Generative Artificial Intelligence (GenAI) Policy

Implementation date:

Control schedule

Approved by	Policy and Sustainability Committee
Approval date	March 2026
Senior Responsible Officer	Dr Deborah Smart – Corporate Director of Customer and Corporate Services
Author	Mike Brown – Cyber Security Manager
Scheduled for review	Quarterly for 1st year, then yearly thereafter

Version control

Version	Date	Author	Comment
0.1	September 2024	Mike Brown	Initial Draft
0.2	October 2024	Heather Robb	Minor updates.
0.3	November 2024	Kevin Wilbraham	Minor Updates
0.4	November 2024	Mike Bell	Minor Updates
0.5	November 2024	Alison Roarty	Minor Updates
0.6	December 2024	James McConnell	Minor Updates
1.0	March 2025	Mike Brown	Final version
1.1	October 2025	Mike Brown	Ethics Updates
1.2	October 2025	Mike Brown, Heather Robb and Kevin Wilbraham	Additional updates around Governance and Ethics
1.3	December 2025	Mike Brown	Updates on Definitions, Training and Sustainability
1.4	December 2025	Mike Brown	Minor updates to reflect comments made by technology partners
1.5	January 2026	Mike Brown	Minor updates to outline operational efficiencies, AI Register / Governance. Procurement and Sustainability
1.6	January 2026	Mike Brown, Nicola Harvey	Additional Information added in relation to Disclosure

EDINBVRGH
THE CITY OF EDINBURGH COUNCIL

1.7	February 2026	Mike Brown	Minor changes to definitions for technology and personnel
1.8	April 2026	Mike Brown, Heather Robb & Kevin Wilbraham	Updates from P&S referral report
1.9	June 2026	Mike Brown, Heather Robb & Kevin Wilbraham	Quarterly Review – Minor Updates
2.0	September 2026	Mike Brown, Heather Robb & Kevin Wilbraham	Quarterly Review – Minor Updates

Committee decisions affecting this policy

Date	Committee	Link to report	Link to minute

Internal

Purpose and Background

1.1 The purpose of this policy document is to establish a clear framework for the use of Generative Artificial Intelligence Large Language Model tools (collectively referred to as GenAI), such as M365 Co-Pilot, ChatGPT, Gemini, Bing, and other similar tools, by Council employees, Elected Members, contractors, developers, vendors, temporary staff, consultants, and other third parties, hereinafter referred to as 'Council personnel'.

1.2 This policy aims to ensure that the use of GenAI in the Council is ethical, complies with all applicable laws, regulations, and Council policies, and complements the Council's existing information and security policies. Given the rapid development and application of GenAI, this policy will continuously evolve.

1.3 Generative Artificial Intelligence (Gen AI) will be leveraged to deliver operational efficiencies for the Council by automating routine tasks, accelerating analysis, and supporting faster decision-making across the Council. Its use is intended to enhance productivity, improve the consistency and quality of outputs, and enable Council personnel to focus on higher-value, strategic activities, while allowing frontline professionals to spend more time with the citizens most in need.

1.4 Generative artificial intelligence (GenAI) can generate realistic, human-like text, images, code, and art based on extensive data (usually public) it has been trained on. Generative AI:

- Can produce a range of useful outputs, such as text, audio, images, and code.
- Responds to natural language questions, making it accessible for any employee.
- Can handle or use different types of data, which is beneficial for the Council managing large amounts of unstructured data in various formats.

1.5 The UK's National AI Strategy recognises that AI is the fastest growing deep technology in the world, with huge potential to rewrite the rules of entire industries, drive substantial economic growth and transform all areas of life. It aims to:

- Invest and plan for the long-term needs of the AI ecosystem to continue our leadership as a science and AI superpower.
- Support the transition to an AI-enabled economy, capturing the benefits of innovation in the UK, and ensuring AI benefits all sectors and regions.
- Ensure the UK gets the national and international governance of AI technologies right, to encourage innovation, investment, and protect the public and our fundamental values.

