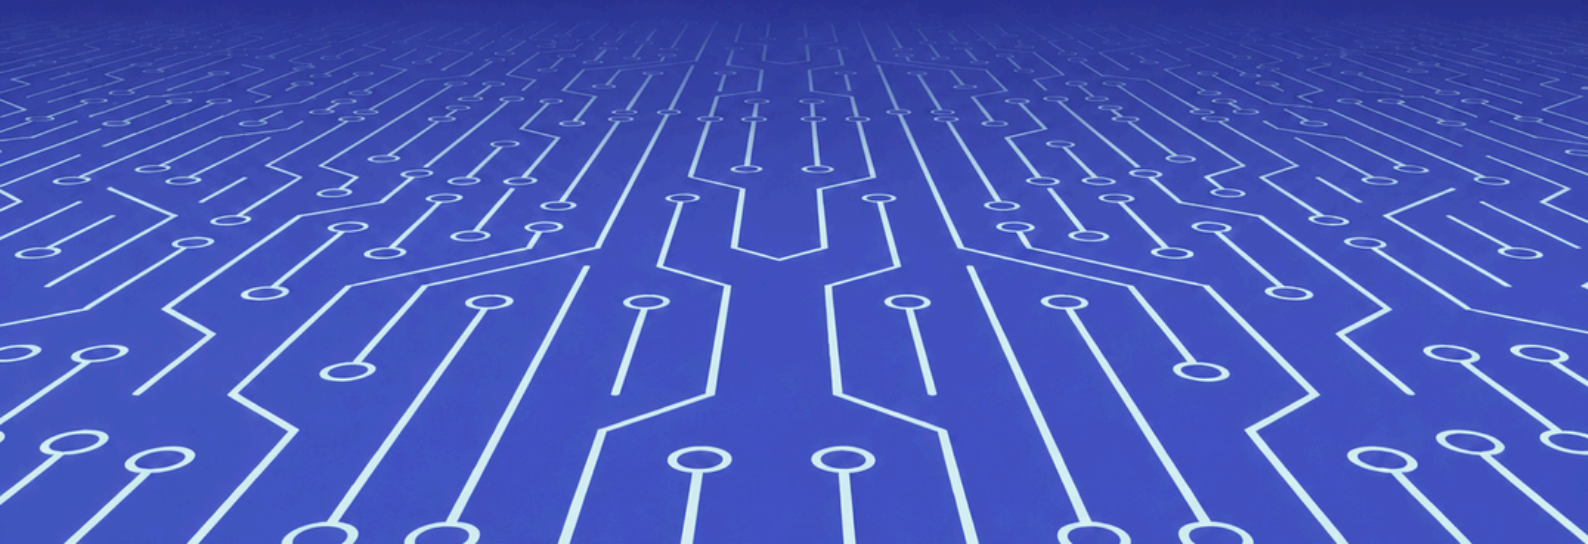
Woodhurst

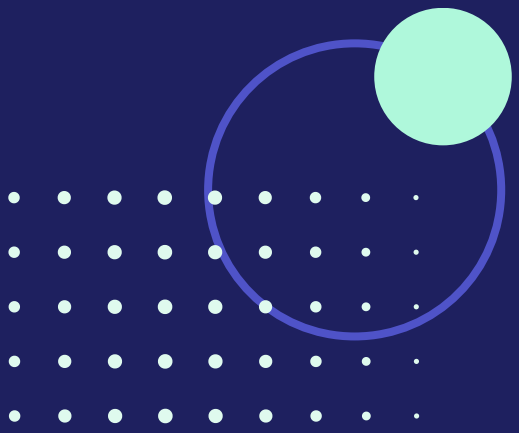
BEYOND EXPERIMENTATION:

BUILDING **AI GOVERNANCE** FOR THE **MUTUAL SECTOR**

A practical whitepaper helping building societies adopt AI safely, responsibly and at scale.



Woodhurst





Contents

Foreword by Steve Hogg 04

Purpose of this whitepaper 06

Governance without paralysis 07

Building AI on trusted data 09

Federated innovation, central accountability 11

From adoption to accountable use 13

Measuring value and managing the hidden AI estate 15

AI governance framework 17

Conclusion 18

About the authors 19

About Woodhurst 20

Foreword

Steve Hogg
Chief Operating Officer at
Enra Specialist Finance



> A strong AI-deployment framework and capability is really the only part of the investment with a shelf life.

