

Nine questions to ask, and one policy to adopt

AI in pension governance

A reference document for trustees, scheme secretaries, and in-house pensions teams.

INTRODUCTION

AI is already in your scheme

AI is now in routine use across the pensions ecosystem. A trustee, adviser, or scheme secretary may already be using public or enterprise AI tools to summarise, draft, or analyse scheme material. Advisers use AI in correspondence triage, first drafts, and scenario analysis. Administrators use machine learning for fraud detection and exception flagging. Member-facing platforms deploy AI in retirement modellers, chatbots, and personalised guidance. Some of this is procured and supervised by the board. Much of it is not.

The Pensions Regulator's AI plan, published in May 2026, is direct on this point. Trustees remain accountable for decisions and outcomes even where activities are delegated to providers or advisers. TPR expects trustees to understand where and how AI is being used by or on behalf of the scheme, to establish governance and accountability for that use, to assure themselves that the same standards apply across their administrators, service providers, and advisers, and to test, monitor, and manage the risks on a continuing basis.

This paper is for the people who carry that duty: trustees, scheme secretaries, and in-house pensions teams. It is not for the vendors.

Responsible AI in a pension scheme is not primarily about the model. It is about whether the board can govern the use case, understand the data, test the output, challenge the answer, evidence the decision, control the supply chain, and reconstruct the record later.

The nine questions in this paper are designed to operationalise that view. They are an initial governance screen, not a substitute for procurement, due diligence, or legal review. They are short enough to ask out loud in a board meeting. Each is designed to produce an answer that is either useful or revealing.

No platform architecture removes the need for trustee judgement, challenge, adviser oversight, and governance discipline. AI amplifies governance quality. It does not replace it. The questions in this paper test the technology. The board still has to test itself.

THE NINE QUESTIONS

An overview

Q01	Is this built for fiduciary work, or is it general-purpose AI in a new wrapper?	DOMAIN EXPERTISE
Q02	Does this match how a trustee board actually operates, including how we work with our advisers?	FIT FOR TRUSTEE BOARDS
Q03	Can the trustees themselves use this, or will we be relying on the scheme secretary to interpret its output?	USABILITY
Q04	Will the output hold up under regulatory, legal, and audit scrutiny, not today but in five years?	PERFORMANCE
Q05	Where does the answer come from, and can we check it?	ACCURACY
Q06	Will this platform keep pace with the regulator, the law, and the profession?	INNOVATION
Q07	Is the platform built to the standard of any regulated business, and is the evidence to support that on the table?	TRUST
Q08	Will this strengthen the scheme's memory across changes of trustees, advisers, and sponsor teams?	ENDURANCE
Q09	Will this connect to the rest of our scheme's stack, and can we get our record back out if we ever leave?	CONNECTIVITY

How to use this paper. Read it once. Then keep the checklist on the last page to hand. Every time AI is pitched, proposed, or already in use in connection with the scheme, work the nine questions. The board's accountability does not change because the AI is convenient, free, or familiar.

Q01

DOMAIN EXPERTISE

Is this built for fiduciary work, or is it general-purpose AI in a new wrapper?

AI cannot replace fiduciary judgement. The question is whether the tool in front of the board has been configured to support that judgement under the specific constraints of trustee work, or whether it is a general-purpose tool being asked to operate in a setting it was not designed for. Modern general-purpose models often understand pensions terminology well. Vocabulary is not the test.

The governance question is what surrounds the model: the scheme-specific records it can reach, the data boundaries it respects, the workflow it sits in, the controls on what it is permitted to do, the citations behind its answers, and the audit trail behind its decisions. Tools used for fiduciary work need either domain-specific design or demonstrably equivalent controls.

WHAT A GOOD ANSWER LOOKS LIKE

- **Built by people who have done the job.** The product team includes trustees, scheme secretaries, and pensions professionals, not only engineers.
- **Domain depth shows in small things:** late papers handled cleanly, actions carry forward across meetings, conflicts of interest are a structural feature not a checkbox.
- **Confidence comes from the combination:** data quality, retrieval design, controls, logging, and human review working together.

