

R O Y C

AI as the Operating Layer of Private Markets:

**From fragmented workflows
to scalable execution**

White Paper
July 2026

Written by Peter Bergenwald, Jonas Lundberg, Mathias Leijon

About ROYC

ROYC is the leading European B2B financial technology company that provides a complete private markets operating system, empowering private equity firms, banks, wealth managers, and multi-family offices to seamlessly access, distribute, and manage private investments at scale. As private markets expand, financial institutions require scalable, technology-driven solutions to manage complexity, optimising fund operations, and delivering exceptional client experiences. ROYC combines state-of-the-art private markets technology with tailored fund structuring and investment solutions. Its intuitive, scalable platform replaces manual processes with automation and real-time data access, transforming how private market investments are managed across the entire fund lifecycle.

Executive Summary

Private markets are entering a phase of sustained growth, driven by increasing participation from both institutional investors and private wealth. Regulatory developments, product innovation, and broader distribution are all contributing to this expansion, and for most market participants, the direction is clear: private markets are becoming a core part of long-term allocation strategies.

At the same time, the way private markets are operated has not evolved at the same pace. Across fund structuring, operations and distribution, workflows remain fragmented, highly manual and difficult to scale. As volumes increase, this gap becomes more visible. The main constraint is no longer access to capital or investment opportunities, but the ability to execute consistently and efficiently across the lifecycle.

Artificial intelligence is often positioned as one solution to this problem. However, in practice and so far, its impact has been limited, particularly in fund and client operations where the need is most acute. The reason is not a lack of technology, but a mismatch between how AI is typically applied and how private markets actually function. This is not a simple data problem; it is an execution environment where context, dependencies and precision are critical.

What is now starting to change is the way AI is being integrated. Rather than being used as a tool layered on top of existing systems, it is increasingly being embedded into the workflows themselves. This allows systems not only to support work, but to carry parts of it out, creating a more consistent and scalable operating model over time.

This paper explores why current operating models are no longer sufficient, why more traditional AI approaches fall short in private markets, and what changes when AI is applied directly to execution. It also outlines what this shift means for GPs, fund administrators and wealth managers, and how ROYC approaches this as part of its operating layer platform.

Table of Content

Private markets at an inflection point	5
The limits of the current technology stack	6
Why AI has not (yet) delivered	7
From assistance to execution	8
From data to operational understanding	9
Responsibility and control	10
Implications for operating models and economics	11
ROYC's approach	12
AI applied to platform workflows	14
Conclusion	16

1. Private markets at an inflection point

Private markets are undergoing a structural shift. What was previously a relatively closed ecosystem, dominated by institutional investors and a limited number of ultra-high-net-worth individuals, is becoming broader and more accessible.

This is not simply an expansion at the margins. It changes how capital is raised, structured and distributed, and introduces a new level of complexity across the value chain. As more participants enter the market and the number of transactions increases, the operational demands grow accordingly.

What is striking, however, is that the underlying operating model has remained largely unchanged. Onboarding processes are still heavily reliant on documentation and manual validation. Capital calls and distributions continue to be coordinated across multiple systems and counterparties, often requiring significant human intervention. Reporting remains slow, fragmented and difficult to standardise across different fund structures and jurisdictions.

As long as volumes were manageable, this model could hold. But as activity increases, the limitations become more apparent. Revenue opportunities scale with demand, but operational costs tend to follow in parallel. In practical terms, many firms still need to add headcount for every additional fund, client segment or distribution channel they introduce.

Over time, this creates a structural constraint. Growth becomes tied to operational capacity, and scaling the business becomes increasingly complex. The next phase of private markets will therefore not be defined by access or product innovation alone, but by the ability to operate at scale without a corresponding increase in complexity.

Exhibit 1

Private markets at an inflection point

THE MARKET	THE OPERATING MODEL
<p>More investors Wealth managers, new client segments and broader private wealth entering the market</p>	<p>Onboarding unchanged Still reliant on manual document review and validation</p>
<p>More transactions Growing volume across funds, geographies and distribution channels</p>	<p>Capital calls still manual Coordinated across multiple systems with significant human intervention</p>
<p>More complexity New fund structures, regulatory frameworks and distribution requirements</p>	<p>Reporting still fragmented Slow and difficult to standardise across structures and jurisdictions</p>

2. The limits of the current technology stack

Technology has played an important role in enabling the expansion of private markets. Over the past decade, platforms have improved connectivity between stakeholders, streamlined certain parts of the onboarding process, and enhanced the client experience through more intuitive interfaces and better access to information.

However, most of these systems were designed with a specific purpose in mind: to record, track and coordinate workflows rather than to execute them. They capture data, provide visibility and support users in navigating complex processes, but they do not fundamentally change how the work itself is performed.

This distinction becomes important when looking at how workflows are handled in practice. Even in environments that are considered highly digitised, manual intervention remains a core part of the process. Onboarding may begin through digital channels, but still requires manual review and validation across multiple documents. Capital call workflows are often tracked within systems, yet coordinated through email and spreadsheets. Reporting outputs may be generated automatically, but still need to be reconciled and interpreted before they can be distributed.

As a result, technology has improved transparency, but not necessarily efficiency at a structural level. It has made workflows easier to manage, but has not removed or simplified them to the extent required for true scalability. In many cases, the underlying complexity has simply been moved rather than eliminated.

3. Why AI has not (yet) delivered

Given these constraints, AI appears to offer a natural solution. The ability to process unstructured data, automate repetitive tasks and generate outputs at scale aligns well with many of the challenges faced in private markets.

Yet adoption has been uneven, and the impact has been more limited than expected, particularly in core operational areas. The reasons for this are often misunderstood.


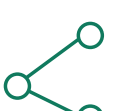

Private markets are not simply lagging behind in technology adoption. They are characterised by a level of operational complexity that generic AI solutions are not designed to handle effectively. Context plays a central role, and even small variations in terminology can have significant legal and financial implications. At the same time, workflows are deeply interconnected, with each step depending on the correct execution of the previous one. Finally, the tolerance for error is extremely low, given the regulatory and fiduciary responsibilities involved.

In this environment, generic AI solutions, typically trained on broad datasets and optimised for horizontal use cases, struggle to deliver reliable outcomes. They can produce outputs, but not always with the level of precision and contextual understanding required to execute processes end-to-end.

The limitation is therefore not intelligence in isolation, but the absence of a deeper understanding of how work is actually performed within the domain.

Exhibit 2

Why generic AI solutions fall short in private markets

 <p>Context determines meaning</p> <p>Small variations in terminology carry significant legal and financial consequences. Generic AI solutions are not designed to handle this level of precision.</p>	 <p>Workflows are deeply interconnected</p> <p>Each step depends on the correct execution of the previous one. An error at any point does not stay contained – it propagates across the entire lifecycle.</p>	 <p>The tolerance for error is extremely low</p> <p>Regulatory and fiduciary responsibilities leave no room for approximation. Every output must be consistent, traceable and precise – regardless of volume or complexity.</p>
--	---	---

4. From assistance to execution

The way AI is applied is now starting to shift. Early use cases have focused on assisting users, helping draft documents, summarise information or suggest next steps. While these applications can improve productivity, they leave the underlying workflow largely unchanged.

