

Responsible AI Governance & Policy Development: A How-to Guide for Executives

Governance, Policy, and Ethical Use of AI



CONTENTS

.....	
Foreward	3
1. Executive Summary	4
2. Core Concepts: AI Governance in the Housing Sector	6
3. AI Governance Maturity Ladder	8
4. Seven-Step Guide to Developing Your AI Governance Policy	9
5. Case Studies	13
6. Implementation Strategy	15
7. What's Next: Trends & Tools	17
8. Appendix	19
Appendix 1 Templates, Tools & Checklists	19
Appendix 2 Resources	22
About Demystifying AI for Social Housing (DASH)	23

Disclaimer: This resource is for general guidance in the social housing sector. It is not legal or professional advice. See our [full disclaimer](#) for permitted use and limitations.

Foreward

When I first started hearing about the rise of AI, the type of AI that we can get our own hands on as end-users, I thought it sounded fun, but probably not something that would impact my work much. I couldn't have been more wrong!

As the possibilities of AI-powered technology increasingly seem almost unlimited, the social housing sector is, understandably, looking to harness that power. When the potential is so vast, though, it can be a challenge knowing where to start. So, we need direction and a starting point, which DASH has provided in this guidance; a clear and pragmatic roadmap to help ensure the use of AI enhances the achievement of the sector's purpose.

While my sphere of expertise does not include coding, algorithms, or technology of most types, it does include data. I specialise in personal data, from a legal and strategic perspective, but my experience and work also expand to other data, including asset data, having worked across several different teams in social housing over the last two decades. One thing I do know about AI is that it is nothing without data. And without purposeful data, AI and its outcomes lack purpose too. Without clarity of purpose when launching new technologies, we risk spending precious time and resources going down roads that take us in the wrong direction, away from, instead of closer to, our goals.

This DASH guide won't tell you exactly what you can or can't do with AI, or what you should or shouldn't do, in a dogmatic or unnecessarily restrictive way; instead, it highlights the questions to ask – and answer – and the risks and controls to consider, with helpful and practical suggestions, to ensure you can confidently engage with AI technology, with purpose.

[Clare Paterson](#)

Data Strategist & Data Protection Champion

Social Housing Data Specialist

Author- Practical Guide to Data Protection in Social Housing.

Founder of the Data in Social Housing (DISH) Network

1. Executive Summary

Why responsible AI starts with responsible leadership.

Artificial Intelligence (AI) is no longer a future technology; it's already here, reshaping how housing organisations operate, communicate, and serve tenants.

Across the UK, housing providers are using AI to;

- Predict arrears and prevent eviction,
- Streamline maintenance scheduling,
- Improve tenant communications through smart chatbots and automation.

These tools offer powerful opportunities for **greater efficiency, more personalised services, faster response times**, and smarter resource use. But they also carry very real risks.

The Risk Side of the Equation: AI systems, if poorly governed, can:

- Amplify **bias** and discrimination,
- Breach of **tenant privacy**,
- Undermine **trust** if decisions are opaque or unfair.

Add to this a rapidly tightening regulatory environment, including **UK GDPR**, the **Equality Act 2010**, the forthcoming **EU AI Act**, and the UK **ICO's AI guidance in development**, and the action case becomes clear. In short, housing providers cannot afford to adopt AI without a clear governance framework. The consequences, both legal and reputational, are too high.

Why This Guide, and Why Now?

This guide is designed to help **housing executives and leadership teams** take control of AI not as technologists, but as ethical leaders. It

It gives you:

- A clear seven-step approach to AI governance,
- Practical tools and policy templates you can adapt today,
- Real-world case studies from the housing sector,
- Foresight into the trends and regulations shaping tomorrow.

Whether you're just exploring AI or already using it across your services, this guide will help you **go beyond compliance** and build a governance approach rooted in trust, transparency, and tenant dignity.

Bottom line? AI is already shaping the future of housing, and governance is how we make sure it does so fairly, safely, and for everyone.

Quick Start: Begin Your AI Governance Journey

3 things to do this month. 2 key risks to assess. 1 success metric to track.