WARNING SIGNS

- ! The vendor cannot name a pensions professional who shaped the product.
- ! The output reads as articulate but generic, and could equally apply to a charity board or an NHS trust.
- ! Demonstrations rely on document summarisation, not on tasks specific to fiduciary work.

Q02

FIT FOR HOW TRUSTEE BOARDS WORK

Does this match how a trustee board actually operates, including how we work with our advisers?

Pension governance is collaborative by design. Trustees deliberate together. They take advice from investment consultants, scheme actuaries, covenant advisers, legal counsel, and the in-house pensions team. They sign off in committee, often by sub-committee first. The work is asynchronous, distributed, and shared between people who do not all sit in the same firm, the same country, or the same week.

Many AI deployments, including capable enterprise tools, still fail to reflect the multi-party, adviser-led nature of trustee decision-making. A platform that treats the board as a single user, in one organisation, working on one document, will quietly create more friction than it removes.

WHAT TPR EXPECTS

TPR's AI plan (May 2026) expects trustees to assure themselves at all stages that their administrators, service providers and advisers have similarly robust governance arrangements in place for AI use. The board's duty extends across the supply chain, not only to AI it procured directly.

WHAT A GOOD ANSWER LOOKS LIKE

- **Reflects the meeting cycle and the adviser relationship.** Built around quarterly cycle, agenda-driven, paper-led, sub-committees, advisers.
- **Selective collaboration with advisers and the sponsor.** Controlled access for EBCs, professional trustee firms, legal counsel, and sub-committees, with an audit trail of who saw what and when.
- **Context-rich answers grounded in the scheme's own documents** at the point of retrieval, not on a model trained on them.

WARNING SIGNS

- ! Single-user productivity tool with no native concept of trustees, advisers, sponsor-side teams, or committees.
- ! No audit trail of who accessed what, and when.
- ! Collaboration costs or access models discourage proper adviser participation.

- ! Data leaves the scheme's environment in ways the board cannot see or control.

Q03

USABILITY

Can the trustees themselves use this, or will we be relying on the scheme secretary to interpret its output?

Trustees and the scheme secretariat are different kinds of user. Trustees are intermittent and high-stakes: a handful of board meetings a year, deciding on questions with significant consequences. The secretariat is operational and continuous: running the cycle, preparing the pack, holding the record together between meetings. A good platform serves both.

A platform the trustees themselves do not feel confident using is a platform whose outputs the trustees cannot meaningfully challenge. Confidence to challenge the machine is the floor of fiduciary use of AI.

WHAT A GOOD ANSWER LOOKS LIKE

- **Designed for both audiences.** Trustees can read a paper on a phone, search the record in three taps, and ask a plain-English question that returns a cited answer on the same screen.
- **Onboarding fit for each role.** Trustees get a short induction scheduled around board meetings. The secretariat gets the operational depth they need.
- **Embedded in the existing workflow.** Removes friction from the existing cycle rather than asking the board to migrate to a new productivity philosophy.

WARNING SIGNS

- ! The platform is designed for one audience only, and the other is treated as an afterthought.
- ! The training plan does not distinguish between intermittent high-stakes users and daily operational users.
- ! Reading and searching the record on a phone is awkward, or limited to a desktop experience.

Q04

PERFORMANCE

Will the output hold up under regulatory, legal, and audit scrutiny, not today but in five years?

In fiduciary work, the failure that matters is not the obvious mistake. It is the answer that is almost right: the summary that misses one of five key points, the action logged against the wrong meeting, the risk flagged but never resurfaced. These survive the meeting and only surface when something else has gone wrong.

WHAT A GOOD ANSWER LOOKS LIKE

1. Consistent output across users, quarters, and advisers.
2. Performance demonstrated on representative trustee work, not on a generic benchmark.
3. Named references the board can speak to directly.
4. Honest under partial information. A serious platform tells the board when it does not know, with the same confidence it uses when it does.

