Mae'r ddogfen hon hefyd ar gael yn Gymraeg

This document is also available in Welsh

GIG CYMRU NHS WALES | Iechyd Cyhoeddus Cymru Public Health Wales	**Reference Number:** PHW-SCD20 **Version Number:** 1 **Date of Next review:** February 2029

GENERATIVE ARTIFICAL INTELLIGENCE (AI) GUIDANCE

Introduction and Aim

This document provides guidance on the safe and effective use of generative AI tools within Public Health Wales.

The guidance seeks to:

- Provide information on approved generative AI tools within PHW
- Provide guidance on the safe and effective use of those tools
- The process for access generative AI tools

This procedure ensures that PHW:

- Has a clear process for the use of Generative AI
- Clearly defines the principals & expectations in the use of generative AI products
- Clearly outlines the process for requesting alternative generative AI tools

Linked Policies, Procedures and Written Control Documents

All corporate policies and procedures are available on the Public Health Wales website

- Information Security Policy
- Information Governance Policy
- Acceptable use Policies

Scope

This guidance is relevant to:

- The workforce of Public Health Wales, including staff, Independent (Non-Executive) Board Members, students, trainees, secondees, volunteers, contracted third parties and any other persons undertaking duties on behalf of the organisation where the use of generative AI tools are required.

Equality and Health Impact Assessment	An Equality, Welsh Language and Health Impact Assessment has been completed and can be viewed on the policy webpages.
Approved by	Leadership Team
Approval Date	19/02/2026
Review Date	19/02/2029
Date of Publication:	29/04/2026
Accountable Executive Director/Director	Director of Research, Data and Digital and Senior Information Risk Owner
Author	Lead Cyber Security Manager

Summary of reviews/amendments				
Version number	**Date of Review**	**Date of Approval**	**Date published**	**Summary of Amendments**
1		19/02/26	29/04/26	New

Contents

1. Purpose

This document outlines the guidance for the use of Artificial Intelligence (AI) tools within PHW, including the following:

- PHW's position on the use of Generative AI (GenAI) tools by staff
- Guidance on the safe and effective use of GenAI
- The process for accessing Gen AI tools

2. Background

UK Government guidance describes Generative Artificial Intelligence (GenAI) as 'a broad label used to describe any type of artificial intelligence (AI) that can be used to create new text, images, video, audio, or code. Large Language Models (LLMs) are part of this category of AI and produce text outputs.'

Gen AI has become a transformative force across various sectors, including the public sector. The ability of AI to analyse vast amounts of data, generate insights, and automate tasks presents significant opportunities for enhancing efficiency, innovation, and service delivery.

Generative AI has shown promise in areas such as content creation, data analysis, and decision support. However, the use of AI also brings challenges related to data privacy, security, and ethical considerations. Balancing innovation with safety and compliance is crucial for leveraging AI's potential while safeguarding organisational integrity and public trust.

This guidance intends to balance the need to mitigate the specific risks of Generative AI with a strong desire not to inhibit innovation and increased productivity. To benefit, safely, legally and ethically from the use of Generative AI tools, staff need to be aware not just of the guidance, but of the underlying risks that require mitigation through following that guidance.

3. What is Generative AI?

Generative AI (GenAI) is a type of artificial intelligence that creates new content — such as text, images, code, or audio — by learning patterns from large datasets. It works by predicting what comes next in a sequence, allowing it to generate human-like responses, summaries, designs, and more.

GenAI can be used in a wide range of applications, including drafting documents, answering questions, generating creative content, assisting with coding, supporting customer service, and even helping with clinical documentation or research in healthcare.

It’s a powerful tool for enhancing productivity, creativity, and decision-making — when used responsibly.

What GenAI is:

- It creates content by identifying and combining patterns from its training data.
- It can answer the same question in different ways, depending on how it's asked.
- It’s a support tool, not a replacement for human expertise.

What GenAI isn’t:

- It doesn’t “understand” like a human — it predicts likely responses, not facts.
- Its training data may include biased or unverified content — so you must verify important outputs.
- It should not be used as the sole decision-maker, especially in critical or regulated environments.