The promise

ChatGPT launched in November 2022. 4 years on, trillions of dollars of CAPEX have been poured into AI with the share prices of AI-related firms reaching stratospheric highs. The evangelists proclaim a productivity revolution and rising redundancies. We are told we may also expect 100% genome mapping, personalised medicines for every disease, nuclear fusion… and perfect credit decisioning (!) – all will be available via the miracle of AI (depending on how frothily you take your podcasts).

The current reality

And yet… despite astounding advances at the bleeding edge, the extent to which a productivity revolution is being felt in many industries is debatable. Some redundancies are occurring; some hiring rates are down – but whether this is driven entirely by AI is less certain. The research coming out of the AI labs is incredibly impressive, and the rate of progression is like no other technology in human history; but whilst the hype remains at fever pitch, the reality of delivering practical implementations across recognisable use cases has arguably been more limited to date.

In particular, finished AI "products" pitched at non-tech businesses have only just started to appear. They are not yet mature, adoption is limited and genuine benefits cases are relatively hard to find beyond small-scale mini-agents running fairly basic processes – e.g. meeting notes, agentic chatbots replacing call centres etc. The case for investing in these products is further complicated by the pace of change. Some first-generation AI products look increasingly out of date already in the face of the rapid advances in AI's power and capabilities. "Finished" AI technology products which non-tech companies can buy and deploy are in danger of becoming obsolete faster than the procurement cycle in most businesses can adopt them.

> Whilst the hype remains at fever pitch, the reality of delivering practical implementations across recognisable use cases has arguably been more limited to date.



Accelerating adoption

The term “AI” really represents a set of underlying tools (machine learning, RPA, OCR, NLP, agentic, generative etc.). As with any new tool, using it to good effect requires the selection of the right task, and then creating the conditions for it to be used competently and safely by skilled personnel. A scalpel in the hands of a trained surgeon in a well-staffed operating theatre is a tool for good; in the hands of a 2am Soho drunkard it’s just dangerous – to himself and others!

In practical terms this means companies wishing to deploy AI must first grapple with establishing a safe AI-operating environment. This requires consideration to be given to governance, access management, data security, regulatory compliance and model evaluation and auditability. These elements are not “compliance drag” or nice-to-haves; they are foundational requirements before proceeding to experiment with AI in any meaningful way.

The firms making the fastest progress and deriving meaningful value out of AI are likely to be those that do the governance work effectively, and early on. They can then assess new use cases, or new AI products, in a fortnight rather than reopening a six-month debate every time the technology moves or users suggest a new process transformation candidate. Given how quickly any individual product dates, a strong AI-deployment framework and capability is really the only part of the investment with a shelf life.

Inevitably, every operating business in every industry cannot find, hire, train and retain enough specialist resources to deliver AI adoption internally, so there has to be another way. This whitepaper explores the options for how mid-size firms might create a firm foundation for AI adoption safely, effectively and above all, quickly by using specialist external vendors and resources without needing billion-pound R&D budgets.

Purpose of this whitepaper

Who this whitepaper is for

This whitepaper is intended for CEOs, executive teams and senior leaders responsible for shaping how their organisation approaches AI. It is particularly relevant for COOs, CIOs, CTOs, CROs, Heads of Data, Heads of Risk and Transformation Directors who are balancing the need to encourage innovation with the need to manage regulatory, operational, data and reputational risk.

The whitepaper is designed to help these leaders understand the practical governance choices building societies are beginning to make, the common challenges emerging across the sector and the foundations required to adopt AI safely and responsibly.

It should also support more informed conversations about ownership, investment, risk appetite, operating models and where AI can deliver the greatest value for members, colleagues and the organisation.

Our perspective

Our perspective is grounded in a combination of direct sector insight and practical delivery experience.

It draws on interviews with building society leaders who are actively shaping their organisations' approaches to AI, alongside wider research into emerging AI governance practices and regulatory considerations.

This gives us a view not only of the principles that good AI governance should follow, but also of the practical challenges organisations face when turning those principles into clear ownership, workable controls and sustainable ways of working.