Whether you're just getting started or already piloting AI tools, here's how to take focused, responsible action right now — without getting bogged down in complexity

3 Things to Do This Month

1. **Week 1 - Map your AI activity.**

Find out where AI is already being used or discussed, from arrears prediction tools to staff using, eg Chat GPT.

2. **Week 2- Assign ownership**

Appoint a senior lead (e.g. Director of Transformation) and an operational owner (e.g. DPO or data lead).

3. **Week 3 - Create a basic AUP**

Introduce an Acceptable Use Policy to guide how staff use generative AI tools. It prevents well-meaning misuse.

2 Risks to Assess Immediately

1. **Data leakage**

Are staff putting tenant information into public AI tools? If yes, act now.

2. **Unexplained decisions**

Can you explain how your AI-supported tools make decisions, and who is accountable for them?

Metric to Track: “What % of AI-supported decisions are explainable to staff and tenants?”

If the answer is “not many,” that’s your governance gap.

2. Core Concepts: AI Governance in the Housing Sector

Setting the foundations for ethical, effective, and human-centred AI

Before you can lead the responsible use of artificial intelligence (AI) in your housing organisation, it's essential to understand what governance means in this context, and how it applies to both tenants and staff. Let's break it down.

What Is AI Governance?

This refers to the structures, principles, and processes that ensure AI systems are used **safely, fairly, legally, and transparently**. It's about assigning responsibility, managing risk, and making sure AI serves people, not the other way around.

In the housing context, this means protecting tenant rights, guiding staff behaviour, and aligning with organisational values like trust, inclusion, and accountability.

AI Governance Framework Comparison

Framework	Focus	Key Contribution
ISO/IEC 42001	AI management systems, oversight, risk	Sets governance requirements across the AI lifecycle.
OECD AI Principles	Human values, fairness, transparency	Ethical foundation adopted by 40+ countries.
NIST AI RMF	Risk mapping, explainability, trust	A practical approach to AI risk and trustworthiness.

Rather than picking one framework, housing providers should **blend the best elements** of each to meet sector-specific needs.

Policy vs AUP vs Risk Governance: What's the Difference?

AI governance isn't a single document – it's a set of tools working together. Let's clarify the roles of your key policy components:

Term	What It Does	Housing Example
AI Policy	Sets your organisation's position and values around AI use	"We commit to transparency, fairness, and human oversight."
Acceptable Use Policy (AUP)	Defines how staff can and cannot use AI tools	"Staff may not enter tenant data into public generative AI platforms."
Risk Governance Framework	Details on how AI risks are assessed and mitigated	"All AI tools undergo a DPIA and bias audit before rollout."

Think of it this way:

- Your **AI policy** is the compass.

- Your **AUP** is the user manual.
- Your **risk framework** is the safety harness.

Why Tenant-Centred Ethics Matter

In housing, AI decisions affect real lives, from property allocation to arrears escalation. Ethics must be baked into every step:

- Human Oversight – AI should support, not replace, human judgment.
- Transparency – Tenants must know when AI is involved
- Fairness – Avoid bias by design, especially for vulnerable groups.
- Privacy – Respect and protect tenant data in line with GDPR and the Equality Act.

The Role of Independent Oversight

Both ISACA and the Responsible AI Institute recommend that organisations establish **independent, cross-functional oversight** for AI systems, especially those impacting people's rights. What that looks like:

- An **AI Steering Committee** with members from IT, legal, housing services, and tenants.
- **Ethical reviews are embedded early** in project design.
- **Peer audits or external checks** for high-risk use cases (e.g. automated tenant profiling).

Think of it as your organisation's **moral firewall** — catching ethical blind spots before they become real-world harm.

Key Takeaways

- AI governance is about **boundaries and trust** with tenants, staff, and regulators.
- **AI policy, AUPS, and risk frameworks** each play a distinct, essential role.
- Ethical, **tenant-centred design** must be embedded from the start.
- **Independent oversight** isn't optional — it's best practice and a future regulatory requirement.

Ready to make this real? In the next section, we'll walk through a practical, step-by-step process for building your own AI governance policy.