A more structural change occurs when AI is applied directly within the workflow itself. In this model, the system does not simply support the user, but participates in the process by interpreting inputs, applying rules and carrying out specific tasks. This can include reading and structuring documents, triggering subsequent steps in a process, or executing actions across systems.

Over time, this approach allows workflows to be handled in a more consistent and integrated way. Processes that previously required coordination across multiple teams and systems can be managed within a single flow, with fewer handovers and less dependency on manual intervention. Documents become inputs that drive actions, rather than static sources of information, and reporting can be generated and validated as part of the same process.

Importantly, this does not remove the need for human oversight. Rather, it changes where and how that oversight is applied, allowing people to focus on judgement and exception handling rather than routine coordination.

5. From data to operational understanding

A common response to the challenges outlined so far is to focus on improving data, making it cleaner, more structured, and easier to access. While this is clearly important, it does not in itself resolve the underlying issue.

Private markets do not suffer from a lack of data. In fact, most organisations already have access to large volumes of information across documents, systems, and counterparties. The difficulty lies in how that information is structured, interpreted and used within the context of real workflows.

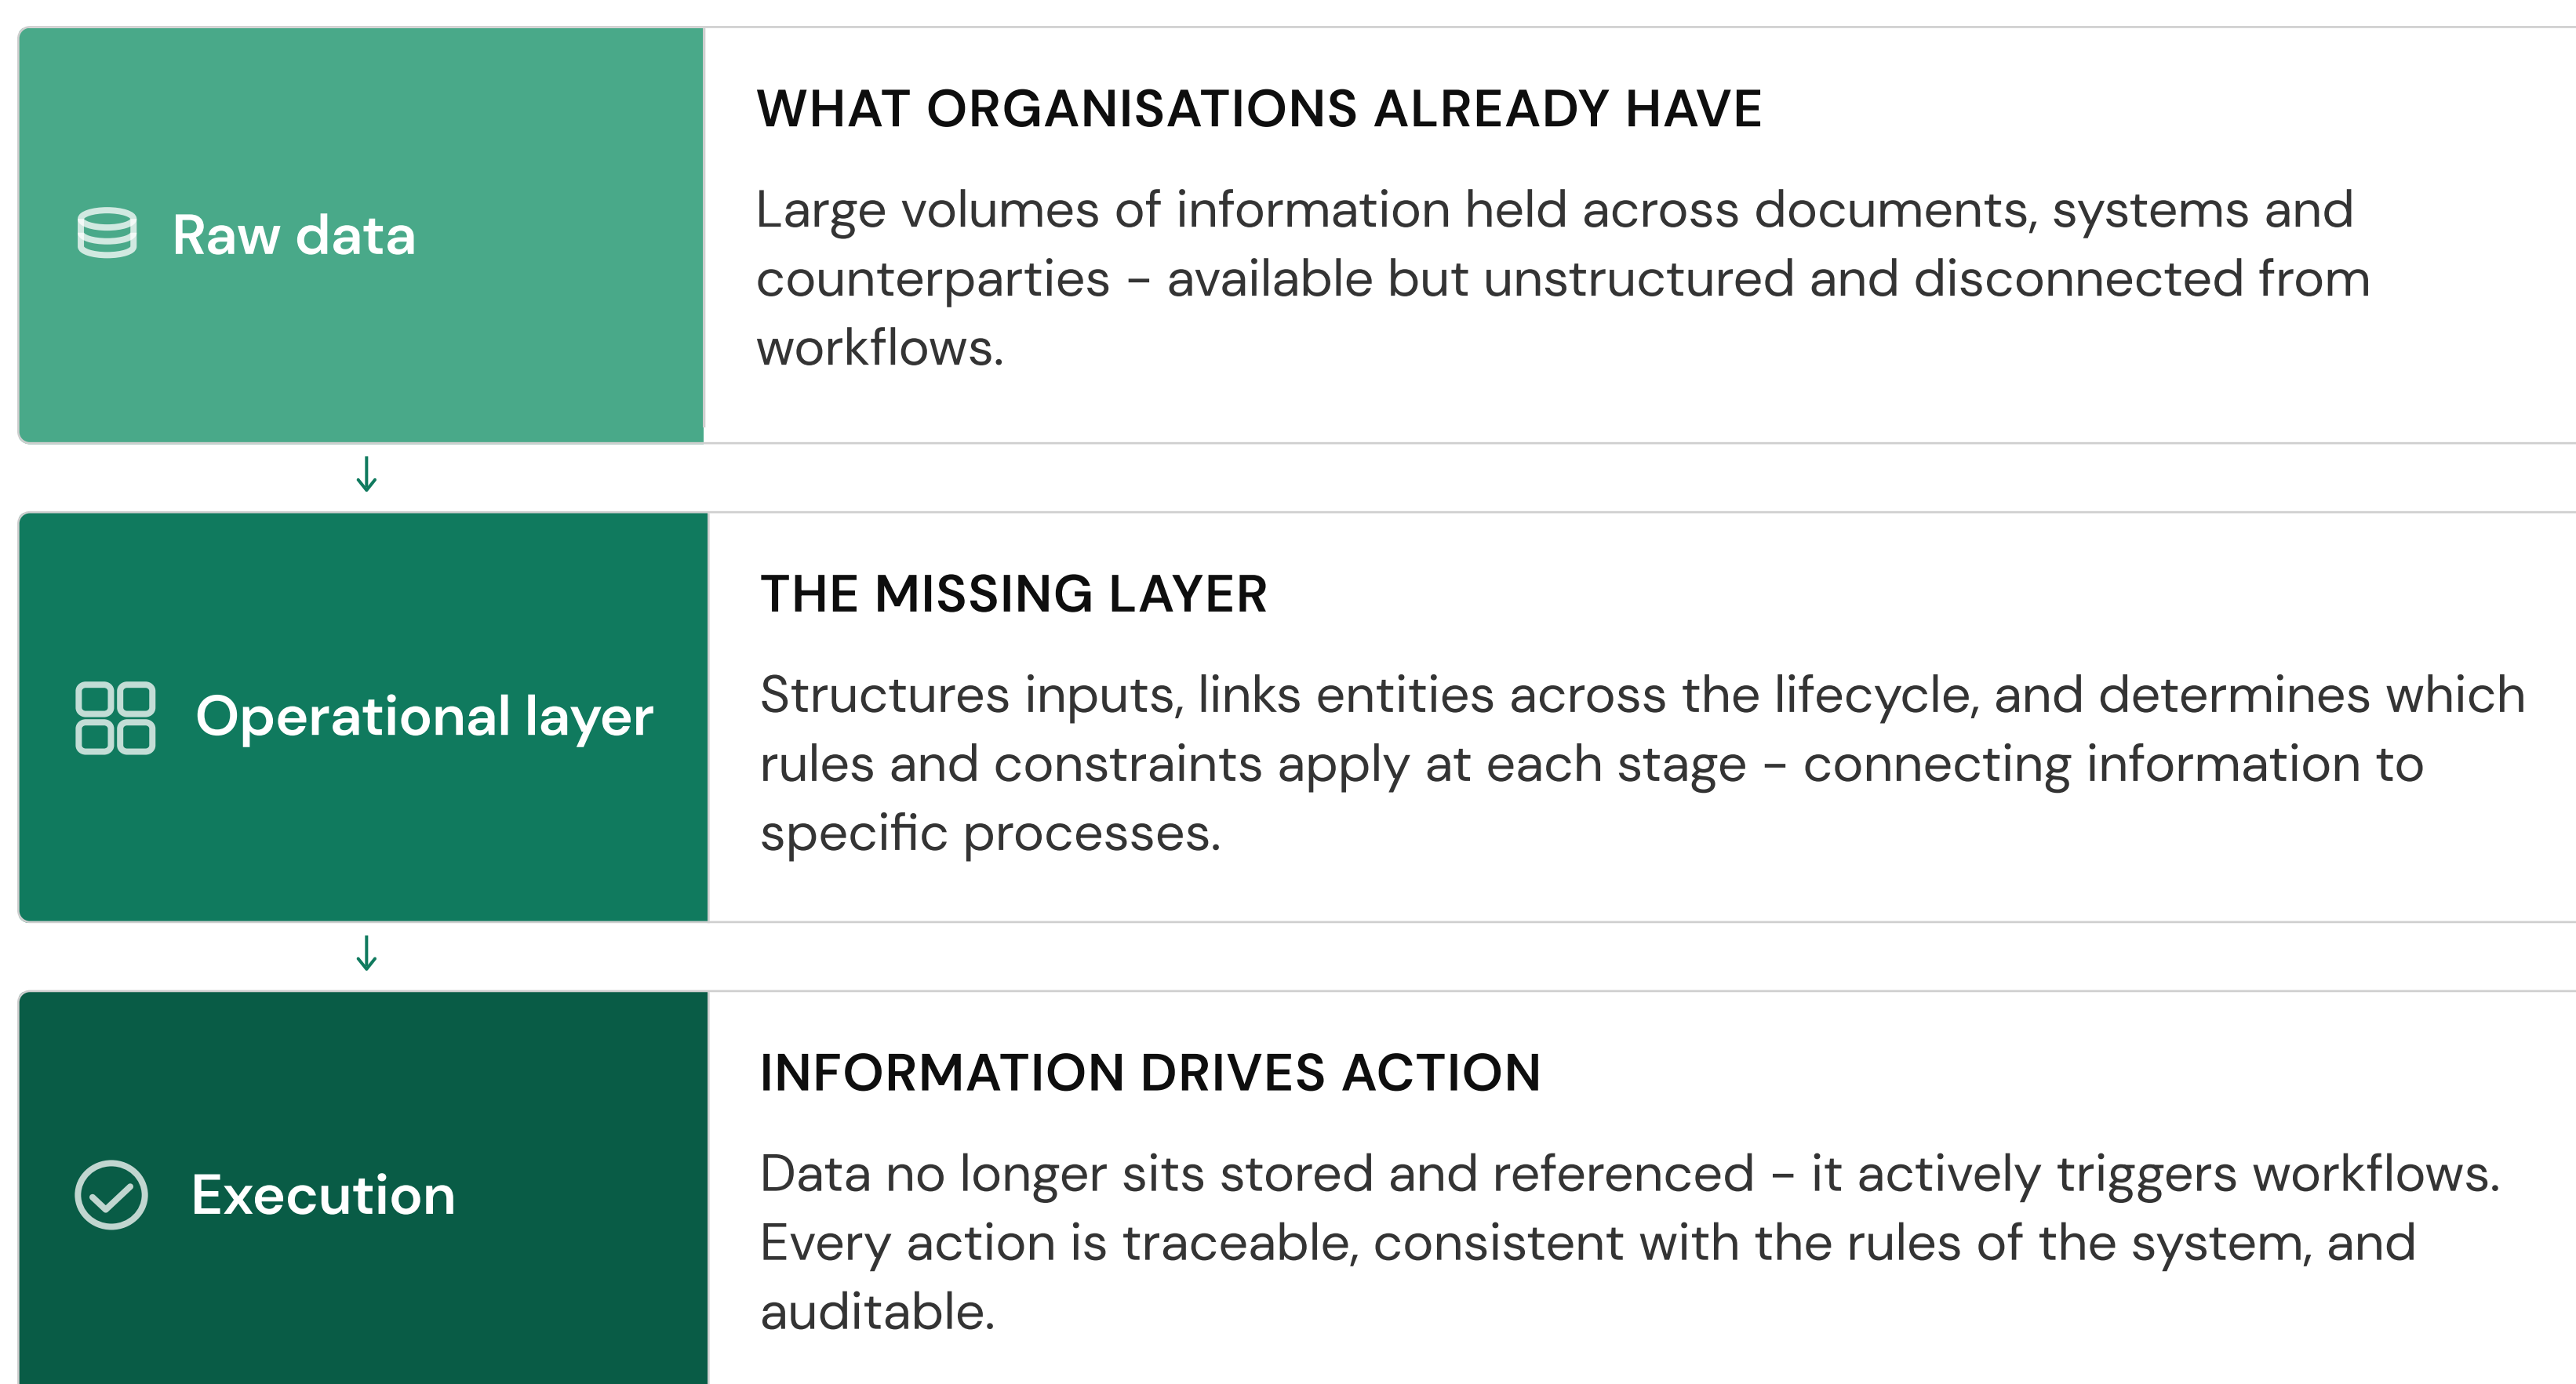
What is required is not simply better data, but a more consistent way of connecting data to an actual action. This means building systems that understand how information relates to specific processes, how it should be handled at different stages of the lifecycle, and what needs to happen next.