Governance without paralysis

Governance should make explicit what colleagues can do safely, not just what they cannot do

Publish clear guidance for colleagues regarding approved tools, the information that can be shared with those tools and the tasks that AI can be used for. Colleagues should also understand when human checks are mandatory within the AI workflow. Clarifying these guidelines enables colleagues to act confidently around AI without having to repeatedly navigate policy or seek approval for routine uses.

Where approval is required, offer teams a fast, visible route from idea to controlled experimentation. If the formal route is slower or more challenging to navigate than finding a workaround, governance risks driving AI use out of the building society's sight rather than making it safer.

Unmanaged AI use creates a larger risk than visible experimentation

An abundance of personal accounts, browser tools and supplier-enabled features should be interpreted as signals that the approved route is not meeting colleagues' needs. Monitor these patterns to understand what people are trying to achieve with AI and where existing tools or approval processes are creating friction. Then, respond by approving, constraining or blocking activity according to its risk.

Report the gap between approved capability and observed demand to whoever is the executive AI owner at the building society. The objective should be to bring useful experimentation into a controlled environment by making the sanctioned route more accessible and responsive than the workaround. Otherwise, unmet demand risks becoming invisible, unmanaged AI use.

> We had stifled our innovation with governance. The latter should be an enabler, not a blocker; they support each other, not get in each other's way.

- Charlie Moran, Skipton Building Society (Head of Data Capability)

Control intensity should rise with autonomy, reach and irreversibility

Routine activities like drafting emails or summarising meetings should be able to proceed quickly within published guardrails. Meanwhile, assurance should increase as AI becomes more autonomous on employees' behalf, reaches more customers, accesses sensitive data or becomes harder to reverse.

For higher-risk use cases, define clear stop conditions in advance and identify where human review and decision-making must remain within the workflow. This can allow AI to support more consequential activities while ensuring a human remains accountable for the final decision. These proportionate controls allow low-risk experimentation to move quickly while concentrating governance effort where the potential consequences justify doing so.

NatWest demonstrates how ethical governance can be applied without making every use case a Board-level decision. It uses ethical AI impact assessments for significant deployments, which are reviewed by a dedicated ethics team. A multidisciplinary AI and Data Ethics Panel considers higher-risk cases.[1]

What building societies can learn:
Routine, low-consequence uses should sit within published guardrails and manager oversight. Uses that influence member outcomes, regulated activity, sensitive data or material operational decisions should receive structured challenge and documented approval.

> For any uses beyond the boundaries that were originally established for inbuilt tools - just come and have a conversation. The answer is not a no by default. Let's talk it through and work out the best way forward.

- Ian Long, Suffolk Building Society (Head of Technology and Transformation)

Controlled experimentation should mature the framework

Offer controlled experimentation as the easiest way for colleagues to test AI ideas. Explicitly assign each use case a clear risk tier and accountable owner, then use pilots to test approval, monitoring, explainability and human oversight. The lessons from these experiments should continuously simplify and strengthen the safe route for the next use case.

Proportionate Board oversight and clear delegated authority help prevent approval bottlenecks. The Board should have visibility over significant AI decisions, supplier exposure and emerging demand, while allowing lower-risk experimentation to move quickly within established guardrails.

[1] Source: https://www.natwestgroup.com/news-and-insights/latest-stories/ai-and-data/2025/may/upholding-ethical-use-of-ai-and-data-management.html

Building AI on trusted data

Data governance should support defensible outcomes

Perfect data should not be the entry requirement for deploying AI. Instead, define the essential standard of data quality according to what the AI is being asked to do and the consequences of getting the answer wrong. A low-risk internal workflow can work with bounded, well-understood data and surface gaps or inconsistencies for human review.

Standards around data quality should rise with the stakes. AI influencing member decisions, regulatory activity or high-impact operational processes requires stronger evidence of accuracy, completeness and reliability. Therefore, building societies' objective shouldn't be to immediately make all organisational data AI-ready; instead, prioritise making the right data sufficiently reliable for the specific use case.

Every AI use case needs a data accountability chain

Before an AI use case pilot begins, name the data owner, process owner and AI-use owner, and establish who is responsible for determining whether the available data is suitable for what the AI is being asked to do. Record the relevant retention requirements and decisions about who or what can access the data.

