

CHECKLIST

Is your fund-ops team ready for AI?

Ten questions on IR pipelines, custodian reconciliation and document workflows. Where AI saves partner hours, and where it only adds risk.

For COOs, CFOs and heads of operations at trust companies, family offices and wealth managers.

Most AI projects in regulated firms fail on the same thing: someone automated a process that needed a human in it.

These are the questions I ask first, before anything gets built. They are not a test of how modern your team is. They show where AI would save partner hours, and where it would only add risk.

Answer each one honestly. A “not yet” is useful. It tells you what to fix before you spend money on tools. No call required.

How to use it. Score each question: 2 if your answer looks like “ready”, 1 if it is partly there, 0 if it looks like “not yet”. The scoring guide is at the end.

SECTION A

Investor reporting and IR pipelines

WHERE AI SAVES PARTNER HOURS

Drafting reports and letters. Summarising and routing investor requests.

WHERE IT ONLY ADDS RISK

Sending anything without a person's approval. Numbers with no traceable source.

01 Can you trace every figure in an investor report back to its source?

WHY IT MATTERS

An investor report is only as good as its weakest number. If a figure is copied by hand from a spreadsheet that was itself copied from a portal, AI will repeat the error faster, and with more confidence.

READY

Each figure has one named source system. Someone can show the path from source to report without opening five files.

NOT YET

Figures are assembled by hand each quarter. The path lives in one person's memory. Different reports carry different versions of the same number.

YOUR SCORE

2

1

0

02 Does a named person approve every investor communication before it goes out?

WHY IT MATTERS

Drafting is where AI saves partner hours. Sending is where it adds risk. A model that drafts a quarterly letter is useful. A model that sends one is a liability.

READY

Every outbound investor message ends in "submit for approval". The approver is a person, and the approval is recorded.

NOT YET

Approval is informal, a "looks fine" in a chat. Some messages go out without anyone signing for them.

YOUR SCORE

2

1

0

03 Do investor requests arrive in one place, with an owner and a status?

WHY IT MATTERS

Requests from investors and their advisers are where hours disappear quietly. AI can sort, summarise and route them well. It cannot fix a process where nobody knows which requests exist.

READY

Requests land in one queue or shared inbox. Each has an owner and a status. You can see what is open today.

NOT YET

Requests sit in personal inboxes. Follow-ups depend on memory. Nobody can say how many are open.

YOUR SCORE

2

1

0

SECTION B

Custodian reconciliation

WHERE AI SAVES PARTNER HOURS

Matching routine items. Explaining breaks.
Preparing the exceptions list for a person.

WHERE IT ONLY ADDS RISK

Resolving breaks or posting adjustments without a person deciding.

04 Does custodian data reach you in a consistent, machine-readable form?

WHY IT MATTERS

Reconciliation automation stands or falls on its inputs. Consistent files make matching routine work for a machine. Statements that change layout every quarter make every run a new problem.

READY

You receive structured files or feeds from each custodian, on a known schedule, in a format that rarely changes.

NOT YET

Statements arrive as PDFs or emails in different layouts. Someone re-keys them before any matching starts.

YOUR SCORE

2

1

0

05 Are your matching rules and tolerances written down?

WHY IT MATTERS

If the rules for what counts as a match, and what difference is acceptable, live in one person's head, you cannot automate them and you cannot audit them. Writing them down is worth doing even if nothing is ever automated.

READY

The rules and tolerances are documented, agreed by finance and operations, and reviewed when something changes.

NOT YET

The rules are whatever the person doing the reconciliation has always done. When they are away, the work waits.

YOUR SCORE

2

1

0

06 When a break appears, is it clear who resolves it, and by when?

WHY IT MATTERS

AI is good at finding and explaining breaks. Resolving them is a judgement, often with a custodian or counterparty on the other end. Automation without an owner for exceptions only produces a longer list of exceptions.

READY

Each type of break has an owner and an expected resolution time. Ageing breaks are visible and escalated.

NOT YET

Breaks surface at month-end and are cleared under pressure. Old breaks are carried forward without anyone deciding to.

YOUR SCORE

2

1

0

SECTION C

Document workflows

WHERE AI SAVES PARTNER HOURS

Classifying, naming and filing documents.
Pulling out key fields such as entity, date and document type.

WHERE IT ONLY ADDS RISK

Silent misfiling, when nobody has decided what happens if the system is unsure.

07 Can you name your main document types, and where each one should end up?

WHY IT MATTERS

Classifying and filing documents is one of the clearest places AI saves time. But a model can only sort documents into categories you have defined. If the team disagrees on where a document belongs, so will the system.

READY

A short, agreed list of document types (investor updates, statements, KYC documents, legal correspondence and so on), each with a destination and an owner.

NOT YET

Folder structures differ by person or by entity. The same document is filed in two places, or in none.

YOUR SCORE

08 Have you decided what happens when the system is not sure?

WHY IT MATTERS

No classifier is right every time. What makes document automation safe is a confidence threshold: above it, the system files; below it, a person reviews. Without that rule, errors are silent.

READY

There is an agreed threshold, a named reviewer for low-confidence items, and a way to see what was filed automatically.

NOT YET

The plan is to let the AI handle it and check if someone complains.

YOUR SCORE

SECTION D

Before anything is built

WHERE AI SAVES PARTNER HOURS

Nothing directly. This is what makes the rest safe.

WHERE IT ONLY ADDS RISK

Skipping it.

09 Do you know which data a model may see, and where it may be processed?

WHY IT MATTERS

Client confidentiality and data residency are where AI most often adds risk in a regulated firm. KYC files, beneficial-owner details and personal data need rules before any tool touches them.

READY

Data is classified. Sensitive fields are masked before a model sees them. Access is read-only by default. You know where your data may be processed, and your vendors' terms match.

NOT YET

People paste client documents into public AI tools to save time. Nobody has checked where that data goes.

YOUR SCORE

2

1

0

10 Have you written down what stays manual, and why?

WHY IT MATTERS

This is the first question I ask. Supervision, client confidentiality, data residency and the audit trail are all good reasons to keep a step with a person. That list is the first deliverable, not an afterthought.

READY

A short list of steps that stay manual, each with a reason. Every automated step leaves an audit trail showing what happened, who approved it, and when.

NOT YET

The list of things to automate is long. The list of things to keep manual is empty.

YOUR SCORE

2

1

0

SCORING

Scoring guide

Add up your scores. The maximum is 20.

16–20	Ready for a focused start. Pick one workflow where you scored 2 on every related question, and start there. One workflow, one owner, one measure agreed before anything is built.
9–15	Partly ready. Some workflows can be automated now. Others need foundations first. Fix the questions you scored 0 before you spend money on tools.
0–8	Not yet. AI would mostly add risk today. The useful work is writing down sources, rules, owners and what stays manual. It costs less than a failed pilot, and it is worth doing anyway.

Whatever your total, a 0 on question 2, 9 or 10 is a stop sign. Fix that first.

Twenty minutes is enough to know.

If your answers raised a question you would like a second view on, book a 20-minute call. One call about your team, your stack, and the one decision that is stuck.

If there is something to do, we confirm the scope of the first process on the call. If there isn't, you walk away with a written memo.

Start with one process.

Book a 20-minute call at alarifa.com · eg@alarifa.com

I'm Evgeny Goncharov, founder of Alarifa, in Madrid. I decide what to automate in financial-services operations, build it inside your operations, and stay accountable for the outcome. I hold a fractional CTO seat with a regulated European trust and corporate services provider, and I am CTO for a global investment firm.

Alarifa provides technology and operations advisory and does not offer investment, tax or legal advice.