In practical terms, this introduces the need for an additional layer within the architecture, one that sits between raw data and execution. This layer is responsible for structuring inputs, linking entities across the lifecycle, and assessing what rules and constraints to apply. It allows information to move from being something that is stored and referenced to something that actively drives workflows.

The importance of this becomes particularly clear in a regulated environment. Every action must be traceable, and every output must be consistent with the underlying rules of the system. Without this level of operational understanding, automation quickly breaks down or introduces unacceptable risk.

Exhibit 3

From data to operational understanding



6. Responsibility and control

The introduction of AI into private markets inevitably raises questions around data security, governance, and control. These are not theoretical concerns, but practical considerations that need to be addressed from the outset.

In this context, AI cannot be treated as an experimental layer that sits alongside existing systems. It has to be integrated into the core infrastructure and subject to the same standards that apply to any other part of the operating model.





This includes clear separation of client data, strict control over how information is used, and full auditability of actions taken by the system. It also requires a defined role for human oversight, particularly in situations where decisions have financial, legal, or reputational implications.

Approached in this way, responsibility does not slow down adoption but rather enables it. Firms that establish the right governance frameworks early are able to move more confidently, extend automation further across workflows, and build trust with both clients and regulators.

Over time, this becomes a differentiating factor. The ability to combine efficiency with control and oversight is not just a technical advantage, but a commercial one.