Reassess that standard as the use case evolves. A dataset that is sufficient for a contained experiment may no longer be sufficient when the AI reaches more users, informs more consequential decisions or moves into a core organisational workflow. Owners should therefore reconfirm data suitability as the use case changes or scales.

AI can expose weak data practices, not merely wait for them to be fixed

Building societies do not need to complete a society-wide data transformation before they can start using AI. In fact, one of the most practical early applications of AI is to help identify and improve poor-quality data itself.

Within a bounded workflow, AI can compare sources, flag inconsistencies, identify missing or duplicated information, highlight conflicting definitions and surface records that require human review. This can help societies target data-quality effort where it matters most, rather than trying to cleanse the entire estate before any AI activity begins.

The important distinction is that AI should support the identification and remediation of data issues, not silently correct or infer its way around them. Material outputs should remain traceable to source information, with clear ownership for reviewing issues, approving corrections and addressing weaknesses at source.

Data access controls are central to AI governance

Define what data each AI use case actually needs rather than giving tools broad access simply because the information is available. Specify the datasets, fields and personal information the AI may access, and limit access to what is necessary for its intended purpose.

As a pilot scales or its purpose changes, revisit those permissions alongside the required data-quality standard. Broader access should be a deliberate governance decision, ensuring that the AI's access to information grows only where the expanding use case genuinely necessitates it.

> It would be good practice to complete data-retention and cleansing exercises before exposing data to AI, reducing noise and improving reliability - for example, using 1,000 clean rows rather than 20,000 inconsistent ones.

- Jonathan Hedley, Darlington Building Society (Data Governance Manager)

> Before getting too far into it, you need to understand exactly what the data is, where it is, its accuracy and have faith that it is properly managed and complete.

- Ian Long, Suffolk Building Society (Head of Technology and Transformation)

Federated innovation, central accountability

Business teams should define the problems AI addresses and the solutions it offers

AI opportunities should originate with the people who understand firsthand which processes can be automated. For example, underwriters, risk analysts and operations teams are often best placed to identify repetitive tasks, inefficient hand-offs and areas where AI could improve member or colleague outcomes. Rather than requiring a central AI team to find opportunities for them, offer business teams a simple route to bring problems forward and establish one standardised process for deciding if and how AI should be used to solve them.



This becomes particularly important in smaller societies, where a large central innovation function is unlikely to be practical or desirable. Ideas must feed into a consistent governance route so that local innovation does not spiral into unmanaged experimentation.

Nationwide's work with ARIS, a process intelligence platform, illustrates a foundational governance point: an organisation must understand its processes before it decides to automate them. Nationwide established a formal process-ownership framework, created a common process library and connected process models with risks and controls. It has also created visibility of the data, systems and teams involved in each process. Nationwide now sees this connected process intelligence as the foundation on which it can layer AI.[2]

What building societies can learn: AI readiness starts with a clear view of the process being changed. A society should use AI only where it understands the process well enough to govern it.

> If you centralise AI and it becomes one department, then how do you get innovation? Innovation with AI has got to come from people on the ground who say, 'I could use something different here.'
>
> **- Keith Edwards, Leek Building Society (Head of IT, Change and Data)**

Local champions should build capability within a controlled model

Use local AI champions to connect central governance with the teams closest to the opportunities. Champions can identify demand, share approved practices and direct colleagues towards the right intake and approval route, while being trained to recognise important data-handling requirements and accountability boundaries.

The champion model can itself scale with the society. A smaller building society may need only a handful of colleagues performing the role alongside their existing responsibilities, while larger societies can develop more formal networks supported by dedicated AI expertise. Their collective insight also provides a valuable view of recurring demand for AI.

> Our ideas and innovation are federated across the organisation, coming from colleagues who understand the problems and opportunities in their own areas. Governance is then applied centrally, with representatives from across the business bringing the right domain knowledge, expertise and challenge.
>
> **- Charlie Moran, Skipton Building Society (Head of Data Capability)**

[2] Source: https://aris.com/resources/nationwide-process-intelligence-ai-ready/ (Nationwide)

Lloyds describes an approach that combines a dedicated AI Centre of Excellence with responsible AI principles and embeds governance, risk and model-lifecycle processes across the business. It has also integrated AI risks into existing control and oversight structures. The result is an approach designed to enable responsible innovation at scale.[3]

What building societies can learn: the team that expects the operational or member benefit from utilising AI should own the outcome and first-line controls. A central coordinating role—such as an extension of an existing data, risk or change function—should set standards and define what the use case looks like.