WARNING SIGNS

1. Performance claims rest on generic AI benchmarks.
2. Output quality varies between users or between quarters.
3. The platform never says 'the record does not support this.'

Q05

ACCURACY

Where does the answer come from, and can we check it?

Accuracy in fiduciary AI has a specific meaning. It is not whether the output sounds right. It is whether every substantive claim can be traced to a primary source, and whether the trustee can verify it without leaving the platform. A generic AI tool can summarise a meeting pack in seconds. It will also, with the same fluency, invent a clause that does not exist in the deed, attribute a position to an adviser who did not take it, and confidently misquote the regulator.

Citation discipline is the operational expression of that test. So is refusal: a platform that will not answer questions outside its grounded record is doing the trustee a service, not failing them.

WHAT TPR EXPECTS

TPR flags inaccuracy of AI-generated information, and the boundary between advice and guidance, as specific risks. TPR expects boards to ensure scheme data is of high quality, and to understand how AI models use and process that data.

WHAT A GOOD ANSWER LOOKS LIKE

- **Citations at the sentence level.** Not 'here are some documents you should read.' Citations on every substantive claim, linked back to the specific paragraph in the underlying document.
- **Refusal as a feature.** The platform tells the trustee when the record does not contain the answer. It does not improvise from general knowledge.
- **Reviewability built into the workflow.** The auditor, three years later, can reconstruct what the board was told and on what basis.

WARNING SIGNS

- ! The platform produces fluent answers with no source visible to the user.
- ! Citations are at document level ('see the SIP') rather than at paragraph level.
- ! The platform never refuses to answer.

Q06

INNOVATION

Will this platform keep pace with the regulator, the law, and the profession?

The UK pensions environment is moving quickly. The General Code reset trustee obligations. Mansion House and productive finance are reshaping investment expectations. TPR has committed to publishing AI guidance for pension schemes in 2026.

A platform built for this market needs to evolve at the cadence the market evolves at. Annual release cycles are usually too slow. Updates measured in months rather than years are typically what good regulatory responsiveness looks like.

WHAT A GOOD ANSWER LOOKS LIKE

1. The platform should know more about the scheme in year three than in year one. If the value proposition is the same every year, the board has bought a static tool.
2. The platform should resolve new challenges imposed by changes in guidance or legislation (ESOG, ORA etc) or market trends (schemes buying in and winding-up).

WARNING SIGNS

1. The roadmap reflects general AI trends, not pensions-specific change.
2. The platform sits on a single underlying model with no path to substitute it.

Q07

TRUST

Is the platform built to the standard of any regulated business, and is the evidence to support that on the table?

Any AI sitting inside the trust chain between members, trustees, advisers, sponsors, and regulators has to meet the standard the rest of the chain is held to. In fiduciary AI, trust is not a marketing claim. It is a structural property of the platform: visible in how it handles data, who can see what, what gets logged, and how those properties have been independently tested.

The same security, permissions, audit, and data-handling discipline applies whether the platform serves legal teams, charity trustees, NHS boards, or pension trustees. A platform that is robust enough for a regulated business is robust enough for a pension scheme.

WHAT TPR EXPECTS

TPR's principles for AI in pensions include fairness, accountability, transparency, and explainability. Boards are also expected to comply with data protection legislation, including on automated decision-making, and to understand how AI models use and process scheme data.

WHAT A GOOD ANSWER LOOKS LIKE

- **Governance-grade security as a baseline.** ISO 27001 is the floor, not the ceiling. UK and EU data residency, encryption in transit and at rest, penetration testing on a regular cycle.
- **Data handling specified, not asserted.** The vendor specifies subprocessors, hosting arrangements, model call destinations, telemetry, retention periods, and whether customer data is ever used to train models.
- **Permissions, audit, and accountability built in.** Role-based access at the level the work actually needs. A full audit trail granular enough to support a regulator enquiry.