Exhibit 4

Responsibility and control: governance requirements for AI in private markets

REQUIREMENT	WHY IT MATTERS
 Client data separation	Each client's data must be held in strict isolation. Mixing information across counterparties introduces legal exposure and undermines investor trust.
 Control over information use	Strict governance must define how data is accessed, processed and shared within the system. AI cannot operate without clearly bounded permissions.
 Full auditability	Every action taken by the system must be traceable. This is a prerequisite in any regulated environment.
 Defined human oversight	Decisions with financial, legal or reputational implications must involve human sign-off. AI can drive workflows, but accountability must remain with people.

7. Implications for operating models and economics

Today, most of the work is still done by people, with systems helping along the way. When AI is applied properly, parts of that work are handled by the system itself, which means firms can handle more without adding more people.

Exhibit 5

How AI changes the economics across roles



Wealth Managers

The ability to onboard clients and manage positions without a corresponding increase in manual work opens up new possibilities for distribution. Private markets can be offered more broadly, without significantly increasing the cost-to-serve.



General Partners

Reducing the operational burden associated with fund launches and ongoing management shortens time-to-market and allows for more flexibility in structuring and distribution strategies.



Fund Administrators

Automating document-heavy and coordination-intensive processes improves throughput and consistency, making it possible to handle higher volumes without a proportional increase in resources.

Across all of these roles, a similar pattern emerges. Growth is no longer tightly linked to headcount. Instead, it becomes possible to scale activity within a more stable operational framework.

This has a direct effect on the unit economics of the business. Margins improve not only through cost reduction, but through the ability to expand without introducing additional layers of complexity.

8. ROYC's approach

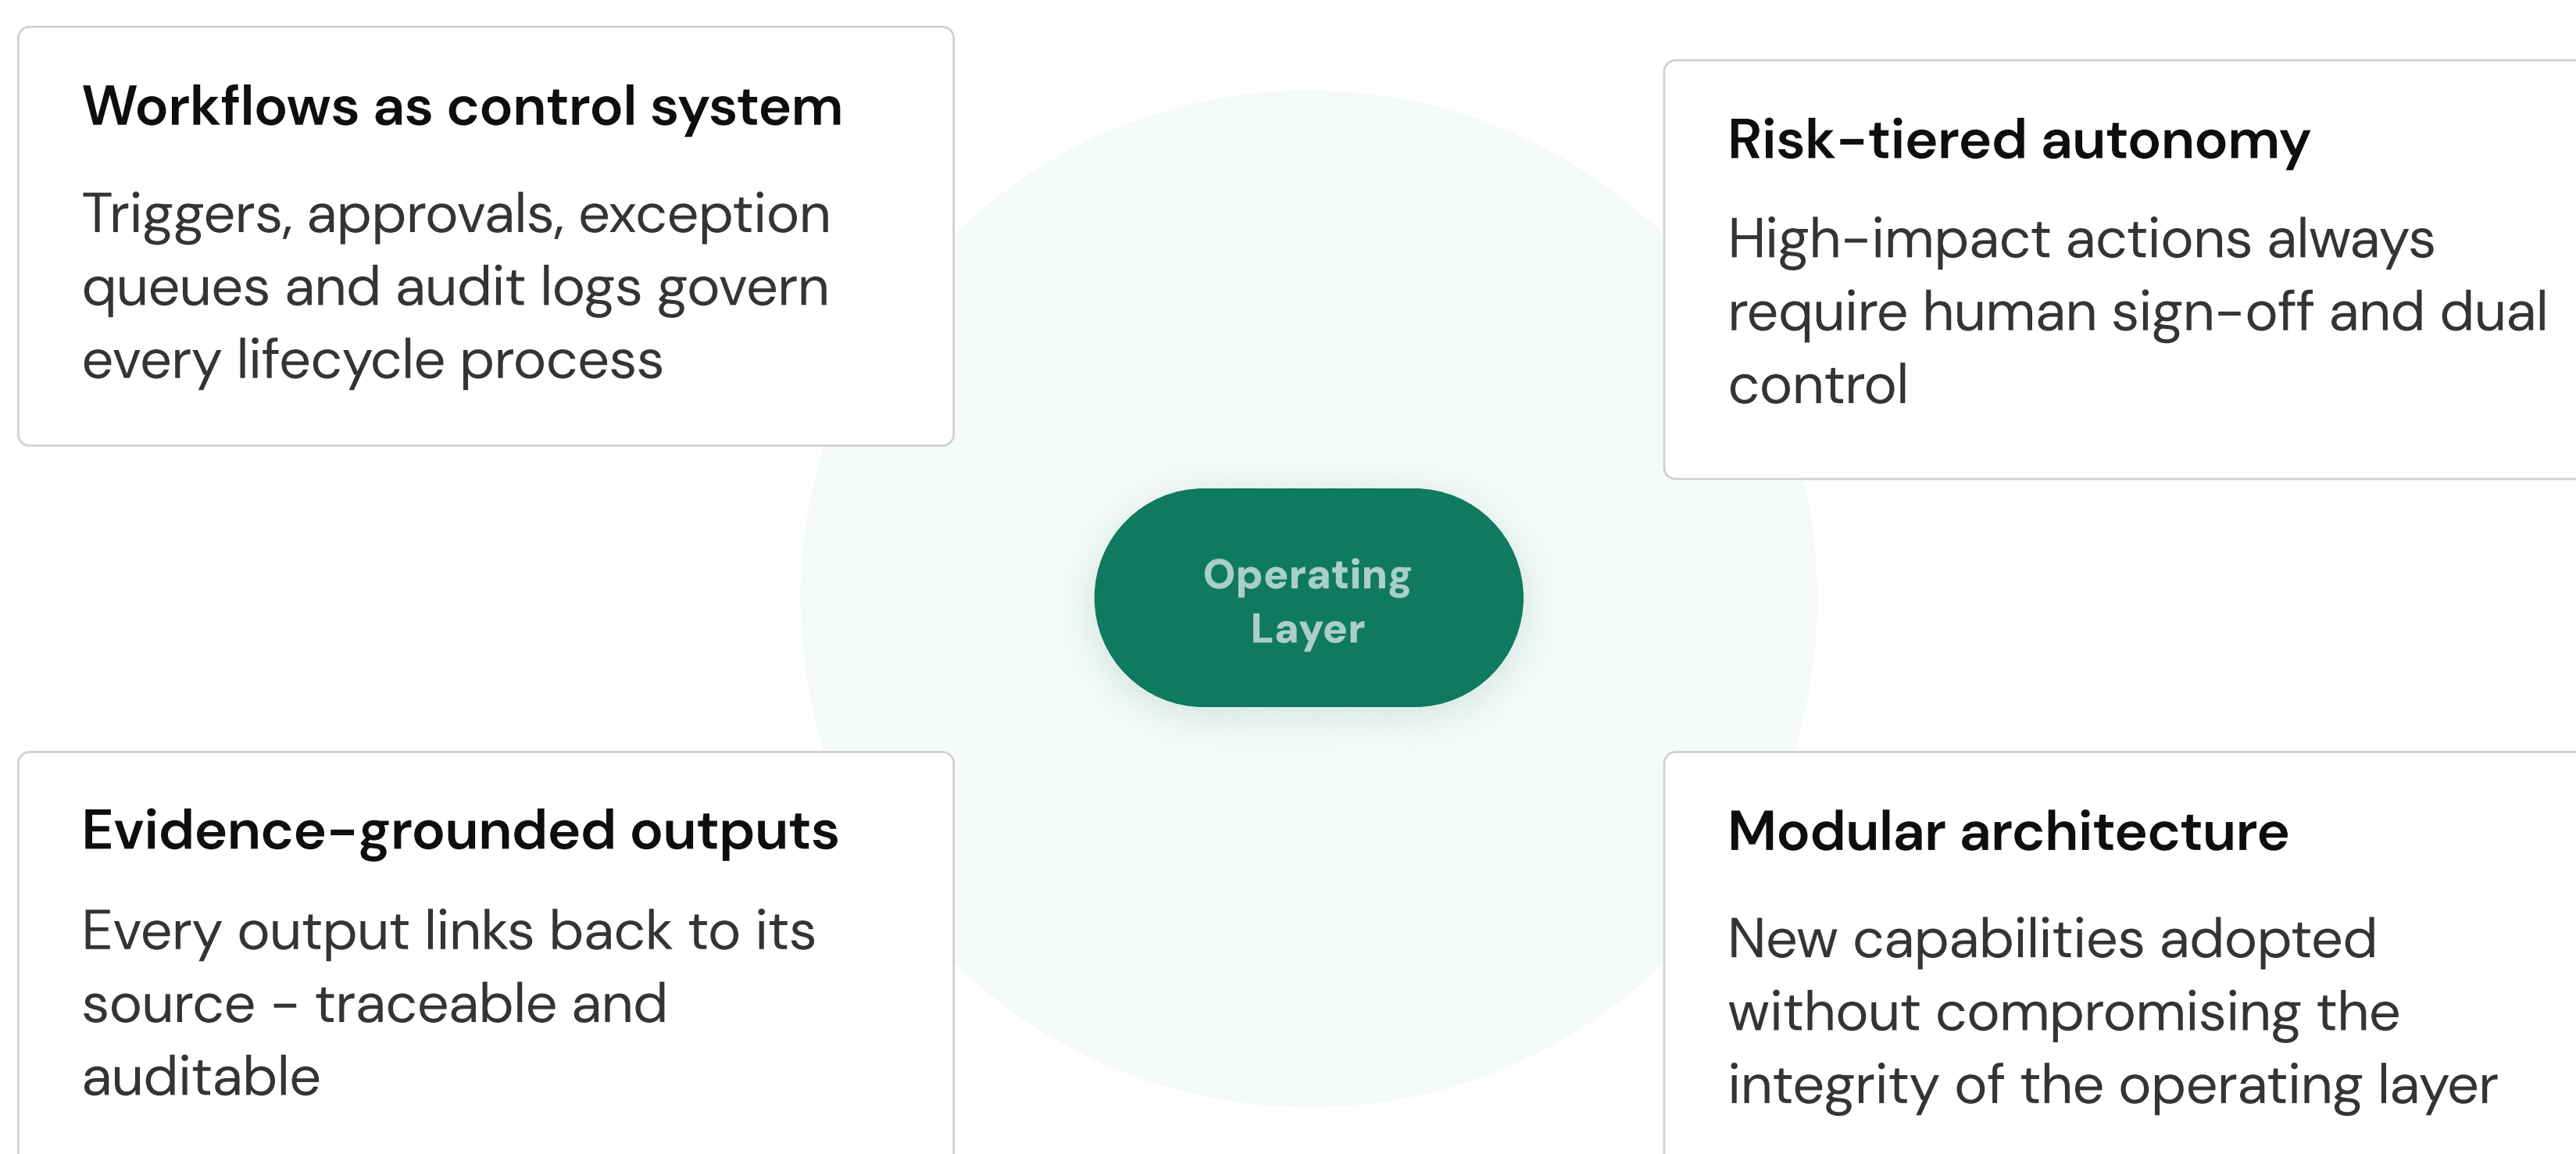
The previous sections describe a set of conditions that a private markets operating model now needs to meet. Workflows must be connected rather than coordinated. AI must be embedded in execution rather than layered on top. Data must be linked to action, and control must be built in from the start rather than added afterwards. ROYC has been designed around exactly these requirements.