The governance forum should provide central challenge and oversight

AI governance does not necessarily require building an entirely new committee. Smaller societies can adapt an existing risk, data, technology or change forum, provided it brings together the necessary perspectives and has clear authority to make decisions. As AI adoption and exposure increase, a dedicated or more frequent forum may become appropriate.

Business teams should remain responsible for identifying problems, pain points and inefficiencies, shaping potential use cases and ultimately owning the outcomes. However, the central AI team also has an important role in building awareness of what AI can and cannot do, so that business teams are equipped to recognise where it could genuinely help. The governance forum can then provide consistent challenge and oversight, drawing on compliance, legal, procurement, operations, security and relevant business expertise. This avoids centralising innovation while ensuring AI activity is governed consistently across the society.

Give the forum authority to approve, constrain, pause or retire AI use cases where necessary. Its effectiveness should be measured through the speed and quality of its decisions and its ability to reduce unmanaged activity, rather than the number of policies it creates.

Executive accountability must remain visible

Every building society should have clear executive accountability for AI, but that does not mean the CIO should automatically own it. AI affects customer outcomes, operations, risk, data and workforce design, so ownership should sit with the executive best placed to balance the opportunity with the associated exposure.

Societies should consider who is closest to the areas where AI is expected to create the greatest value, who has sufficient authority across functions and who has enough understanding of AI to challenge both its potential and its risks. Depending on the organisation, that could be the COO, CIO, CDO, transformation lead or another executive sponsor. What matters is that ownership is deliberate rather than inherited simply because AI is seen as a technology issue.

The Board should then challenge the organisation's overall AI direction, material exposures and governance effectiveness without becoming an approval committee for routine experimentation. Public approaches already vary: the Bank of England co-chairs its AI governance committee through its Chief Data Officer and Chief Information Officer,[4] while Lloyds Banking Group has a dedicated Chief Data & AI Officer reporting to the Group COO.[5]

[3] Source: https://www.lloydsbankinggroup.com/who-we-are/group-overview/artificial-intelligence.html

[4] Source: https://www.bankofengland.co.uk/about/governance-and-funding/staff-codes-and-policies/the-banks-ai-strategy

[5] Source: https://www.lloydsbankinggroup.com/media/press-releases/2026/lloyds-banking-group/sameer-gupta-chief-data-ai-officer.html

From adoption to accountable use



> You've got some people who are happy to play around with AI at home and know what they're doing, and you've got some people saying, 'I'm not touching AI.' Getting everyone to a good level is quite difficult.

- Trevor Tannenbaum, The Cambridge Building Society (Head of IT)

Training should develop sound judgement

Training should teach colleagues not only how to use AI, but how to exercise judgement over its outputs. Use realistic scenarios to show what information can be shared, which tasks are appropriate for AI, what specifically needs to be checked and when an output should be challenged or escalated.

Managers need additional capability to translate those principles into day-to-day work. They should be able to redesign workflows, define where human judgement remains necessary and intervene when AI use creates risk. The goal is for colleagues to understand what they remain accountable for when they do choose to use AI.

Responsible application determines AI value

Measure AI adoption through its responsible use in defined workflows, rather than licence numbers or the available access to tools. Look at whether colleagues use AI repeatedly for appropriate tasks, whether outputs meet the required standard and whether its use improves service, decision-making or organisational capacity.

Crucially, AI adoption is not as simple as colleagues applying AI to all of their tasks. Focus instead on identifying workflows where AI can create genuine value for building societies. From there, assign responsibility for reviewing AI's contribution and signing off the final outcome it produces.

> In the same way you train somebody about your policies, AI needs training too…so view it as an extension of the person rather than a mathematical equation that's just going to magically come and give you the answer.