WARNING SIGNS

- ! Vague answers on data residency, training-data usage, or model behaviour.
- ! No role-based access, or audit trails that do not support real scrutiny.
- ! The vendor implies a track record it cannot substantiate when pressed for named references.

- ! The vendor cannot describe its position on member-facing transparency.

Q08

ENDURANCE

Will this strengthen the scheme's memory across changes of trustees, advisers, and sponsor teams?

Pension schemes are long-duration institutions. Their decisions outlast the trustees who made them, the advisers who recommended them, the scheme secretaries who recorded them. The board may have reached a sound conclusion five years ago, but if the rationale sits in an old meeting pack or the memory of a retired chair, the practical value of that judgement is lost.

AI used in pension governance should not merely answer today's question. It should strengthen the scheme's ability to remember why earlier questions were answered as they were. The platform's value should compound over time.

WHAT A GOOD ANSWER LOOKS LIKE

- **Continuity across people and roles.** Actions, risks, conflicts, decisions, and obligations survive changes in trustees, advisers, chairs, scheme secretaries, and sponsor-side personnel.
- **Decision history, not just document storage.** The platform connects papers to meetings, meetings to minutes, minutes to actions, and decisions to later reviews.
- **Endgame and wind-up durability.** As schemes approach buy-in, buyout, or wind-up, historic decisions become more important, not less.

WARNING SIGNS

- ! The platform stores documents but does not link them to decisions or actions.
- ! A new trustee cannot use the platform to understand the scheme's position without an extended briefing.
- ! Historic context relies on individual recall rather than on the record.
- ! The vendor cannot describe how the platform supports an endgame, wind-up, or post-wind-up scenario.

Q09

CONNECTIVITY

Will this connect to the rest of our scheme's stack, and can we get our record back out if we ever leave?

A pension scheme does not run on a single platform. Documents sit in one system. The actuary's data lives in another. The investment consultant publishes through their own portal. The administrator runs its own stack. Any AI used in connection with the scheme will be more useful, or more limiting, depending on how it relates to all of that.

Connectivity is also the question that decides what happens at exit. The scheme's record belongs to the scheme. A platform that makes it easy to leave is a platform that earns the right to be stayed with.

WHAT A GOOD ANSWER LOOKS LIKE

- **Integrates with the systems the scheme already uses:** dashboards, actuarial valuation tools, third-party governance applications (e.g. ORA), and the systems advisers and administrators use.
- **Open to the wider AI ecosystem.** The platform can be reached by other tools through documented APIs or open standards such as MCP.
- **Portable record at exit.** A full export with metadata, audit logs, permissions history, and decision and action links preserved.

WARNING SIGNS

- ! The platform expects to be the only place the scheme's record lives, and integration with other systems is described as a future feature.
- ! The vendor cannot describe what a clean export of the scheme's full record would look like, or hedges on whether it includes decision history and audit trail.
- ! APIs are limited to data the vendor chooses to expose, or are only available on the highest pricing tier.
- ! The architecture assumes the board's other tools will adapt to it, not the other way round.

CLOSING

The standard you set

These nine questions are deliberately demanding. They are also, we think, the questions the UK pensions market will eventually settle on: through TPR's expectations, through professional standards in the trustee community, through the slower verdict of which schemes ended up adequately served by the technology they adopted, and which did not.

Boards adopting AI now are setting the standard for what the next decade looks like. The tools they choose to trust will, in turn, set the standard for what the next generation of governance technology is built to do.

We wrote this paper because we believe that standard should be set deliberately, not by default.

The right tool raises the floor of the conversation between those who advise and those who must decide. Without that, the rest is secondary.

The nine questions tell a board what to look for. The harder question — what 'governed' has to mean in practice — is the subject of the next paper in this series.

