Microsoft Copilot Extensibility Acceptable Use Policy

(September 2025 – v1.1)

NHS.net Connect: Copilot Extensibility Acceptable Use Policy

Overview

The Copilot Extensibility Acceptable Use Policy provides guidance for organisations and users interacting with Copilot Agents within the NHS.net Connect shared tenant. This includes agents created across the variety of Copilot Extensibility tooling, including Copilot Studio full, Copilot Studio lite, SharePoint Agents and M365 Agent Toolkit.

Copilot Studio allows users (Makers) to design, test, and deploy AI-powered agents that can automate tasks, answer queries, or retrieve information from connected knowledge sources.

The purpose of this policy is to ensure:

- Safe and responsible agent creation and use;
- Compliance with NHS.net Connect security and data governance standards;
- Avoidance of clinical, legal, or regulatory risks

This policy is complementary to the service-wide guidelines found in the NHS.net Connect Acceptable Use Policy.

Important Note:

- Copilot Agents are for non-clinical, administrative, and business use cases only.
- Please refer to and ensure you are compliant with your organisation's policies and guidelines throughout your development and use of Copilot Agents.

Please review the policy below before developing or accessing Copilot Agents, or and any Copilot Extensibility tooling including Copilot Studio full, SharePoint Agents, Copilot Studio lite and M365 Agent Toolkit.

Acceptable Use

- **Do not build or use Copilot Agents in clinical scenarios:** Copilot Agent building tools are provided strictly for administrative and business support purposes. Agents must not be used for any clinical activity, including informing or supporting clinical decision-making, direct patient care, or any activity requiring clinical judgement. Users must not develop, request, generate, or act upon agent outputs in any clinical context.
- **Always require users to label confidential and sensitive information as Official Sensitive:** Where a Copilot Agent has access to a SharePoint library as a knowledge source, it will be able to query information within this data based on the user's own permissions to those files, excluding files that are labelled as 'Official Sensitive'. Where agents are being used against this data,

its recommended to regularly review and update the necessary user permissions and data labels within the SharePoint library.

- **Risk acknowledgement:** Be aware of the risks of misuse, including data breaches and inaccurate results through reviewing the *AI Risk Guidance document* and training material available on the support site and Viva Learning. It is a local organisation's responsibility to conduct its own risk assessment for the enablement of Copilot Studio, considering its specific cyber maturity and operational context. Moreover, organisations are responsible for ongoing management of these risks through ensuring that all Copilot Studio agents undergo appropriate local governance, security and clinical safety assessments, and ensuring all new or updated agents follow the same assessment process.

- **Access and controls:** Make sure makers have access to features and data based on their role and responsibilities. Administrative and configuration level access should be limited to authorised personnel with appropriate training.

- **Introducing new functionality to your Copilot Studio-enabled environment:** Before requesting any changes to your Copilot Studio-enabled environment, such as through DLP exceptions, it is a local organisation's responsibility to conduct your own risk assessment on the risks the requested changes may introduce to your organisation's use of Copilot Studio as a service. Regular engagement with your local IG, Security and Clinical Safety team is strongly recommended prior to submitting requests for new functionality.

- **Guidance to Makers:** It is a local organisation's responsibility to direct makers to existing guidance on the support site, and to supplement this with information regarding local processes for agent development and governance. This includes but is not limited to:

    - **Ethical use:** Do not allow makers to develop Agents that allow or encourage prompts to be sent about illegal matters, requesting edits or images of others, or asking for outputs that are 'based on' or 'in the style' of a specific author or third party. This helps to ensure compliance with legal and ethical guidelines.
    - **Responsible use:** Ensure agents which are developed and used align with Microsoft's Responsible AI Policy and do not engage in actions that may conflict with it.
    - **Testing:** Before sharing an agent with its intended userbase, makers should ensure that they have sufficiently tested the agent to validate that it complies with the policies outlined here.
    - **Permissions for data use:** Ensure you have the rights to use the data generated by agents. For example, copyrighted material may be returned in a generated response. Whenever an agent uses knowledge sources (e.g.SharePoint, databases), the user running the agent must sign in with their own credentials, this ensures that data access always respects that user's individual permissions and security context. When using Tools (previously called "Actions"), configure your agent to require user

authentication. This ensures the Tool executes under the permissions of the signed-in user, not the agent author.
    - **Cross-Organisation Sharing:** Be aware of cross-organisation data visibility risks when publishing agents using shared knowledge sources.
    - **Use Copilot Agents for work purposes with your NHS.net Connect account**: When accessing Copilot Studio or other agent building tools, please use corporate or managed devices that are signed into your NHS.net Connect account to ensure Copilot Agent use is secured by enterprise level protections.