- Keith Edwards, Leek Building Society (Head of IT, Change and Data)

Research published by Zango describes **Allica Bank** embedding AI adoption in its Company Scorecard so that each Executive Committee member is measured on how effectively they adopt AI in their function. It also uses an AI Risk Working Group and weekly company-wide AI "show and tell" sessions, allowing colleagues to build their confidence around AI and share practical uses with one another.[6]

What building societies can learn: leaders should be accountable for ensuring that newly released capacity improves the building society's customer service or colleagues' professional development.

Material AI use requires transparency

Human accountability is only meaningful if the building society can identify how AI contributed to an outcome. Require disclosure where AI contributes materially to customer communications, decisions or regulated activity, and retain sufficient evidence to understand how the final outcome was reached.

Employees should be able to clearly identify what the AI produced, what the human reviewed and who ultimately approved the outcome. This transforms "human-in-the-loop" from a broad principle into explicit guidelines.

Building society leaders must manage the implications of AI for their workforce

Leaders should plan for how traditional roles change as AI takes on more work. Identify which skills and forms of judgement become more important to upskill and where released capacity will be reinvested.

Avoid making commitments to widespread AI implementation without a credible plan for role redesign and capability development. Leaders should be able to understand where in their organisation AI does the work, where humans add their own judgement to AI workflows and who ultimately remains accountable for the outcome.

[6] Source: https://www.zango.ai/research-the-future-of-ai-governance---compliance-in-financial-services (Allica Bank)

Measuring value and managing the hidden AI estate



The AI estate includes every capability that can shape an outcome or expose data

Maintain a living inventory that extends beyond the AI tools the building society has formally purchased. Include internally built tools, pilots, browser services, personal-use tools and AI functionality embedded within existing supplier products. Anywhere AI can access organisational data or shape an outcome ought to be noted in the inventory.

The inventory should make the estate governable: record what each capability does, who owns it, what data it accesses and its supplier dependencies.

According to the Bank of England, **1/3** of AI use cases in UK financial services are **third-party implementations**, but **46%** of firms report only a **partial understanding** of the AI technologies they use.[7]

Third-party AI governance must continue after procurement

Approving a supplier once does not mean its potential future AI risk remains static. We recommend that building societies require suppliers to disclose relevant AI functionality, model hosting and assurance arrangements. Establish contractual rights to notification and assessment before material new AI capabilities are introduced.

Continue that oversight throughout the relationship, including monitoring supplier release notes and product changes. AI can arrive unknowingly via an update to software the building society already uses, so they need a way to identify when the risk profile of an existing supplier has changed.

> A tool that we've recently started using has highlighted that there was a lot of hidden AI around the business through third parties, which we are blocking. Having more granular controls has allowed us to roll out more AI tools confidently.
>
> **- Trevor Tannenbaum, The Cambridge Building Society (Head of IT)**

[7] Source: https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

Zango's research found that in several regulated firms, risk and compliance functions had limited insight into which AI tools were being used across the organisation. One compliance leader shared they could not answer where AI was being used across the organisation. The report identifies decentralised deployment, embedded supplier capability and limited central coordination as key causes.[8]

What building societies can learn: maintain a living inventory that covers approved tools, pilots, browser-based services, internal workflows, supplier-enabled functionality and AI used within control functions.

Technology and vendor choices should follow the use case

When making choices around selecting AI, start with the business problem, then assess whether an existing enterprise capability, specialist vendor or internal workflow build provides the most appropriate tool. Decisions should consider data access, architecture fit, total cost, expected value and risk appetite.

Set architecture standards, integration routes, approval checkpoints and handover requirements before production deployment. This helps prevent an accumulation of disconnected tools and ensures additions to the AI estate both solve a defined problem and can be governed effectively.

Leaders should manage AI as a portfolio of investments and exposures

Before implementation, establish relevant measures such as service quality, turnaround time, decision quality, control performance and colleague capacity. Assess benefits against the full cost of the use case.

Use this view of AI as a portfolio to make active choices about the AI estate over time. Scale use cases that demonstrate value within the society's risk appetite, while pausing or retiring those that have weak adoption. Leaders' objective should be to build an estate in which each use case has a clear purpose, owner and acceptable risk level.

Supplier concentration is an AI governance issue

Visibility of the AI estate should extend to the infrastructure beneath individual tools. Map where multiple important processes ultimately depend on the same model, as seemingly separate AI applications may create a shared point of failure. Test what happens if a supplier becomes unavailable, service degrades, a model is withdrawn or an unapproved feature is introduced.