3. AI Governance Maturity Ladder

A self-assessment model for housing organisations

Use this ladder to understand your current position and identify practical steps toward responsible AI maturity. You don't need to leap to the top — progress is what matters.

Level	Description	What You Might See	Next Step
1. Ad Hoc	AI tools are used informally without governance. No policies or oversight exist.	Staff using Gen AI (e.g. Chat GPT) without guidance; AI features built into supplier systems with no review.	Map where AI is being used. Introduce a basic Acceptable Use Policy (AUP). Assign a named lead.
2. Emerging	Early awareness of AI risks and potential. Some policies or pilots are in place, but fragmented.	One or two AI projects with informal DPIAs or ethical reviews; training is ad hoc.	Form a working group. Draft an AI policy aligned with your values and risk appetite.
3. Defined	Governance structures exist. AI policy, risk assessments, and staff training are active.	AUPs, ethics principles, and risk checklists are in use. Tenants are informed about some AI tools.	Conduct regular bias audits and DPIAs. Involve tenants in ethical oversight.
4. Embedded	AI governance is integrated into business-as-usual. Oversight is cross-functional and proactive.	AI tools are logged and monitored; a governance board is in place; procurement includes AI clauses.	Align AI impact to ESG targets and KPIs. Performance report.
5. Leading	AI governance is sector-leading. Practices are shared externally. Community and tenant voices drive innovation.	Transparent decision registers, ethics boards with tenant reps; regular audits and public reporting.	Support others in the sector. Contribute to policy development and housing innovation networks.

Tip for Use

- Include this table in team away days or leadership workshops.
- Let departments self-assess and compare results.
- Use it to shape your **AI roadmap** over the next 6–12 months.

Remember: It's not about being perfect, it's about being intentional, transparent, and tenant-centred.

4. Seven-Step Guide to Developing Your AI Governance Policy

Build confidence. Reduce risk. Protect tenants.

AI is moving fast, but governance shouldn't be rushed. This seven-step process gives you a clear, practical way to build a policy that protects your organisation, supports your staff, and ensures AI tools are deployed responsibly across your housing services.

Step 1: Establish Governance Structures

Purpose: Lay the foundation for accountability, oversight, and cross-departmental coordination.

Actions:

- Set up an **AI Steering Committee** for strategic oversight and risk appetite.
- Form an **Operational AI Working Group** to handle day-to-day use cases.
- Assign an **Executive sponsor** to lead governance.
- Include a **Data Protection Officer** to ensure personal data is lawfully used, controlled.
- Include **tenant representatives** on ethics boards or advisory panels.

Tools & Templates:

- [Committee Terms of Reference – Responsible AI Institute](#)
- [Role Descriptions for AI Sponsor & DPO – ISO/IEC 42001 Annex](#)

Housing Note: Tenant input is essential, not optional. If AI affects access to homes or services, the lived experience must help shape the rules.

Step 2: Define AI Principles & Organisational Commitments

Purpose: Set the ethical compass that guides your AI decisions.

Actions:

- Adopt global AI principles: **fairness, explainability, privacy, transparency, and accountability.**
- Add sector-specific values like **tenant dignity, non-discrimination, and digital inclusion.**
- Align to frameworks like **OECD, NIST, and ISO 42001.**

Tools & Templates:

- [AI Ethics Charter Template – RAI/](#)

Housing Note: One provider used “dignity” to guide the tone of its chatbot, proving principles aren't just theoretical; they shape real experiences.

Step 3: Create Staff Use & Acceptable Use Policies (AUPs)

Purpose: Empower staff to use AI confidently and safely.

Actions:

- Draft a **Generative AI Acceptable Use Policy (AUP)** covering approved tools, restrictions, and staff responsibilities.
- Integrate AUPs with **HR, IT, and safeguarding policies**.
- Provide **training and escalation pathways** to build awareness and reduce misuse.

Tools & Templates:

- [AUP Template – ISACA GenAI Policy \(2023\)](#)
- [Staff Guidance Checklist – UC Berkeley GenAI Playbook](#)

Housing Note: An officer once pasted tenant data into ChatGPT to rewrite a letter. AUPs stop well-meaning mistakes from becoming privacy breaches.

