ARTIFICIAL INTELLIGENCE (AI) POLICY

DOCUMENTATION MASTER SHEET

Amendments to this Document are Detailed Below.

call 01452 424344
email customer.experience@gch.co.uk
visit www.gch.co.uk

Pride. Quality. Integrity. Innovation.

Version Number	Date Amended	Comments	Date Approved	Author	Approved By
01		New policy	1/4/2026	Paul Haines	SLT
02					
03					
04					
05					
06					
07					
08					
09					
10					

Pride. Quality. Integrity. Innovation.

Business Information provided by Gloucester City Homes

OVERALL AIM/INTRODUCTION

This policy aims to ensure the responsible, ethical, and transparent use of artificial intelligence (AI) technologies in delivering housing services, supporting tenants, and managing operations. We aim to support better, more efficient services, not replace personalised human services.

SCOPE OF POLICY

This policy applies to all staff, contractors, Board and Committee members, and third-party providers using, or considering, AI tools or systems on behalf of GCH, as well to the AI tools themselves.

LEGAL AND REGULATORY EXPECTATIONS

1. Social Housing (Regulation) Act 2023
   - Transparency & Accountability: AI systems used in customer communications or decision-making must support clear, auditable processes.
   - Tenant Voice: AI must not replace meaningful customer engagement. Automated systems must allow for human escalation.
   - Performance Monitoring: AI tools used in repairs or arrears must align with new consumer standards and be subject to scrutiny by the Regulator of Social Housing

2. UK GDPR and Data Protection Act 2018
   - Lawful Basis for Processing: AI systems must have a clear legal basis for using personal data.
   - Automated Decision-Making: If AI makes decisions without human input, customers have the right to explanation and challenge.
   - Data Minimisation & Security: Only necessary personal data should be used, and it must be securely stored and processed.
   - Special Category Data: Under the UK GDPR, "sensitive data" is defined as special category data, which requires more protection and includes personal data revealing an individual's race, ethnic origin, political opinions, religious beliefs, trade union membership, genetic data, biometric data (for identification), health data, sex life, or sexual orientation.

3. Equality Act 2010
   - AI must not discriminate: The rules that prevent discrimination by humans and organisations apply equally to AI systems.

GUIDING PRINCIPLES

- Fairness: AI must not discriminate against customers or colleagues based on protected characteristics.
- Transparency: Decisions made or supported by AI must be explainable and documented.
- Accountability: Human oversight is required for all critical decisions.
- Privacy: AI systems must comply with the UK GDPR and Data Protection Act 2018 if they have access to personal data.
- Security: AI systems must meet organisational cybersecurity standards.

USE OF LARGE LANGUAGE MODELS (LLMs)

LLM AI tools include CoPilot, ChatGPT, Gemini, and many others. These tools are designed to understand natural language and generate text, pictures and video based on patterns learnt from vast amounts of data.

Colleagues are encouraged to use these tools to improve the delivery of services to both internal and external customers, but this must be done thoughtfully, with an understanding of the potential risks and flaws.

1. Data Security: LLMs continue to train themselves on new data. For most models this includes any data that you provide to them. So, anything that you give to these models can and will become part of its 'knowledge' and can then be freely given out to anyone who asks the right question. If you wouldn't share the data with a stranger on the Internet, don't share it with an LLM.
2. Hallucinations: Despite being highly advanced systems, LLMs are not always right. There is no guarantee that the information they provide, conclusions that they draw, or suggestions that they give are based on sound data and reasoning. Whilst extremely useful as a 'thought partner', their output should not be relied upon without applying your own critical thinking and fact checking.
3. Bias: LLMs can only be trained on the data available to them. This data will include opinions and conjecture and is highly likely to under-represent specific groups. This can lead to bias, where the answers or suggestions provided by the model are based on information that isn't fact, or does not fully take account of all groups of people. Think critically…
    a. Does the response truly reflect our customer demographics?
    b. Does it account for all cultures?
    c. Whilst each individual response looks fine, does a bias appear when you look at a wider series of responses?

Permitted Use

- Drafting non-sensitive content such as general communications, reports, pictures, videos, and policy templates.
- Supporting colleagues with research and summarisation tasks where data is not personal or sensitive*.
- Analysing data or generating test data where that data is not high risk*.
- Generating code, excel formulas, or other complex compute instructions provided they are well tested before deployment.
- *If you are working with data, please see the section below 'which LLM should I use?'.

Prohibited Use

- Entering information that is personal, confidential, sensitive, or any information that could identify a person into LLMs.
- Using LLM outputs for decisions that directly affect individuals without human review.
- Deploying LLMs without following GCH procurement rules – particularly those concerning data privacy and security.

Which LLM Should I Use?

Each model has their own strengths, weaknesses and personality – some deliver better results than others for any specific task, and this is often down to personal preference. To understand which LLM you should use, you must first assess the data that you will be providing – either by uploading documents or typing in information...

High Risk Data

Applies to any data defined as "special category" (sensitive) by law or any other information that a person might not even tell their friends. This includes any data about...

- Racial or ethnic origin
- Political opinions
- Religious or philosophical beliefs
- Trade union membership
- Genetic data
- Biometric data
- Health data
- Sex life or sexual orientation
- Criminal offence data
- Whilst not legally defined as sensitive, financial data should also be treated as high risk.

Examples might include an EDI data table, customer or colleague salary or bank account information, or a list of trade union members.

Medium Risk Data
Applies to any data that could identify a living person (personal data), or data that is commercially sensitive. Examples might include a tenancy agreement that contains names or addresses, a list of employees, or a suppliers pricing matrix.

Zero Risk Data
Applies if you are not sharing any data, or you are only sharing data that is not personal, sensitive, or commercially sensitive. Examples might include a tenancy agreement template with no names or addresses, a job description, or a table of repairs SoR codes.

Do not share High Risk Data with any LLM. If you feel that there is a specific use case or reason to do so, you must talk to your line manager or a data protection specialist (in the Governance Team) to seek approval first.

Medium Risk Data can only be shared with a Microsoft CoPilot account linked to our corporate Microsoft subscription. Doing so ensures that the data never leaves the boundaries of our Microsoft estate.

When dealing with Zero Risk Data, you may use whichever LLM delivers the best results for you. Colleagues are encouraged to experiment, collaborate, and research to find the best solution.

USE OF SYSTEMS WITH SPECIALISED AI (INCLUDING AGENTS)

As well as LLMs, specialised AI can be found within other systems. These systems are often used to determine insights and actions from large datasets, to implement specific actions as part of a business process, to handle, triage, or summarise customer contacts, or to summarise meetings, amongst other things.

Any specialised AI utilised by GCH must meet the guiding principles of this policy. New systems including specialised AI must be procured within GCH's procurement rules – with particular attention paid to data privacy and security. Should such functionality be introduced without the need for procurement (for example as a feature update to a system already in place), data privacy and security should be given the same scrutiny as if it were the procurement of a new system.

Changes to AI capable systems (such as a new agent or change to an existing agent) should also be measured against the guiding principles set out in this policy, to ensure that those changes do not introduce unwanted flaws.

REVIEW

This policy will normally be reviewed annually to ensure that it reminds fit for purpose, unless an earlier review is required due to regulatory, statutory, or other requirements, including significant technical advancements.

BREACH AND ENFORCEMENT

Non-compliance with this policy may result in disciplinary action and/or termination of contracts. Serious breaches will be reported to the ICO and relevant regulators.