> Assurance and risk management approaches for third-party vendors in relation to the use and supply of AI solutions are difficult to get right. The more challenging areas are the products integrating AI where we may not be aware of it.
>
> **- Ian Long, Suffolk Building Society (Head of Technology and Transformation)**

[8] Source: https://www.zango.ai/research-the-future-of-ai-governance---compliance-in-financial-services

AI Governance Framework

What you will learn by completing this checklist:

- Where your strongest governance foundations already exist.
- Which of the five dimensions needs attention first.
- The practical actions that would make an immediate difference over the next 90 days.
- Whether existing risk, change, data, procurement and operational-resilience processes can be adapted, rather than creating new layers of governance.
- The questions your executive team and Board should be asking as AI use expands.

How your results compare:

Scan the QR code to submit your completed AI Governance Framework online. If you share your completed checklist with Woodhurst, we can provide a confidential, anonymised view of how your responses compare with the wider building society sector. **We will not name individual societies or disclose their results.** Instead, we will show where your approach appears broadly aligned with peers, where you may be ahead of the emerging norm and where a small number of targeted improvements could strengthen your confidence in your AI governance. You can also email your results to **Ahmad Al-Tahhan (ahmad.altahhan@woodhurst.com)**, **Ben Nadel (ben.nadel@woodhurst.com)** or **James Silk (james.silk@woodhurst.com)**.

SCAN HERE

Dimension 1: Governance without paralysis

- ☐ We have published the AI tools colleagues may use, the information they may share and the cases that require human review.
- ☐ We have a clear route for teams to propose, test and approve an AI use case.
- ☐ We assess use cases according to their potential impact (e.g. member reach, data sensitivity and reversibility).
- ☐ The Board receives information on material AI decisions and any emerging demand for specific AI tools or workflows to be automated.

Dimension 2: Building AI on trusted data

- ☐ Each pilot has a defined data scope, named data owner, process owner and AI-use owner.
- ☐ We have determined whether the data is sufficiently accurate and complete for the specific AI use case.
- ☐ We define what a reviewer must do when source information is missing, inconsistent or unsuitable.
- ☐ We reassess data quality, retention and access permissions when a use case changes or scales.

Dimension 3: Federated innovation, central accountability

- ☐ Teams have a simple route to bring forward operational problems that may be suitable for AI.
- ☐ We maintain a single register of proposed, approved, restricted, paused and retired AI uses.
- ☐ We have identified the existing forum or role responsible for coordinating AI standards across the building society.
- ☐ We have emphasised to colleagues that owners remain accountable for outcomes where AI is used.

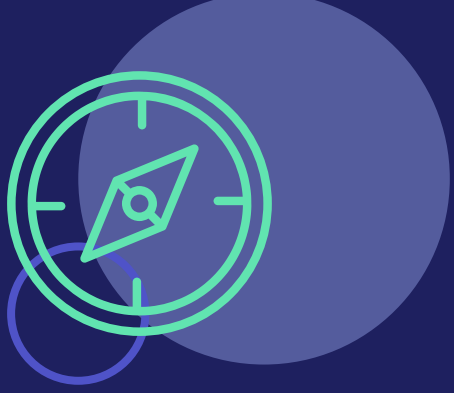
Dimension 4: From adoption to accountable use

- ☐ AI training uses realistic examples from the building society's own processes rather than generic demonstrations of AI tools.
- ☐ For each AI use case, we define what the human reviewer must check, what remains their decision and when they must escalate.
- ☐ We retain sufficient evidence to show what the AI produced, what the human reviewed and who approved the final outcome.
- ☐ We measure AI adoption through use case quality and newfound colleague capacity, rather than licence numbers alone.

Dimension 5: Measuring value and managing the hidden AI estate

- ☐ Our AI inventory includes purchased tools, supplier-enabled features, pilots, internal builds and browser-based services.
- ☐ Third-party suppliers must disclose material AI functionality, data flows, hosting arrangements, sub-processors, assurance evidence and planned changes.
- ☐ We have a process for reviewing supplier releases and assessing new AI functionality before it is activated within the building society.
- ☐ We have contingency arrangements in place for dependencies, such as supplier outage or unapproved feature release.