Key Considerations:

- Bias: Be vigilant — GenAI can reflect or amplify biases in its training data.
- Data Privacy: Know where your data and questions go. Some platforms may store inputs or use them for training; others may not.
- Human Oversight: Always keep a human in the loop. Whether the decision is critical or not, professional and reputational accountability requires review and judgment.

What you should watch out for:

- Sensitive or confidential info: Avoid sharing personal, proprietary, or sensitive data unless you’re sure it’s safe.
- Verification: Use GenAI as a helpful assistant, not a final authority. Cross-check anything important.

Additional guidance on understanding Generative AI from the “Centre for Digital Public Services” (CDPS) is located here.

4. PHW’s Supported Tools

PHW currently supports and promotes the use of generative AI tools that have been vetted and approved by the organisation.

These tools are selected based on:

- Compliance with data protection regulations (e.g., GDPR)
- Adherence to security and privacy standards
- Seamless integration with existing systems

- Staff are encouraged to use these tools to enhance productivity and innovation, provided they follow the guidelines outlined in this document.

4.1. Approved AI Tools

The following generative AI tools are approved for use within the organisation:

- Microsoft Copilot (Enterprise Version)
- ChatGPT (Team Version – Currently in Pilot)
  - ChatGPT is in a controlled Pilot (due to end November 2025, after which controls, guidance, governance and restrictions will be embedded in its use)
- Google Gemini (Enterprise Version –Currently available within the NHS Wales Google Cloud Platform – for use with NDR, NDAP)

These tools are deployed in controlled environments where:

- Data retention policies are managed/ enforced by the organisation. Data entered is stored for a set amount of time as agreed by the organization and any inputs are not used to train the large language model
- Security controls such as encryption, access management, and monitoring are in place.
- Prompt data (entered data) is not used for training large language models (LLMs), ensuring confidentiality and compliance.
- Usage is auditable, and tools are integrated with organisational identity and access management systems where possible.

These approved tools can only be utilised within the confines of our NHS Wales Microsoft Tenancy (Copilot), NHS Wales Google Cloud Platform (Gemini) and PHW ChatGPT environment, the free versions available outside of this context should not be used for corporate purposes.

NHS Wales versions of these tools can be confirmed by checking the user profile that is logged in. The free version of co-pilot in NHS Wales is enabled for use by default for all staff.

PHW have an established AI register, this register maintains a list of AI tools that have been reviewed and approved, the register can be found here.

Use of any other AI tools requires approval through the AI request process detailed below.

4.2. Process for Accessing approved AI Tools

NHS Wales M365 Copilot Enterprise (Free Tier)
Instructions on accessing and utilising Co-pilot can be found in the below Copilot Guidance Document.

NHS Wales M365 Copilot Enterprise (Licensed Version)
If staff wish to request and utilise a licensed version of Copilot, the following process can be followed:

1. Request Approval: Discuss the business justification with your budget holder and confirm funding is available.
2. Service Desk Review: Submit call to the service desk advising on the need for a copilot license (include cost centre for your dept/ team).
3. Approval and Training: When approved, staff must review and agree to abide by the Copilot guidance set out below.
4. Ongoing Monitoring: Usage is monitored for compliance and risk management.

ChatGPT (Team Version)
ChatGPT is currently under a closed and controlled pilot, no further additional licenses will be procured or utilised until pilot finish (Nov 2025).

Guidance on the safe & effective use of permitted AI tools
Please find the following guidance for accessing and using Copilot here.
Please find the following guidance for ChatGPT here.

4.3. Requesting Alternative AI Tools

If the approved AI tools do not meet the specific needs of a use case for a project or department, staff may propose the use of an alternative generative AI tool.

This process requires:

1. Submission of a Formal Proposal: A detailed paper must be submitted outlining the business case, intended use, and justification for the alternative tool to AIDA.
2. Security and Compliance Review: The proposed tool must meet all UK data protection regulations, cybersecurity standards, and organisational policies.
3. Funding Approval: Any associated costs must be approved through the appropriate budgetary channels.

4. Demonstrated Value: The tool must clearly enhance organisational capabilities, efficiency, or innovation.

If unapproved generative AI tools are in use, caution should be used and the below general principles should be adhered to, any unapproved GenAI tools are to be submitted for review and approval via the AIDA process. Further guidance on the use of newly established & tools must be documented and published to the AIDA SharePoint hub.