Our platform brings fund structuring, operations and distribution together within a single infrastructure. This is a deliberate architectural choice. By operating on one shared backbone, the system can hold a consistent view of investors, funds, commitments, positions, documents, cash flow events and approvals, and AI can be applied directly within this environment rather than against fragmented sources. Information moves consistently between participants, and actions are triggered within the system rather than coordinated externally.

Within this structure, four design choices shape how AI is applied on the platform.

Exhibit 6

ROYC's approach: four design choices



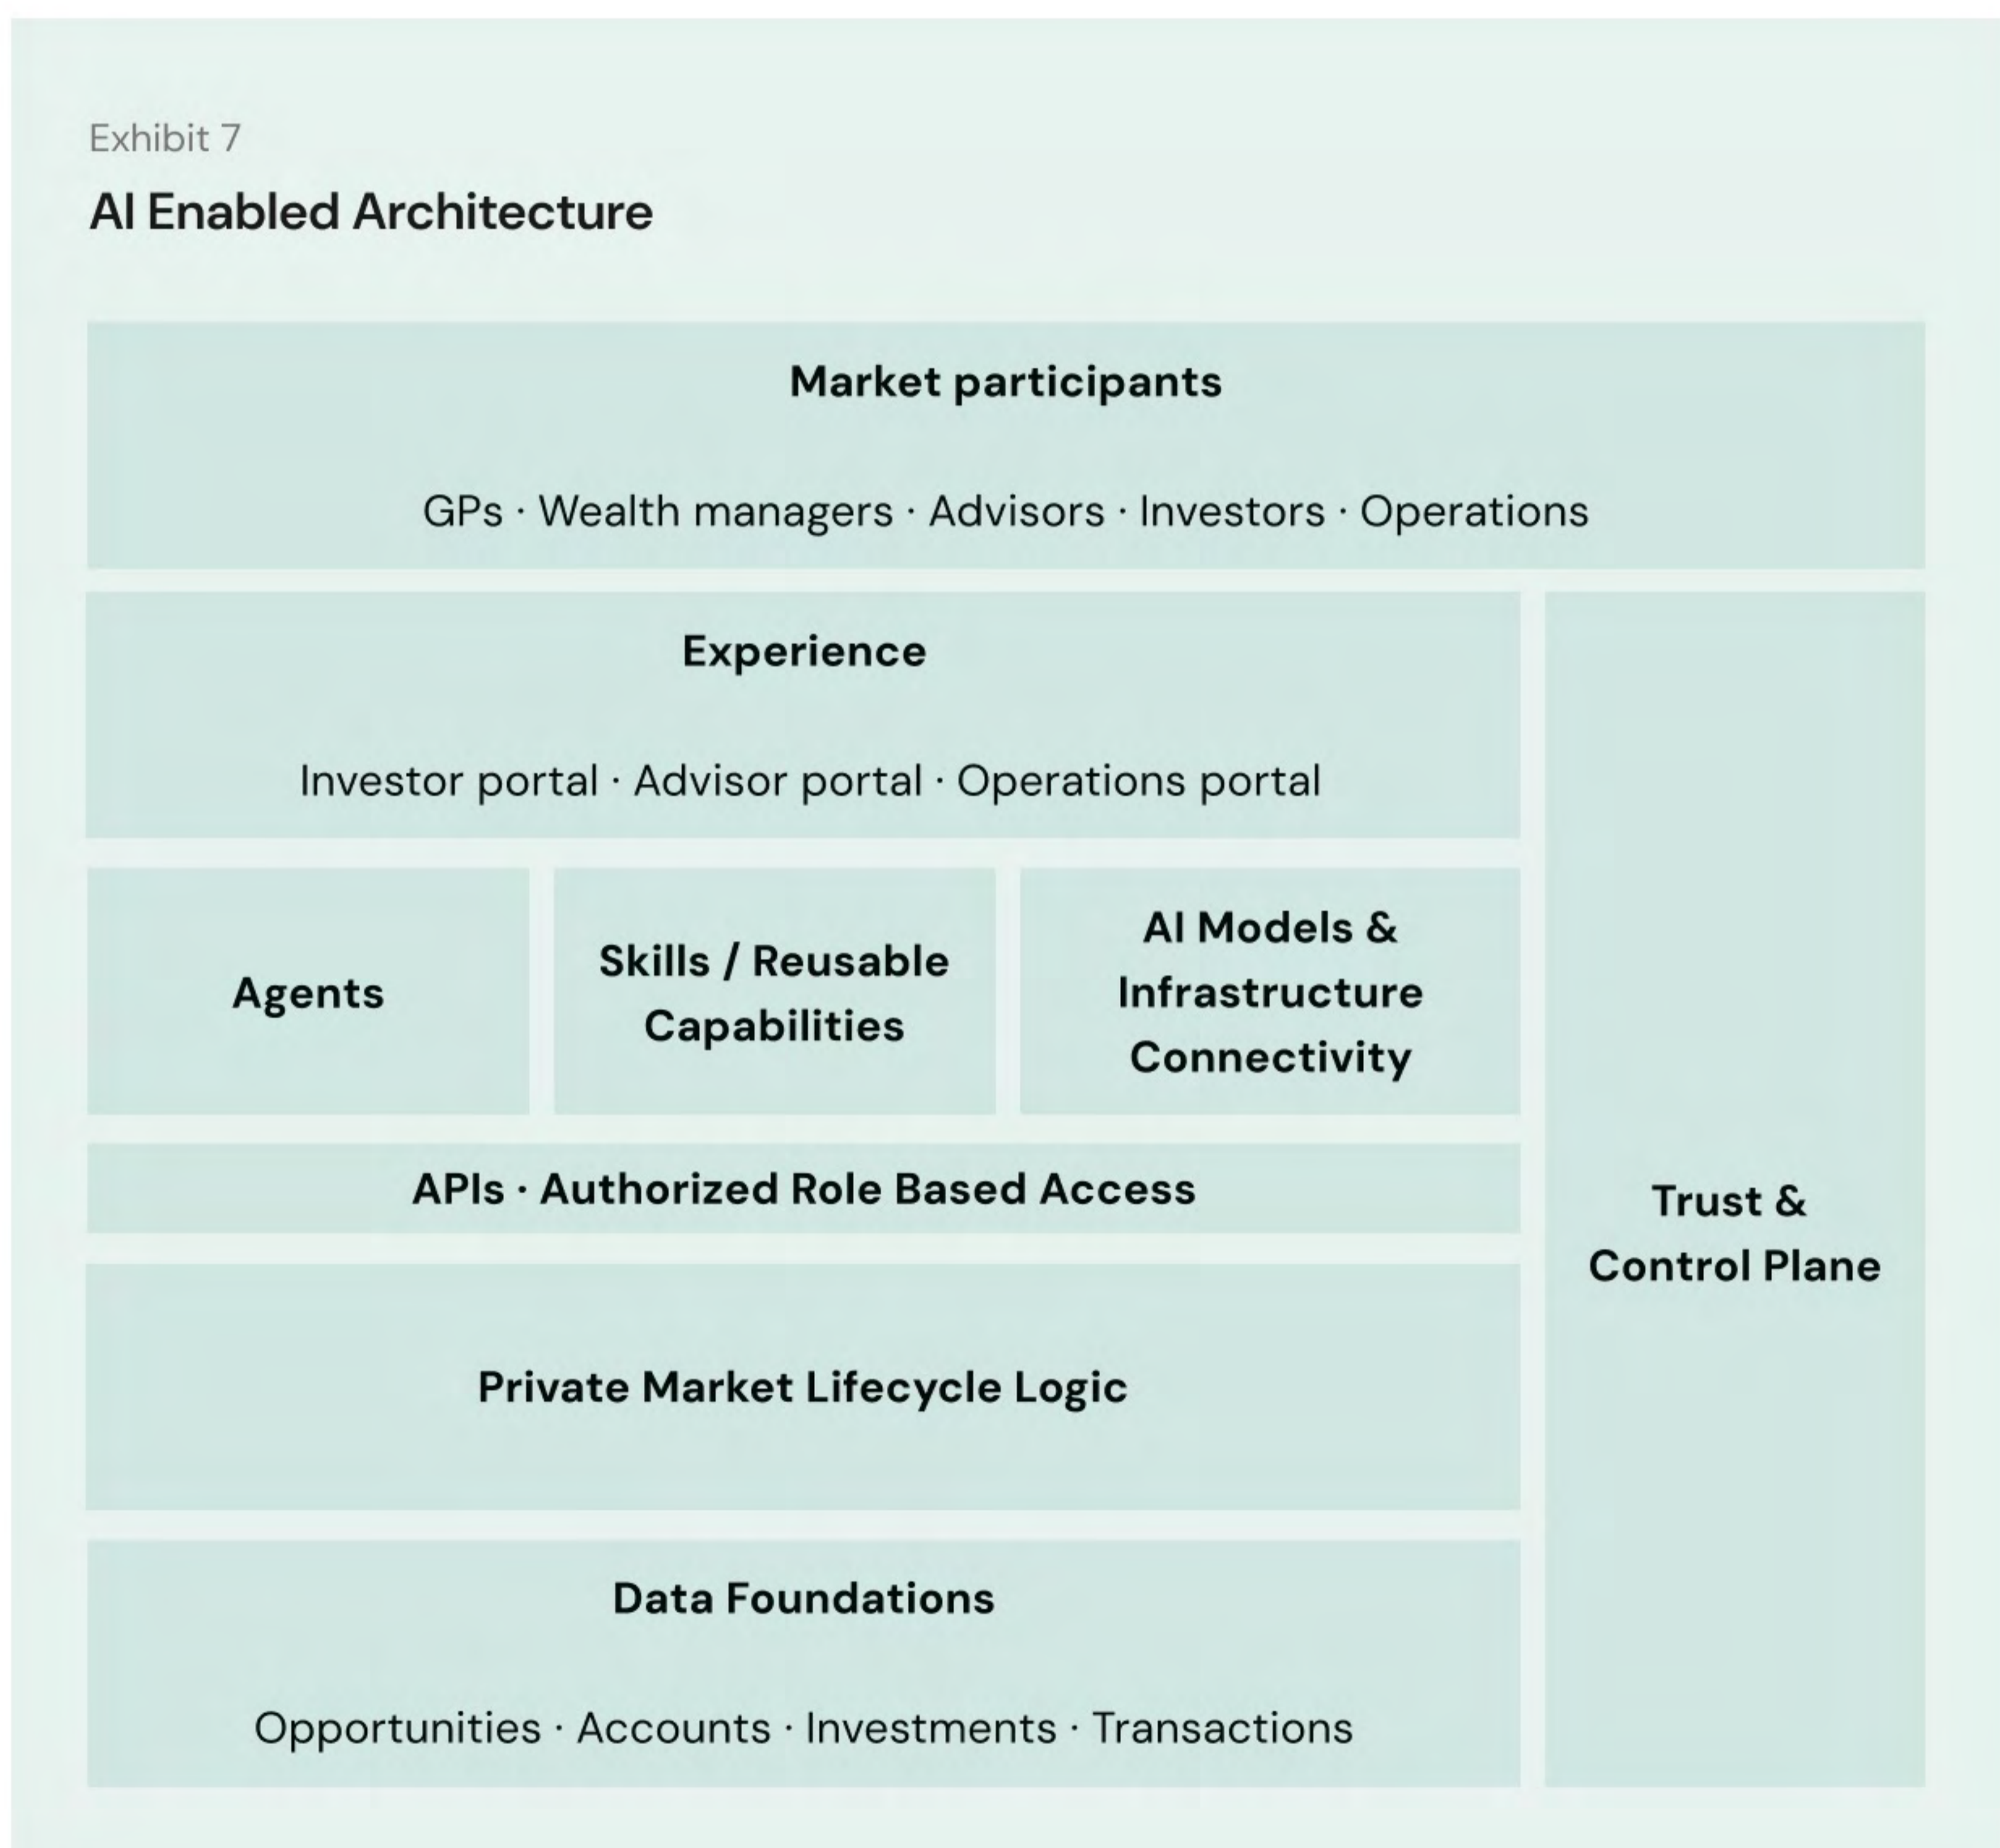
The first is that workflows function as the control system. Each lifecycle process, including onboarding, subscription, capital calls, distributions and reporting, is modelled with defined triggers, approvals, exception queues and audit logs. AI accelerates the work inside these workflows, but the workflow itself governs sequence, controls and accountability. This is what allows automation to extend across the lifecycle without introducing additional operational risk.