Internal

1.6 The Scottish AI Strategy states that Scotland should be a leader in AI technologies and have identified the following actions to achieve this. It aims to:

- Be the collective leadership of the Scottish AI alliance
- Create foundations for success
- Build an AI powerhouse

1.7 GenAI systems can require significant energy consumption, contributing to carbon emissions. The growth in the use of AI requires more data centres which use significant amounts of energy, water, rare earth metals and land.

Scope

2.1 The Policy will apply to all Council employees, Elected Members, contractors, vendors, and third parties using GenAI for Council business. Use in schools by pupils will be covered separately.

2.2 The purpose of this policy is to provide a clear framework to be applied by the Council which governs the utilisation of AI technologies.

Definitions

3.1 Artificial intelligence (AI) describes computer systems which can perform tasks usually requiring human intelligence. This could include visual perception, speech recognition or translation between languages. Below is an overview of key AI technologies

Generative AI (GenAI): Generative AI refers to systems that can create content, such as text, images, or even software code, based on patterns learned from training data.

Basic Machine Learning (ML): Machine learning involves algorithms that allow systems to learn from data and improve over time without being explicitly programmed.

Robotic Process Automation (RPA): RPA involves automating routine, rule-based tasks typically performed by humans.

Natural Language Processing (NLP): NLP is a branch of AI that enables machines to understand, interpret, and generate human language.

Agentic AI: Agentic AI refers to autonomous systems capable of independently planning, deciding, and taking actions to achieve defined goals with minimal human intervention

Computer Vision (CV): This technology allows computers to interpret and make decisions based on visual data, such as images or videos.

Large Language Models (LLMs): LLMs, like those underlying ChatGPT, are advanced AI models that can understand and generate human-like text. LLM's are a subset of NLP and Gen AI

Internal

Use

4.1 This policy applies to all Council personnel using any GenAI tools, whether through Council-owned devices or personal devices used for Council activities. These tools can be embedded in other tools – such as email clients or video conferencing tools. For example, Microsoft 365 includes many GenAI tools – such as Teams transcription and M365 Co-Pilot.

4.2 The use of GenAI and the systems that support these should foster fairness reduce the risk of bias to prevent discrimination, ensure equal treatment, and align positively with the Council's policies, goals and values.

4.3 Subject to the necessary approvals detailed in section 5.1, Council personnel may utilise GenAI for work-related tasks, provided they comply with this policy. This encompasses, but is not limited to, the generation of text or content for reports, emails, presentations, images, and customer service communications, whilst continuing to adhere to the Council's Acceptable Use Policy.

4.4 Special consideration must be given to governance, vendor practices, copyright, data storage/use, accuracy, bias, quality, security, fraud, crime, confidentiality, disclosure, and integration with other tools.

Governance

5.1 To ensure robust governance of AI use, all AI proposals will undergo a data ethics assessment to ensure ethical considerations are documented and understood. Proposals will be triaged by Information Governance before initial approval or rejection by the Data Ethics Board, which provides oversight, challenge, and assurance that AI technologies are used ethically and responsibly.

The Data Ethics Board, chaired by an independent external member, will include Council officers, external organisations, to ensure diverse expertise and perspectives.

Following initial agreement to proceed by the ethics board, AI projects will follow existing governance processes, including data protection impact assessments, security assurance statements, and integrated impact assessments, to evaluate impacts and mitigate security or privacy risks.

After these assessments, the project will return to the Data Ethics Board for a final decision, which will inform a recommendation to the relevant Information Asset Owner (Service Director) on whether to proceed.

Following approval and implementation, all AI systems will undergo regular review, at least annually by the Data Ethics Board. Regular updates will be provided to the Information Board and Digital Governance Board. Information Governance will oversee the maintenance and publication of an approved register of AI tools authorised for use within the Council, supporting transparency around AI adoption. The board will also be responsible for addressing complaints or concerns arising from the use of Gen AI within the Council

Internal

5.2 Deployment of AI technologies across the Council must use the following guidance and guidelines:

NCSC AI and cyber security: what you need to know
NCSC Secure AI system development
ICO Artificial intelligence

Training

6.1 All Council personnel will be provided with e-training covering the key elements of the Council's approved AI Policy.