Step 4: Embed Tenant-Centric Data Protection

Purpose: Respect tenant rights and meet legal duties around privacy and data use.

Actions:

- Complete a **Data Protection Impact Assessment (DPIA)** for every tenant-affecting AI system.
- Apply **privacy-by-design**: collect only what's needed, anonymise where possible, and limit access.
- Ensure compliance with **GDPR (Article 22)** and the **Equality Act 2010**.
- Use clear, accessible language to explain AI use to tenants.

Tools & Templates:

- [DPIA Template for AI – ICO AI Toolkit](#)
- [Camden Council's Data Charter](#)

Housing Note: Transparency builds trust. Open communication with tenants about AI reduces complaints and improves service confidence.

Step 5: Mitigate Bias & Ensure Fairness

Purpose: Prevent discrimination and ensure equal treatment in automated decision-making.

Actions:

- Use **bias detection tools** (e.g. SHAP, AI Fairness 360, Fairlearn) to audit algorithms.
- Apply **synthetic data** to address underrepresented groups.
- Ask vendors to disclose:
 - Training data sources
 - Bias mitigation methods
 - Human oversight processes

Tools & Templates:

- [Bias Audit Log – IBM, Microsoft Fairlearn](#)
- [Vendor Fairness Questionnaire](#)

Housing Note: Together Housing reduced evictions by 18% using fairness audits and community consultation to shape AI support models.

Step 6: Build Transparent Procurement & Risk Processes

Purpose: Ensure third-party AI tools meet your ethical and legal standards before deployment.

Actions:

- Use a **procurement checklist** aligned to ISO 42001 and UK AI procurement guidance.
- Require **documentation** on model design, risks, and explainability.
- Build contracts that include:
 - **Audit rights**
 - **Liability clauses**
 - **Exit terms**
- **Pilot** first, evaluate performance, and gather feedback before full rollout.

Tools & Templates:

- [AI Procurement Checklist – UK Government, ISO 42001, WEF](#)

Step 7: Monitor, Review & Report

Purpose: Make AI governance a continuous, living process — not a one-time policy.

Actions:

- Build a **Responsible AI Dashboard** to track:
 - AI tool usage
 - Complaints and interventions
 - Training completion
- Schedule **annual governance reviews**.

- Stay aligned with updates from **ICO, Ofcom, EU AI Act**, and adjust your policy accordingly.

Tools & Templates:

- [AI Risk Management Framework](#)

Housing Note: Like safeguarding, your AI policy should evolve with practice, regulation, and real-world impact.

Key Takeaways

- **Establish Ownership:** Create clear structures with C-suite leads and tenant representation.
- **Lead with Principles:** Use fairness, privacy, and dignity to guide all decisions.
- **Support Staff:** AUPs and training protect people and build digital confidence.
- **Protect Tenants:** Use DPIAs and transparent communications to safeguard rights.
- **Mitigate Bias:** Audit and question your data before it harms.
- **Procure Wisely:** Vet your vendors — and pilot before scaling.
- **Monitor Continuously:** Governance must adapt to tech, law, and lived experience.

Bottom line? A good AI policy doesn't slow you down — it gives you the confidence to move fast and *responsibly*.



5. Case Studies

What's working, what's a warning, and what's worth watching?

Real-world examples bring governance to life. These four cases show how AI can either enhance or endanger trust, depending on how it's handled. Each illustrates the power of clear oversight, ethical design, or the consequences of getting it wrong.

Together Housing – Predictive AI, Ethically Applied

Challenge: Prevent tenancy breakdowns while avoiding bias in predictive tools.

Action: [Together Housing](#) developed an AI model using arrears and support data, ran a DPIA, and blended synthetic data to reduce bias. They set up an internal ethics committee and involved frontline staff to interpret results.

Outcome:

- 18% reduction in evictions
- Fewer false positives
- Improved tenant and staff trust

Key Lesson: Governance built into the design phase — including fairness checks and human oversight — makes AI both safer and more effective.