A STARTING POINT FOR YOUR POLICY

Ten questions to shape your AI policy

A guide. Adapt to the scheme. Adopt what serves the board.

The nine questions in the first half of this paper are designed to assess the tools the board is being shown. They are necessary, but not sufficient. A board can choose well-controlled AI and still govern it badly. The historic governance failures in UK pensions have rarely been technological. They have been failures of ownership, escalation, documentation, and challenge.

The second half of this paper is a starting set of questions a board might use to shape, review, or stress-test its AI governance policy. A question on this list that the scheme's policy framework cannot yet answer is, at minimum, a useful conversation to have at the next board meeting.

P01 OWNERSHIP

Who owns AI governance internally, what is their authority, and how do they report to the board?

P02 APPROVED USE

What gets approved by the board, and what controls apply to each approved use case?

P03 LEGAL REVIEW

What AI decisions require legal review before adoption, and which can be approved within existing delegations?

P04 PROHIBITED USE

What is prohibited, and how is the boundary made unambiguous to trustees, advisers, and administrators?

P05 TESTING CADENCE

What testing cadence exists for material AI use, what evidence is produced, and what triggers an early re-test?

P06 SIGN-OFF ON MATERIAL CHANGE

Who signs off material AI changes (new use cases, change of model or supplier, expanded access), and at what threshold do they come back to the board?

P07 INCIDENT AND ESCALATION

What is the incident and escalation process, what counts as a reportable AI incident, and what triggers board, sponsor, legal, or regulator notification?

P08 TRUSTEE CHALLENGE

What is the trustee challenge process for AI-derived outputs, and how is the challenge documented and resolved?

P09 AI AUTOMATIONS

Which governance tasks are partly or fully automated by AI, who has approved each automation, and how does the board maintain oversight of what is being done on its behalf?

P10 ASSURANCE MODEL

What is the assurance model for AI use, what is the independent component, and how does the board satisfy itself that the controls are real?

How to use these questions: Treat them as discussion starters. Some boards will find their existing risk and compliance framework already covers most of them. Others will identify gaps worth closing. The questions are a way for a board to satisfy itself that its AI governance is keeping pace with the use of AI in connection with the scheme.

BOARD CHECKLIST

Nine questions to ask, and one policy to adopt

Work these every time AI is pitched, proposed, or already in use — by a vendor, by an adviser, by an administrator, or by a colleague at the keyboard.

Q01**DOMAIN EXPERTISE**

Is this built for fiduciary work, or is it general-purpose AI in a new wrapper?

! WARNING

Cannot name a pensions professional who shaped the product.

Q02**FIT FOR BOARDS**

Does this match how a trustee board actually operates, including how we work with our advisers?

! WARNING

No audit trail of who accessed what.

Q03**USABILITY**

Can the trustees themselves use this, or will we be relying on the scheme secretary?

! WARNING

Designed for one audience only.

Q04**PERFORMANCE**

Will the output hold up under regulatory, legal, and audit scrutiny, not today but in five years?

! WARNING

Performance claims rest on generic AI benchmarks.

Q05**ACCURACY**

Where does the answer come from, and can we check it?

! WARNING

Fluent answers with no source visible.

Q06**INNOVATION**

Will this platform keep pace with the regulator, the law, and the profession?

! WARNING

Roadmap reflects general AI trends, not pensions change.

Q07**TRUST**

Is the platform built to the standard of any regulated business, and is the evidence on the table?

! WARNING

No role-based access or real audit trails.

Q08**ENDURANCE**

Will this strengthen the scheme's memory across changes of trustees, advisers, and sponsor teams?

! WARNING

Stores documents but does not link them to decisions.

Q09**CONNECTIVITY**

Will this connect to the rest of our scheme's stack, and can we get our record back out if we ever leave?

! WARNING

Cannot describe a clean export of the scheme's record.



The governance platform for UK pension schemes.

1,000+ boards. £200bn+ AUM. Three-time UK Pensions Awards winner.

knowa.co