The second is that AI operates with risk tiered autonomy. Not all AI actions warrant the same level of control. On the ROYC platform, read only retrieval, drafting, recommendations and bounded operational tasks are clearly separated from high impact actions such as investor approvals, capital movements or regulatory conclusions, which always require human sign off and, where relevant, dual control. AI can drive flows but are not allowed to e.g. do calculations. Automation scales with confidence, while accountability stays with people.

The third is that outputs are evidence grounded. Every extracted field, recommendation and generated answer links back to its source, whether that is a clause in an LPA, a page in a capital call notice or a record in the system. Outputs are trusted because they are traceable and reviewed. Combined with tenant isolation, role-based entitlements and immutable audit logs, this provides the level of auditability that regulated environments require.

The fourth is that the architecture is modular and replaceable. Models, skills and integrations sit behind controlled interfaces, which means the platform can adopt new capabilities over time without compromising the integrity of the operating layer or the controls applied to it.

Taken together, these choices produce an environment in which AI carries real operational load, including interpreting documents, validating data, reconciling cash flows, drafting communications and routing exceptions, while humans retain responsibility for judgement and the system retains a complete record of what happened, why, and on whose authority. The aim, as set out earlier in this paper, is not to introduce new complexity, but to reduce it. By aligning how workflows are designed and how they are executed, it becomes possible to improve both efficiency and consistency over time, and to scale activity within a more stable operational framework.



9. AI applied to platform workflows

Some clear examples of how AI is applied to platform workflows:


Embedded report import augmented with AI

Custom reports are imported directly, with embedded AI presenting the extracted data alongside confidence levels and highlighted anomalies rather than being re-keyed and only reconciled by hand.

This speeds up fund operations and lets the user catch potential errors before the data enters the platform, with discrepancies surfaced at the point of import rather than later downstream. The AI acts as an additional layer of review within a human-in-the-loop workflow, while the user decides what is accepted.