This training will support Council personnel in understanding their responsibilities, appropriate use of AI, data protection considerations, and the operational controls that apply to AI tools. The training will also raise awareness of climate and ethical considerations regarding the use of AI.

6.2 As part of the Council-wide deployment of Generative AI technologies, including tools such as Microsoft Copilot, key information, usage guidelines, training materials, and instructional videos will be made available on the Council's ORB. Council personnel are expected to access these materials to ensure they use GenAI tools effectively and in compliance with Council policy and governance requirements.

6.3 Where departments implement AI systems specific to their operational needs, responsibility for ensuring appropriate Council personnel training rests with the relevant system owners. System owners must ensure that all Council personnel using departmental AI tools are trained in their correct use, understand the approved controls and guardrails, and operate these systems in line with both departmental procedures and the Council's overarching AI Policy

Vendors

7.1 Any use of GenAI technology for Council activities must be carried out with a complete understanding of the policies, practices, impact of its use and terms & conditions set by the tool's developers or vendors. All AI purchases must be managed through CGI or the Council's procurement team.

7.2 Prior to any utilisation of AI through existing suppliers or the procurement of new technology managed via the Council's Procurement Team or CGI Change Process, approval must be obtained from the Data Ethics Board, as outlined in section 5.1 of this policy.

Internal

Copyright and Licencing

8.1 Council personnel must adhere to copyright laws when utilising GenAI. It is prohibited to use GenAI to generate content that infringes upon the intellectual property rights of others, including but not limited to copyrighted material. If a user is unsure whether a particular use of GenAI constitutes copyright infringement, they should contact Legal Services before using GenAI. For example, using GenAI to produce a logo could produce something based on a copy of a logo that is a trademark or is copyrighted.

8.2 Additionally, Council personnel must ensure that they comply with the licensing terms of the GenAI tools and platforms they are using. This includes verifying whether the generated content can be used commercially, modified, or distributed, depending on the licensing restrictions of the AI service provider. Advice should be sought from Legal Services and/or procurement.

8.3 The use of GenAI tools is explicitly prohibited in the creation, development, or production of any creative, artistic, or cultural outputs.

Accuracy

9.1 GenAI can produce text that appears factual but may not be accurate and may generate AI hallucinations, outputs that seem credible or authoritative but are fabricated, misleading, or incorrect. All final decisions made within the Council will remain the responsibility of humans (Council officers) and cannot be delegated to Artificial Intelligence (AI) systems. AI may be used to support analysis, insight, or recommendation; however, human judgement and accountability will underpin all decision-making processes.

The use of AI in any decision, support capacity will be subject to the necessary assurances, including ethical review, data protection compliance, and security and governance approvals, to ensure that outcomes are transparent, fair, and aligned with the Council's statutory and ethical obligations.

9.2 The use of AI introduces risks, including the potential for errors, bias, or lack of transparency. These risks must be actively managed, therefore all AI-generated content must be reviewed and verified for accuracy before use.

Council officers are accountable for confirming the accuracy and appropriateness of AI outputs before relying on, sharing, releasing or publishing them. Where there is any uncertainty or risk of errors, bias, or lack of transparency, the information must not be used unless it has been corrected and approved.

Internal

Confidentiality

10.1 Confidential and personal data must not be entered into public GenAI tools (like ChatGPT). This is because such information would then be accessible in the public domain and potentially used for further training of the tool, leading to a data breach, and resulting in potential censorship and fines from regulatory bodies such as the ICO. Council personnel must adhere to all applicable data privacy laws and organisational policies when using GenAI. For example:

- Council personnel must **NOT** use unauthorised GenAI tools to draft letters to customers containing personal details (e.g., 'Mr. A N Other at 123 Princes Street') as the data will be ingested and stored by the GenAI for reuse.
- Council personnel must **NOT** use GenAI apps on personal phones to record and summarise work meetings or to use translation services.
- Council personnel must **NOT** upload spreadsheets containing customer data for GenAI analysis.