Mobysoft – Repairs Scheduling with Tenant-First Design

Challenge: Optimise repair scheduling while maintaining fairness and transparency.

Action: [Mobysoft](#) used explainable decision trees, integrated tenant feedback, and allowed staff to override AI decisions. Piloted the model across providers before scaling.

Outcome:

- 92% tenant satisfaction
- Fewer missed appointments
- Strong staff confidence in the system

Key Lesson: Transparency and explainability drive adoption. When frontline teams understand the logic, they trust the system, and so do tenants.

Air Canada – When Chatbots Go Rogue

Challenge: [An AI chatbot](#) gave incorrect fare advice with no oversight or disclaimers.

Action: The chatbot incorrectly promised bereavement fare discounts. There was no process to validate its outputs or alert users that it was AI-generated.

Outcome:

- Air Canada was held legally liable
- Financial penalty and reputational damage
- Public exposure of governance failings

Key Lesson: AI must be accountable. Always disclose when a customer is interacting with AI, and ensure there's human oversight for high-stakes issues.

PwC – Generative AI at Scale, Done Right

Challenge: [PWC](#) Enable staff-wide use of generative AI without compromising data security or ethics.

Action: PwC deployed ChatGPT Enterprise in a secure, private instance with role-based access. They implemented internal governance processes and trained staff on responsible use.

Outcome:

- Safe adoption across 75,000+ users
- No data breaches or misuse
- Positive cultural shift toward AI as a supportive tool

Key Lesson: You can scale genAI responsibly — if you invest early in policies, controls, and culture.

6. Implementation Strategy

Bridging the gap between policy and progress.

You've built the policy. You've defined the principles. Now it's time to **bring it to life** — not just in documents, but in how your people work, how tenants are served, and how your organisation evolves. A governance framework is only as effective as its execution. This section outlines how to roll out your AI governance strategy with purpose, care, and momentum.

1. Build Internal Support

Why it matters: People don't trust what they don't understand. And they won't use responsibly what they don't feel ownership of.

How to do it:

- Deliver **awareness sessions** tailored for different roles — leadership, frontline staff, IT, and compliance.
- Appoint **AI champions** across departments to model best practices and field local questions.
- Embed AI governance into induction, safeguarding training, and team meetings.
- Make it part of your culture, not just your compliance regime.

Tip: Use real examples (good and bad) to show what responsible AI looks like in action.

2. Start Small — Then Scale

Why it matters: AI doesn't need a grand launch. It needs a careful, confident introduction.

How to do it:

- Choose low-risk, high-value use cases to pilot your governance model, such as:
 - **Arrears prediction** (support over sanction)
 - **Repairs triage** (efficiency meets equity)
 - **Tenant communications** (chatbots or smart FAQs)
- Apply your full governance process — risk assessment, DPIA, oversight — even for pilots.
- Monitor performance, collect feedback, and refine your tools.

[Lesson from early adopters?](#) Pilots reveal gaps in policy and perception — *that's a feature, not a failure.*

3. Educate Tenants

Why it matters: Tenants deserve to know when AI is involved in decisions that affect their homes, families, or support. **How to do it:**

- Publish **plain-language explanations** of how AI is used (e.g. “How we prioritise repairs using technology”).
- Hold **tenant workshops or drop-ins** for major changes.
- Include AI transparency in your privacy notices, tenancy agreements, and annual reports.
- Invite tenant reps to sit on your AI ethics or oversight groups.

Respectful transparency isn't just good ethics — it builds trust, reduces complaints, and pre-empts escalation.

4. Use Iterative Feedback Loops

Why it matters: AI evolves. So should your policy. How to do it:

- Create **feedback channels** for staff and tenants to flag AI concerns or unintended effects.
- Log and review **AI incidents** (e.g. unfair decisions, unexpected outputs, system errors).
- Set **annual policy reviews**, or faster if legislation or usage changes.
- Treat your AI governance pack as a **living document** — designed to grow with you.

The best policy is the one that adapts without losing its principles.