Exhibit 8

How AI reviews data before it enters the platform

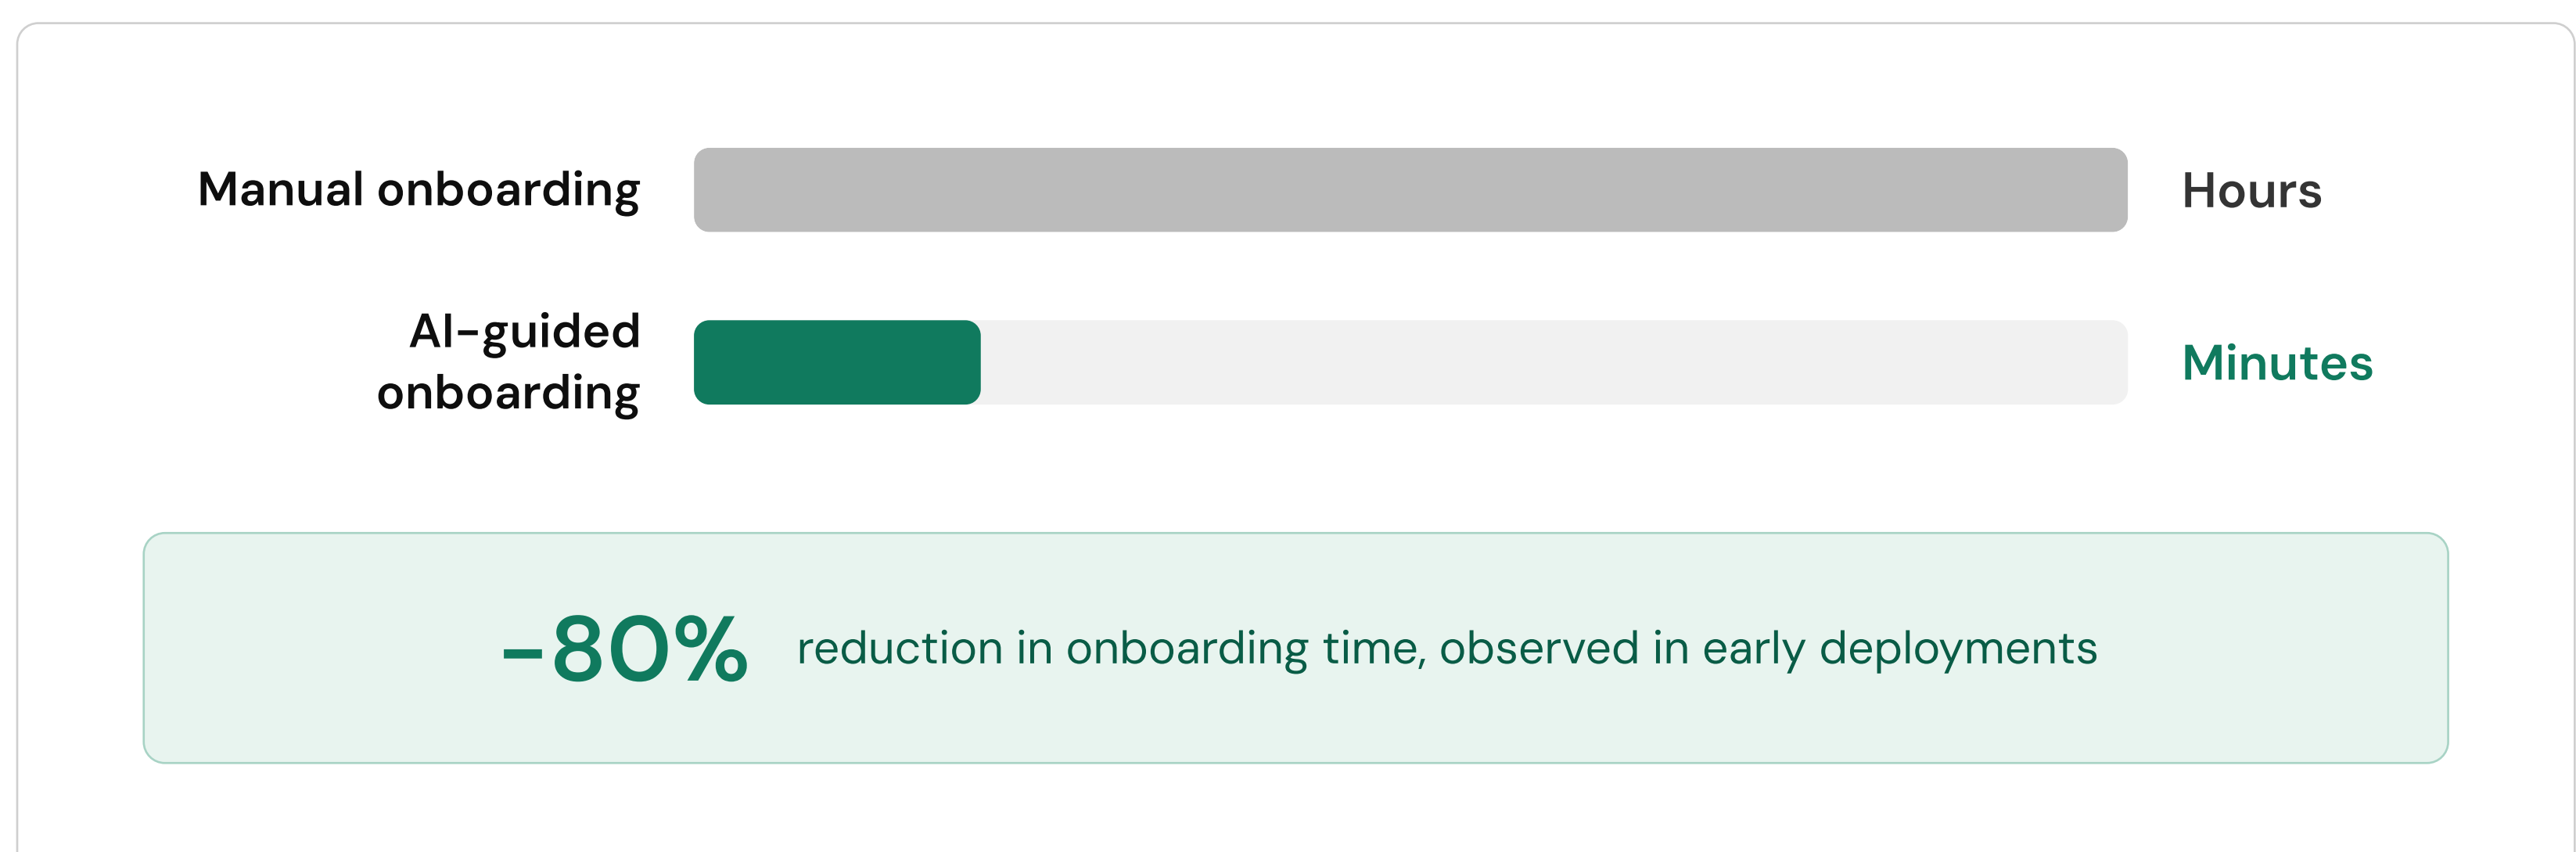


AI Guided Onboarding based on documents

On the ROYC platform, investor details can be read directly from uploaded supporting documents rather than keyed by hand. This reduces manual work, and in early deployments this has cut onboarding time by more than 80%, from several hours to minutes, while reducing the risk of incorrect or inconsistent data entering the system. On the GP side, discrepancies and missing documents are surfaced during investor approval, so the fund administrator reviews a structured, evidence-grounded picture rather than reconstructing it manually.

Exhibit 9

Onboarding time, before and after AI



Natural-language data access through the ROYC Operations MCP

GPs, fund administrators and investor services teams can query platform data such as subscription details and customer due diligence records directly in natural language through an in-platform agent, or another connected MCP-compatible AI tool. For example, this can be used to create ad hoc reports, removing the lag that ad hoc queries have traditionally introduced, and compressing what could take days or weeks into a single conversation. The judgement stays with the user, but they work from live, structured answers.

10. Conclusion

There is no shortage of demand in private markets. More investors are entering, both institutional and private wealth, and that trend is expected to continue.

The constraint limiting scalability lies elsewhere. It sits in the way the industry operates, in the workflows, the systems, and the coordination required to manage increasing levels of activity.

Improving access and distribution has been an important first step. The next phase is about improving execution.

AI will play a very important and transformational role in this transition, but only where it is applied in a way that reflects the realities of the domain. This means focusing less on isolated tools and more on how workflows are structured and carried out end-to-end.

Firms that are able to make this shift will not simply become more efficient. They will be able to scale in a way that is more predictable, more controlled, and less dependent on manual intervention.

About the Authors



Peter Bergenwald, CTO
peter.bergenwald@roycgroup.com



Jonas Lundberg, Chief Architect & AI Lead
jonas.lundberg@roycgroup.com



Mathias Leijon, Co-Founder & President
mathias.leijon@roycgroup.com

Disclaimer

This white paper is provided for general information purposes only and does not constitute, and should not be relied upon as, legal, tax, investment, or other professional advice within the meaning of applicable European Union or national laws. The content reflects solely the opinions and interpretations of the authors and the contributing firms at the time of publication. It is not intended to address the circumstances of any particular individual or entity, and may not reflect the most current legal or regulatory developments in the European Union, the European Economic Area, or specific Member States.

Readers should not act upon the information contained herein without obtaining appropriate professional advice from qualified legal or other advisers licensed in the relevant jurisdiction. While every effort has been made to ensure the accuracy of the information at the date of publication, no representation or warranty, express or implied, is given as to its accuracy or completeness. To the fullest extent permitted by applicable law, the authors and contributing firms disclaim any and all liability for any loss or damage suffered as a result of reliance on this document or its contents.

R O Y C

roycgroup.com



@ROYC

AI as the Operating Layer of Private Markets:
From fragmented workflows to scalable execution