AIDA maintains an AI register; this register can be used to identified what AI tools are currently approved and which have been previously requested.

There are several risks and concerns relating to the use of unapproved products such as:

1. Data Privacy & Confidentiality
   - Inputs may be stored or used to improve models unless you're using a version with strict data controls (like Azure OpenAI).
   - Risk of exposing sensitive business data if not properly managed.
2. Accuracy & Reliability
   - GenAI can produce incorrect or misleading information (hallucinations).
   - It may sound confident even when wrong, which can lead to poor decisions if unchecked.
3. Compliance & Legal Risks
   - Potential issues with data residency, GDPR, or industry-specific regulations.
   - Unclear IP ownership of generated content in some jurisdictions.
4. Bias & Ethics
   - Outputs may reflect biases present in training data.
   - Risk of generating inappropriate or non-inclusive content.
5. Integration & Oversight
   - Lack of proper governance or human oversight can lead to misuse.
   - Integration with internal systems may expose vulnerabilities if not secured.

5. General principals of safe and effective use of GenAI

To ensure the safe and effective use of generative AI tools, staff must adhere to the following guidelines:

1. Purposeful Use: Use AI where it adds clear value and aligns with organisational goals.
2. Public good: AI use should align with societal wellbeing and public trust
3. Human Oversight: Always maintain human oversight. GenAI should support—not replace—professional judgment. Regardless of the decision's impact, human review is essential for accountability, ethical responsibility, and protecting professional reputation.
4. Data Minimisation: Use only necessary data; avoid using personal, business or sensitive data.
5. Bias: Remain aware and be vigilant of potential algorithmic bias of GenAI.
6. Transparency: Clearly communicate when and how AI is used.
7. Security: Apply strong access controls, encryption, and monitoring.
8. Training and Awareness: Provide staff with training on responsible AI use.
9. Incident reporting: Please report all identified or suspected incidents to Cyber Security via the linked process here.

5.1. Prohibited Use in Healthcare Pathways

Generative AI tools must not be used in the design, recommendation, or decision-making processes related to healthcare pathways without strict prior approval from AIDA.

This includes, but is not limited to:

- Clinical diagnosis or treatment planning
- Patient-specific medical advice or triage
- Automated decision-making in healthcare delivery

This restriction is in place to ensure compliance with medical regulations, protect patient safety, and avoid unintended consequences from AI-generated outputs in sensitive health contexts. All healthcare-related decisions must remain under the direct supervision of qualified medical professionals.

AI use in healthcare pathways can be categorised as AI as a Medical Device (AIaMD) under the Medical Devices Regulations 2002 (UK MDR 2002) governed by Medicines and Healthcare products Regulatory Agency (MHRA).

5.2 Use of Copyright Protected and Intellectual Property Material in AI Tools

Staff must not input copyright protected, licensed, or proprietary material into AI tools unless they are explicitly authorised to do so and such use complies with UK intellectual property law, including:

- The Copyright, Designs and Patents Act 1988 (CDPA).
- The Intellectual Property Act 2014.

This includes, but is not limited to:

- Copyrighted text, images, data, software, and training materials.
- Licensed or subscription-based publications and databases.
- Unpublished internal documents, reports, policies, or intellectual property.
- Third party information shared in confidence or subject to contractual restriction.

Uploading copyrighted or proprietary material into AI tools without appropriate rights may constitute copyright infringement and a breach of organisational policy and contractual obligations.

Where protected material is required for legitimate work purposes, staff should:

- Use original summaries or high-level descriptions rather than copying verbatim text.
- Avoid uploading full documents or substantial extracts.
- Ensure that any use falls within permitted exceptions under the CDPA, where applicable.

5.3 Disclosure and Referencing of AI Generated or AI Assisted Content

Where AI tools are used to generate, draft, edit, analyse, or substantially contribute to content, this use must be transparent and appropriately disclosed.

Disclosure is required when:

- Content is published externally (e.g. reports, research, guidance, presentations).
- AI has materially influenced wording, structure, analysis, or conclusions.
- Disclosure is required by organisational, contractual, regulatory, or academic standards.

6. Other support and information

Guidance to civil servants on use of generative AI