5. Link AI to KPIs and Strategic Goals

Why it matters: AI should support your mission, not just your operations. How to do it:

- Align AI outcomes with existing organisational priorities, such as:
 - **KPI: Reduce evictions** → Use AI to predict and prevent tenancy failure
 - **ESG: Cut emissions** → Use AI for predictive maintenance and energy optimisation
 - **Digital inclusion targets** → Use AI to personalise digital services, not gatekeep them
- Report on AI impact through your board, ESG dashboards, or tenant panels.
- Highlight ethical governance in funding bids, regulatory reports, and annual reviews.

When AI delivers both social value and performance gains, everyone wins.

Key Takeaways

- **Make it real:** Training, awareness, and local AI champions are essential to embed policy in practice.
- **Start smart:** Pilot AI in manageable areas like arrears or repairs — test governance early, scale responsibly.
- **Respect tenants:** Be transparent about AI use, involve tenants in oversight, and communicate in plain language.
- **Stay agile:** Use feedback loops, incident logs, and policy reviews to keep your governance framework responsive.
- **Align to strategy:** Connect AI to organisational KPIs, ESG goals, and service improvements — governance becomes value, not overhead.

Bottom line? AI governance isn't a one-off launch — it's a leadership habit. Start where you are, and grow with purpose.

7. What's Next: Trends & Tools

Foresight meets follow-through: Prepare for what's coming and act with what you've got.

As AI evolves, so too must your governance approach. This section gives you a double advantage:

- **What to expect** in terms of policy, regulation, and innovation
- **What to use now** to implement your AI governance policy with confidence

Emerging Trends: What to Watch (and Prepare For)

Regulation Is Tightening

- [EU AI Act \(2024–2026\)](#): Requires risk classification, transparency, and human oversight, especially for “high-risk” systems like housing allocation.
- [UK ICO AI Guidance \(in development\)](#): Focuses on AI under GDPR and the Equality Act.
- [Ofcom Chatbot Standards \(2026\)](#): Will mandate disclosure and escalation pathways for AI in customer communications.

What this means: Housing providers will need to document how AI is used, who is accountable, and what safeguards are in place.

Technology Is Getting Smarter — and Riskier

- **Predictive welfare analytics:** Used to identify tenants at risk — powerful, but sensitive.
- **Digital twins & smart asset management:** AI models of housing stock to optimise investment and energy use.
- **Gen AI in tenant services:** Chatbots, draft letters, case triage — often helpful, but need controls.

What this means: These tools must be transparent, explainable, and fair — or they risk harming those they aim to help.

Governance Toolkit: What to Use Now

The following tools and templates are available to help you put your AI policy into action.

Toolkit Element	What It's For	Source
AI Use Case Risk Assessment	Triage risk before starting a project	UC Berkeley, ISACA
Generative AI Acceptable Use Policy (AUP)	Sets staff expectations and rules	ISACA
AI Ethics Charter	State your values clearly	Responsible AI Institute
DPIA Template for AI	Assess tenant data risks	ICO Toolkit
Bias & Fairness Checklist	Prevents discrimination in models	IBM, Microsoft, SHAP
Procurement Checklist & Due Diligence Questions	Evaluate vendor systems and risk posture	UK Gov, WEF, ISO 42001
Performance Dashboard & Policy Tracker	Monitors AI usage, reviews, and incidents	RAII, NIST RMF

You can find these in the **Appendix & Resources** section — ready to adapt and embed.

Key Takeaways

- **Regulators are catching up,** and housing will be in scope. Prepare now for the upcoming rules on transparency, decision-making, and AI ethics.
- **Transparency is the new normal** — automated decision registers and audit logs will become best practice.
- **AI in housing is expanding fast,** from predictive analytics to smart energy and tenant support tools. Governance must keep pace.

8. Appendix

Appendix 1 Templates, Tools & Checklists

Your practical toolkit for building AI you can trust.

You've seen the strategy. You've read the steps. Now let's get hands-on.

This section gives you the **ready-to-use tools** you need to apply AI governance in your housing organisation, whether you're launching your first chatbot or reviewing an existing predictive model. Each template is drawn from **recognised frameworks, real-world case studies, and best practice guidance**.