10.2 If users are uncertain about the confidentiality of the information or how the data they input will be handled, they must refrain from using that GenAI tool.

10.3 Confidential or personal data should only be entered into a GenAI tool if specifically approved, developed or procured for City of Edinburgh Council use, where the data entered is restricted to the Council's sole use and the tool has been expressly approved for that purpose by the Information Governance and Council Cyber Security Teams. For instance, using Microsoft Teams with a Council login to transcribe meetings is authorised subject to appropriate controls and processes, including those detailed in section 5.1 of this policy. However, using a free tool downloaded to a personal phone for transcribing a work meeting is not authorised and would constitute a data breach.

10.4 All data breaches, technical breaches, and operational incidents, including data entered into the wrong tool or generation of harmful outputs, must be reported immediately using only the appropriate channels on the ORB.

Social Impact and Equality

11.1 Council personnel must be aware of how the use of GenAI may impact different groups of people in different ways as it may have inherent social or political bias or have been trained on stereotypes. It may have inappropriate cultural values or display insensitive content. For example, GenAI must not be allowed to solely determine which customers should have access to services; Humans must be involved in such decision-making where needed, and there must be an appeal processes for any automated or AI-informed decisions. This process will be undertaken by the Digital Services team in consultation with the Information Governance team. If there are any doubts about the appropriateness of using GenAI in a particular situation, Council personnel must consult with their line manager, Information Governance and the Cyber Security team

Internal

11.2 Data protection legislation also gives individuals certain rights in relation to how their personal data is being used or created, particularly for automated decision-making. As such, the use of AI carries inherent risks that must be considered through the data protection impact assessment process.

Ethical Use

12.1 In line with the Alan Turing Institute guidance on AI ethics published on the UK government website, the Council is committed to responsible governance, deployment, and oversight of Artificial Intelligence (AI) systems, ensuring that their use remains fully aligned with our statutory obligations, ethical standards, and the Council's duty to act in the public interest

AI will be guided by the following values

- **Respect for individuals** - upholding dignity, autonomy, rights, and privacy.
- **Inclusivity and openness** - engaging stakeholders in decision making.
- **Wellbeing** - maximising public benefit and minimising potential harm.
- **Justice and fairness** - promoting equality and supporting the public interest.

In applying AI, the Council will adhere to the following principles:

- **Fairness** - reducing unjust bias and ensuring equitable outcomes.
- **Accountability** - maintaining clear responsibility and audit trails.
- **Reliability** - ensuring AI systems are robust and safe.
- **Transparency** - providing clear explanations and openness about AI use
- **Environmental Sustainability** – ensuring use is responsible and efficient

These commitments will be embedded through governance processes, including ethical governance boards, stakeholder engagement, continuous monitoring, and transparent reporting.

12.2 GenAI must be used ethically and in compliance with all applicable legislation, regulations and Council policies. Council personnel must not use GenAI to generate content that is discriminatory, offensive, or inappropriate. If there are any doubts about the appropriateness of using GenAI in a particular situation, Council personnel must consult with their line manager, Information Governance and the Cyber Security team.

Disclosure

13.1 Content created using GenAI must be clearly labelled and disclosed as containing GenAI-generated information.

Footnote example:

Internal

"Note: This document includes content generated by Artificial Intelligence (AI). The AI-generated content has been reviewed and edited for accuracy by the author. The author takes responsibility for this content".

13.2 Privacy information will also need to be updated or developed to ensure absolute transparency around how personal data will be processed in the AI system. This will need to include (at a minimum) details on your purposes for processing an individual's personal data; retention periods for that data; and who you will be sharing it with.

13.3 Where automated or Generative AI systems are used to generate responses, provide information, or interact with users, the use of such systems must, where practicable, be clearly and prominently disclosed. Users must be advised that they are interacting with an AI system and not a Council officer.