Appropriate use of generated content

- **Accuracy and bias**: Inform users of your agent to review all outputs for inaccuracies, bias, confidential information or offensive content to avoid sharing inaccurate, inappropriate or misleading content generated by an agent. Ask an agent to cite the source, review the generated content and discard or revise any problematic statements.
- **Transparency & Human Oversight:** Label all AI-generated content (e.g. with a sentence, icon or watermark) and include a disclaimer notifying the recipients when they are interacting with AI. Ensure humans review or supervise AI generated outputs when used in critical decision making or patient and public facing content as it may be inaccurate, unreliable or offensive. Include any know limitations that users should be made aware of and allow overrides or corrections on generated content.
- **End user understanding:** Make sure that you provide context to help users interpret content correctly. Ensure that users are educated on the limitations of the Agent to avoid a lack of trust, or over reliance.
- **Purposeful work:** It is advised that you use Copilot Agents and any agent-generated content only for work-related purposes. Additionally, do not use any output relating to a person for any purpose which could have a legal or material impact on the individual.

Data Privacy and Processing

- **Scope of Data Processing**: The roll-out of Copilot Extensibility does not make any changes to existing UK GDPR roles and responsibility arrangements to NHS.net Connect organisations, within the Microsoft tenant. NHS England and each organisation implementing M365 Copilot are Joint Data Controllers for their respective roles:
    - NHS England are the Data Controller for service configuration and provision.
    - Each organisation on the tenancy is Data Controller for the data they enter into NHS.net Connect.

    - Accenture is Data Processor acting upon instruction from NHS England and Microsoft are Sub-Data Processor.
- **Data residency**: As part of the NHS.net Connect Microsoft 365 tenant, Copilot Studio services operate within the Microsoft 365 UK Sovereign Cloud. While Copilot primarily maintains data handling within the UK, in certain scenarios, such as service fallback or capacity management the processing may occur in EU-based data centres. Additionally, when Bing Search is invoked as part of Copilot functionality, a limited excerpt of data in the form of a generated search query may be sent to and processed in the United States. Specific details on data residency for Copilot Studio are documented by Microsoft here. Further details around data processing for Bing Search are documented by Microsoft here.
- **Auditability**: Copilot Studio maintains full audit trails of agent creation, modification, publishing, and usage events. These are available in the **Unified Audit Log** and can be ingested into **Microsoft Sentinel** for monitoring and incident response.
- **Scope of processing**:
    - **Allowed**: Data already accessible to the user (files, emails, SharePoint, OneDrive, Graph API data).
    - **Not Advised**: HR Records and highly confidential datasets unless formally approved.
    - **Not allowed**: Patient data and clinical records
    - **Blocked by default**: Direct Line Channel, autonomous agent triggers, and external LLM connectors are disabled in the NHS.net Connect tenant.
- **DLP enforcement**: Environment-specific **Data Loss Prevention (DLP)** policies are enforced to block non-business and unsanctioned connectors, reducing the risk of inappropriate data sharing or exfiltration.
- **Sensitivity labels**: Where present, Copilot Agents observes NHS.net Connect sensitivity labels when processing data from knowledge sources. A technical pre-requisite of onboarding to services such as Copilot Studio is for the organisation to be auto-onboarded into the Global Sensitivity Label Policy. However, as labels are currently applied manually, makers must be aware that unlabelled data may be surfaced in agent responses, with the risk of sensitive information being unintentionally exposed.
- **Data minimisation**: Users should avoid entering sensitive, personal, or data that is otherwise unadvised into prompts unless it is strictly necessary and within local IG policy.

Training and escalation

Please review the *Copilot Agents guidance* available on the support site and Viva Learning to understand how to make and use Copilot Agents effectively.

If a Copilot Agent:

- Surfaces data that a user should not have access to,
- Behaves in a way that could compromise data security,
- Generates biased, inaccurate, or inappropriate content, or
- Violates DLP or sensitivity label policies,

Makers and Users must immediately:

1. **Stop using the agent.**
2. Report the incident to their Primary Local Admin (PLA).
3. The Primary Local Admin (PLA) will escalate to the **NHS.net Connect team** via the standard service request or incident management process.
4. Where applicable, incidents should also be reported through the organisation's **Information Governance and Security team** in line with UK GDPR and NHS incident handling requirements.
5. In serious cases (e.g., data breach), escalation will follow the **NHS England CSOC (Cyber Security Operations Centre)** process.