1. AI Use Case Risk Assessment Template

Purpose: Quickly assess whether an AI project poses low, medium, or high ethical and operational risk before development or procurement.

What It Covers:

- Impact on tenant rights and services
- Sensitivity of the data involved
- Level of automation (human-in-the-loop vs fully automated)
- Transparency and explainability concerns

Source: [UC Berkeley Responsible Use of Generative AI Playbook](#) & [ISACA Risk Guidance \(2023\)](#)

How to Use It: Use this template in discovery workshops to prioritise which AI projects need deeper review (e.g. DPIA, bias audit). Add it to your business case checklist.

2. Generative AI Acceptable Use Policy (AUP) Template

Purpose: Define how employees can use tools like ChatGPT, Microsoft Copilot, or similar genAI systems, and what's strictly off-limits.

What It Covers:

- Prohibited behaviours (e.g. pasting tenant data into public tools)
- Approved tools and data sources
- Staff responsibilities for review and accuracy
- Escalation and exception processes

Source: [ISACA Considerations for Implementing a Generative AI Policy \(2023\)](#)

How to Use It: Adopt it as part of your staff code of conduct. Share in induction packs and include in all roles using AI-powered admin tools.

3. AI Procurement Checklist

Purpose: Ensure third-party AI tools meet your governance, ethical, and compliance standards before you sign a contract.

What It Covers:

- Model documentation and training data transparency
- Bias and fairness testing by vendor
- Explainability and auditability of outputs
- Privacy and data protection assurances
- Ongoing support, updates, and exit plans

Source: [UK Government AI Procurement Guidelines](#), [ISO 42001 Annex B](#), [WEF AI Procurement in a Box](#)

How to Use It: Attach to tender specs or supplier onboarding processes. Make it non-negotiable for any AI tool used in tenant-facing services

4. Bias Mitigation & Fairness Checklist

Purpose: Identify, test for, and mitigate bias in AI systems that may unintentionally disadvantage certain tenant groups.

What It Covers:

- Use of representative or synthetic training data
- Protected characteristics testing (e.g. age, race, disability)
- Inclusion of fairness metrics during model evaluation
- Post-deployment monitoring for drift or discrimination

Source: [IBM AI Fairness 360 Toolkit](#), [Microsoft Fairlearn](#), [Together Housing Case Study](#)

How to Use It: Run this checklist before going live, especially for scoring tools (e.g. arrears risk, housing allocation). Maintain a bias audit log as part of your policy documentation.

5. Governance Committee Terms of Reference (TORs)

Purpose: Define the scope, responsibilities, and structure of your AI Steering and Operational Committees.

What It Covers:

- Committee membership (IT, DPO, Housing Ops, Legal, Tenant Reps)
- Roles and meeting cadence
- Escalation and decision-making responsibilities
- Alignment with wider risk and audit functions

Source: [Responsible AI Institute – AI Policy Template V1.0 \(2024\)](#)

How to Use It: Use as a basis for board approval when establishing oversight. Tailor to fit your governance structure (e.g. combine with ESG or Digital boards).

6. DPIA Template (AI-Specific)

Purpose: Assess and document data protection risks for any AI system processing personal or sensitive data.

What It Covers:

- Purpose and legal basis for processing
- Data flows and security measures
- Impact on tenant rights
- Mitigation strategies and residual risks

Source: [ICO – AI and Data Protection Risk Toolkit \(2023\)](#)

How to Use It: Complete for all AI projects involving tenant profiling, behavioural prediction, or decision-making. Link outcomes to your governance committee for approval.

7. Vendor Due Diligence Questionnaire

Purpose: Vet AI vendors thoroughly, before procurement, to ensure ethical, safe, and compliant products.

What It Covers:

- The company's AI governance and risk practices
- Model development process
- Bias testing and fairness assurance
- Data handling, security, and privacy safeguards
- Past ethical incidents or failures

Source: [Responsible AI Institute, ISO 42001 supplier evaluation guidance](#)

How to Use It: Include it as part of your vendor onboarding pack or procurement RFP. Score responses to help compare offers based on more than just cost or performance.