Conclusion



It's a misconception that building societies need to achieve "perfect" data or complete regulatory certainty before making meaningful inroads with AI. Rather, building societies can begin to develop their AI ecosystems with several concrete steps. These include setting boundaries for colleague use, starting with workflows and data that are well understood and assigning responsibility for the outcome from the outset.

Good governance can help catalyse progress, while overly restrictive controls risk creating workarounds or AI workflows that fall outside of the society's visibility. Allow routine, low-risk uses to operate within clear guardrails. Controls should become more rigorous when AI affects members, informs regulated activity, accesses sensitive information or acts with greater autonomy. Critically, accountability remains with the building society; a supplier or model cannot assume responsibility for outcomes.

The most effective approach to AI governance will be proportionate. Small building societies do not necessarily need to build governance foundations from scratch. Instead, focus on obtaining visibility of where AI is being used, providing a route for teams to bring forward ideas for use cases and retaining the ability to cease AI activity when required.

Regardless of their size, all building societies should be able to show what an AI tool was allowed to access, what it produced, what a human reviewed and who made the final decision. They should monitor the value created and the supplier dependencies introduced. This is particularly important as AI becomes embedded in everyday workflows and as building societies oversee an array of third-party tools.

Used well, AI can reduce manual effort, improve service and ultimately enable colleagues to focus on tasks involving skills that cannot be automated: judgement, expertise and empathy.

About the authors



Ahmad Al-Tahhan
LEAD AUTHOR
MANAGEMENT CONSULTANT



James Silk
DIRECTOR



Ben Nadel
CO-FOUNDER

With thanks to our contributors:

KEITH EDWARDS
Head of IT, Change and Data at Leek Building Society

JONATHAN HEDLEY
Data Governance Manager at Darlington Building Society

STEVE HOGG
Chief Operating Officer at Enra Specialist Finance

IAN LONG
Head of Technology and Transformation at Suffolk Building Society

CHARLIE MORAN
Head of Data Capability at Skipton Building Society

TREVOR TANNENBAUM
Head of Information Technology at The Cambridge Building Society

About Woodhurst

Woodhurst is a specialist consultancy focused on helping financial institutions define and turn strategy into tangible change through digital modernisation.

Since 2019, we have worked closely with building societies, banks and specialist lenders on the practical realities of transformation - from shaping strategy and redesigning operating models, through to implementing technology, strengthening data foundations and embedding new ways of working that deliver lasting impact.

In 2026, Woodhurst became part of Finyx Group, an established organisation with nearly two decades of experience, more than 100 consultants and offices in Belfast, London and York. Together, the businesses offer clients a broader range of services and the scale to deliver complex transformation programmes, while retaining a close focus on client outcomes.

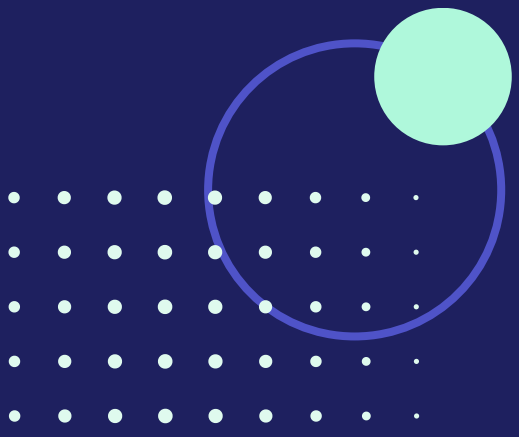
What defines our approach is how we work. We operate alongside our clients, not at a distance, combining hands-on delivery experience with a clear focus on outcomes. As a founder-led consultancy, our teams are deliberately small, highly experienced and deeply embedded in the organisations we support.

Across our work, we see first-hand how quickly the digital landscape is evolving and how challenging it can be to translate ambition into execution. Our role is to help bridge that gap - bringing clarity, pace and structure to complex transformation agendas.

If this whitepaper resonates with the challenges you are facing, we welcome the opportunity to continue the conversation.

Get in touch:

ben.nadel@woodhurst.com
james.silk@woodhurst.com
ahmad.altahhan@woodhurst.com





Woodhurst

We believe that everyone should have access to better financial solutions.

www.woodhurst.com