Integration(s)

14.1 API and plugin tools facilitate access to GenAI and enhance functionality for various services (such as email, Teams, or search engines) to boost automation and productivity. Council personnel must adhere to Open AI Safety Best Practices,

API and plugin tools must undergo thorough testing for:

- Moderation – ensuring the model appropriately handles hate, discriminatory, threatening, etc., inputs.
- Factual responses – providing a ground truth for the API and reviewing responses accordingly.
- Storage retention and capabilities.

Risks

15.1 Use of GenAI carries inherent risks including the categories detailed. Any risks associated with the use or adoption of Artificial Intelligence (AI) will be managed through the governance processes outlined in Section 5.1 of this policy.

- **Climate:** extensive use of Gen AI would contribute to the significant levels of damaging greenhouse gas emissions, water and rare mineral use from the data centres required to generate answers and content for Gen AI queries.

- **Ethics:** ensuring that the technologies currently being developed are used for the common good, rather than for the benefit of a select few.

  Bias: "AI systems may reflect, reproduce or amplify societal biases, potentially leading to unfair or discriminatory outcomes. These biases may adversely affect individuals with protected characteristics under the Equality Act 2010, including age, disability, gender reassignment, marriage and civil partnership, pregnancy and maternity, race, religion or belief, sex, and sexual orientation. Staff must consider and, where appropriate, mitigate

Internal

the risk of bias and discrimination when using or implementing AI solutions, helping to ensure fair and equitable outcomes”

- **Privacy:** The council remain committed to safeguarding personal data. The use of Generative AI presents, but not limited to the following risks

    - Data privacy and leakage
    - Data retention
    - Lack of data control
    - Data misuse

- **Reputational Risk** – Risk to Council’s reputation where decisions, advice, or services rely on AI-generated data that is inaccurate, misleading, or unverified, potentially leading to reputational damage, poor decision-making, and loss of public trust

- **Copyright infringement:** Copyright infringement can happen when a GenAI system produces content that uses someone else’s protected work without permission. For example, text, images, music, or code. This can occur if the system creates something too similar to existing copyrighted material or reproduces parts of it directly. The risk is Council personnel may unintentionally publish or rely on AI-generated content that they don’t legally have the right to use. To avoid this, users should review outputs carefully, avoid asking AI to copy specific copyrighted works, and ensure that any generated material is original or appropriately licensed.

- **Licensing:** The use of AI in business and innovation introduces potential licensing risks that the Council must carefully manage. These risks primarily stem from the use of proprietary datasets, third-party software, and pre-trained AI models, which may have complex licensing terms. For example, many AI tools are developed using open-source libraries that come with specific restrictions on use, modification, and distribution. Failure to comply with these terms could result in legal disputes, penalties, or the loss of rights to use critical technologies.

- **Secondary mining of the Metadata:** Extensible metadata that adds custom data labelling is common in object storage and available in some distributed filesystems. The metadata can be used to track data origin, add labels and even tag data used for different AI models.

- **Consistency of output:** In the context of AI, reproducibility refers to the ability to achieve the same or similar results using the same dataset and AI algorithm within the same environment.

- **Data Protection (UK GDPR):** Giving the right to query automated decisions. The Guidance on AI and Data Protection has been updated after requests from UK industry to clarify requirements for fairness in AI. It also delivers on a key ICO25 commitment, which is to help organisations adopt new technologies while protecting people and vulnerable groups. The UK GDPR has provisions on:

Internal

- Automated individual decision-making (decision made solely by automated means without any human involvement).
- Profiling (automated processing of personal data to evaluate certain things about an individual). Profiling can be part of an automated decision-making process.
- Article 22 of the UK GDPR has additional rules to protect individuals if you are carrying out solely automated decision-making that has legal or similarly significant effects on them.

- **Legal implications**: Vicarious responsibility for Chatbot generated answers and advice.

A comprehensive risk assessment must be conducted for any project or process where use of GenAI is proposed via a data protection impact assessment and Security Assurance Statement (SAS).

The risk assessments consider potential impacts including legal compliance; bias and discrimination; security (including technical protections and security certifications); and data sovereignty and protection.

GenAI may store sensitive data and information, which could be at risk of being breached or hacked. The Council must assess technical protections and security certification of a GenAI tool before use. If Council personnel have any doubt about the security of information input into GenAI, they must not use GenAI.