Where to Start?

If you're just beginning to develop your policy:

- **Start with the Risk Assessment Template** to prioritise use cases.
- Use the **TORs and AUP** to set up governance and staff expectations.
- Then move to **procurement, DPIA, and bias checklists** as you implement new tools.

If your policy already exists:

- Review these tools against your current process and identify gaps.
- Update your policy pack with these new templates to align with emerging best practices and standards.

Appendix 2 Resources

A practical reference to support your AI governance journey.

This appendix brings together the key resources, frameworks, and terminology that underpin your AI governance strategy. Whether you're drafting policies, evaluating tools, or briefing board members, these references ensure you're building on solid ground.

Key Reference Frameworks

Source	Purpose & Link
OECD AI Principles (2024 update)	Global ethical standards for trustworthy AI: fairness, transparency, human-centredness.
ISO/IEC 42001 (2023)	International standard for AI management systems — including risk, ethics, and transparency.
UK Information Commissioner's Office (ICO) – AI & Data Toolkit	Practical tools and legal guidance on AI, GDPR, and data protection risk.
EU AI Act (2024)	Comprehensive AI legislation focused on high-risk systems, rights, and governance duties.
NIST AI Risk Management Framework (RMF) (2023)	US-led framework for mapping, measuring, and mitigating AI risks.

Practical Governance Tools (Reviewed in This Guide)

Tool	Publisher & Access
Responsible AI Policy Template V1.0	Sets out governance roles, policy structure, risk controls, and committee TORs.
UC Berkeley Gen AI Playbook	10 practical "plays" for safe, ethical GenAI use across teams.
ISACA GenAI AUP Guidance	Structured approach for writing Acceptable Use Policies and staff guidance.

Tool	Publisher & Access
ISACA AI Governance Brief	The report provides a practical audit tool for organisations to assess and strengthen their AI governance systems.

Further Reading & Sector Insights

- [Centre for Data Ethics and Innovation – AI Assurance Roadmap](#)
- [World Economic Forum – AI Procurement in a Box Toolkit](#)
- [Camden Council – Data Charter \(Transparency Best Practice\)](#)
- [Alan Turing Institute – Explaining Decisions Made with AI](#)

Glossary of Key AI Terms for Executives

Term	Definition
AI Governance	The structures, principles, and processes used to oversee the responsible use of AI.
DPIA (Data Protection Impact Assessment)	A legal requirement under GDPR to assess privacy risks in data-driven projects, especially those using AI.
Automated Decision-Making	When an AI system decides without meaningful human involvement (GDPR Article 22 regulates this).
Bias Mitigation	The process of identifying and correcting unfair outcomes in AI systems, especially those affecting protected groups.
Explainability	The degree to which the outputs of an AI system can be understood by humans is essential for accountability.
Black Box Model	An AI system whose internal workings are not interpretable makes it harder to explain decisions.
Generative AI (GenAI)	AI models that can produce content, such as text, images, or code, are often based on large language models.
Fairness Audit	A test to identify potential discrimination or bias in AI systems before and after deployment.
Model Fact Sheet	A document summarising an AI model's purpose, data sources, risks, performance, and safeguards — used to promote transparency.
High-Risk AI System	A term used in the EU AI Act to describe systems that impact people's rights or access to services, requiring stricter governance.

About Demystifying AI for Social Housing (DASH)

At DASH, we make AI simple and practical for social housing leaders. We know AI can seem complex, so we focus on clear, actionable insights that help you drive efficiency and make informed decisions, whether in day-to-day operations or long-term strategy. For **board members**, our AI tools enhance governance by streamlining routine tasks, allowing you to focus on impact. For **executives**, we provide the insights needed to drive innovation and stay ahead of change.

Supported by our panel of AI experts with 15+ years of sector expertise and partnerships with AI specialists, we turn complex technology into real-world results. More than just a tech provider, we're your AI co-pilot, helping you harness AI's potential to improve services, strengthen decision-making, and shape the future of social housing.

Team DASH

peter@dashai.co.uk

<https://dash-ai.co.uk/>