception was granted.

The Q&A: Testing the Approach in Practice

Finalix's presentation identified where AI might support reducing repetitive manual efforts and how a KYC process can be supported defensibly with AI. The moderated Q&A widened the lens to the conditions that determine whether such an application helps in practice. The remaining themes in the webinar deck, data and integration, human oversight, and the route from pilot to production, brought the different institutions' experiences into focus. Their constraints mattered as much as the shared interest in saving time.

The Bottleneck Changes with the Institution

The panel broadly recognised KYC as a source of friction, but its practical form differed. In a multi-family office, the greatest burden was not necessarily the volume of account openings. It was the quality and consistency of data arriving from several custodians. A smaller team can sometimes change a process faster than a large bank, yet it has fewer specialists to build and govern new tools. AI-assisted development can make internal automation more attainable without solving every data problem beneath it.

The family office example also showed how ordinary uses of language models sit beside more ambitious process change. Drafting, translation and summarising can save time immediately. Building an internal automation may have greater lasting value, provided the small team understands the workflow it wants to improve and has enough expertise to maintain what it builds.

Banks face a larger and more complex operating chain. Expectations of prompt onboarding sit alongside detailed client-profiling requirements. Information collected at the start of a relationship must be refreshed as circumstances change, and several teams may handle the same record. The participants treated KYC as an end-to-end process in which AI could help, provided the wider workflow is also redesigned.

Institutional policies create another source of repetition. The same client document may prompt different questions from two banks; a family office can be asked to supply a passport more than once. Singapore's MyInfo service showed one way to reuse consented information in some applications, but it did not settle more complex source-of-wealth assessments. Different risk appetites and the facts behind a client's wealth still call for investigation and judgement.

The adviser also acts as the connection between the client, internal departments and outside providers. One participant described how this coordinating role becomes especially visible in KYC: a document may already be on file, yet the adviser still has to resolve different requests and interpretations across institutions. Improving one document-handling task could ease that burden, but the wider gain requires the information and the decisions around it to travel through the whole process.

Data Quality, Governance and the Route to Production

Participants described AI projects across employee productivity, operations and client-facing services. Their priorities varied, but each faced the question of which tools could be used safely and how to keep track of expanding capabilities. Approval processes, the handling of confidential information and clear human responsibility were central to moving a prototype into production.

Fragmented data can make an apparently polished answer unreliable. Records built up across systems or mergers may use different definitions or contain omissions. The firm must know what was extracted, what was inferred and which evidence the reviewer still needs to see. In the discussion, Ronnie warned that poor data can be made to look more convincing when fed through AI. Andrew emphasised the preparation needed to obtain meaningful, repeatable results that can withstand regulatory scrutiny.

An institution should assess the total effort across the process before choosing a solution. Several tasks may sit in separate functions, while a proposed product addresses only one. The panel favoured a clear business requirement, a controlled initial case and an explicit plan for adoption and measurement. Andrew argued that teams should work back from business and regulatory requirements and treat implementation as a change initiative involving the people who will use and review the result.

A small pilot can show a quick saving in preparation time without proving that the whole process has improved. To reach production, its output must feed an approved workflow, staff must rely on it consistently, and the institution must be able to measure what happens to the time released. Ronnie warned that a demonstration built in isolation can become a box-ticking exercise if it cannot be used more widely or if the intended users never adopt it.

The Human Decision Becomes More Demanding

Automating preparation can free staff time, but the benefit depends on where that time goes. Participants saw opportunities to devote more attention to client needs, exceptions and higher-value work. They also recognised a cultural challenge: a quicker task does not by itself show that the whole process improved or that a team has redirected its capacity usefully.

A licence and a login can be counted easily; the harder evidence is whether staff change how they prepare a review, coordinate across teams and use the time released. Participants described training existing employees, encouraging formal study and making domain knowledge central to recruitment and development. Familiarity with consumer AI tools creates interest, but it does not by itself teach a banker to challenge a source-of-wealth account or judge when a generated summary is incomplete.

Andrew identified a further consequence for reviewers: “You also need to equip them to have the ability to decide.” If AI assembles the file, the person receiving it must have the knowledge to test the evidence and make the final call. Training and domain expertise become more important as routine preparation recedes. The human-in-the-loop principle has to be supported by a workforce capable of exercising that judgement.

The client relationship is also changing. Advisers can prepare more quickly and draw on a fuller account of a client’s history, while clients may arrive with their own analysis and ask harder questions. Participants expected those conversations to become more substantive, not to disappear. Even younger clients who want less paperwork may still seek a person who understands their changing circumstances and can offer differentiated advice. The adviser’s role remains to interpret the evidence and build trust over time. In that sense, the measure of an AI programme extends beyond speed: it is whether the institution delivers better-informed service while keeping its decisions accountable.

Disclaimer: This article summarises a webinar discussion and reflects information available as at 30 September 2026. It is for general information only and is not tax, legal, financial, investment or other professional advice or a recommendation. Its contents may not apply to individual circumstances, products or jurisdictions. Hubbis accepts no responsibility or liability for actions taken or omitted in reliance on it. Readers should obtain independent advice from qualified professionals before making any decision.