Legal Compliance & Data Sovereignty and Protection

16.1 Data entered into GenAI may become publicly accessible, potentially disclosing non-public information and violating regulatory requirements, customer or vendor contracts, or compromising intellectual property. Unauthorised release of private or personal data can lead to breaches of relevant data protection laws. Additionally, using GenAI to generate content may infringe upon intellectual property protection regulations. Council personnel must ensure their use of GenAI adheres to all applicable laws, regulations, and Council policies.

16.2 While a GenAI platform may be hosted internationally, under data sovereignty rules information created or collected in the originating country will remain under jurisdiction of that country's laws. The reverse also applies. If information is sourced from GenAI hosted overseas, the laws of the source country regarding its use and access may apply. GenAI service providers must be assessed for data sovereignty practice by any Council personnel wishing to use their GenAI.

Sustainability

17.1 The Council is committed to minimising the environmental impact of its operations, including the use of Generative AI (GenAI), in support of its objective to achieve net zero greenhouse gas emissions by 2030. GenAI systems, through the data centres that

Internal

power them, contribute to greenhouse gas emission levels and environmental impact through their consumption of significant energy, water and rare minerals.

The Council will therefore ensure that the use of GenAI is proportionate, responsible, and aligned with environmental sustainability principles, and is adopted only where it delivers clear organisational or public value.

Council personnel will be encouraged to use GenAI responsibly and efficiently, avoiding unnecessary tasks, duplication, and excessive data processing or storage. Environmental considerations will inform decision-making relating to the selection, procurement, use, and lifecycle management of AI-enabled and cloud-based services.

17.2 Some major cloud and AI providers are investing in renewable energy, energy-efficient data centres, advanced cooling technologies, and energy-saving AI hardware. In procuring AI infrastructure or software, supplier sustainability claims will be considered alongside service requirements, best value considerations, transparency, and demonstrable progress in reducing environmental impact.

The Council will aim to adopt the following where possible and practical:

Use sustainable cloud services:

The Council will prioritise cloud and AI providers that can demonstrate progress in reducing carbon emissions, improving data-centre efficiency, responsible water management, and increasing the use of renewable energy, and whose sustainability commitments align with the Council's.

Choose low-emission locations and services:
Where feasible and consistent with service needs, the Council will select cloud regions, services, and deployment options with lower carbon intensity and use available tools to optimise workload placement from an environmental perspective.

Use energy-efficient AI tools and infrastructure:
Where feasible and consistent with service needs, the Council will select AI models, platforms, hardware, and services that minimise energy consumption while maintaining the performance required to deliver high-quality and reliable services.

The environmental impact of GenAI and related digital services will be kept under review as part of existing governance and assurance processes, informed by available data, supplier transparency, and evolving best practice.

This approach will be reviewed periodically to reflect technological advances, regulatory developments, and improvements in sustainable AI practices

Roles, Responsibilities and Review

18.1 This Policy has been developed by Customer and Digital Services, in collaboration with the Council's Information Governance Team, and with contributions from members of the Council's Cyber and Information Security Steering Group.

18.2 The Policy will be reviewed by the Data Ethics Board quarterly in the first year, yearly thereafter and more frequently if needed, particularly where there are changes to legislation, best practice, guidance from specialist organisations such as the National Cyber Security Centre (NCSC) and the Information Commissioner's Office (ICO), or where the way the technology operates changes significantly.

Internal

18.3 The Information Asset Owner (IAO) is responsible for ensuring that all data under their control is accurate, current, and properly maintained, including information recorded on Council Systems, prior to the deployment or utilisation of any AI solution.

Related documents

19.1 NCSC AI and cyber security: what you need to know
NCSC Secure AI system development
ICO Artificial intelligence
ICO25
National AI Strategy
Scottish AI Strategy
Open AI Safety Best Practices
Understanding artificial intelligence ethics and safety - GOV.UK
Green Industrial Strategy

Integrated Impact Assessment

20.1 Integrated Impact Assessment completed

Internal