Clinical Safety Case Report

The Clinical Safety Case Report (CSCR) summarises and reviews the clinical safety activities performed to support the implementation, deployment and use of product(s), relative to their phase of the product lifecycle.

Document Management

Document filename: Clinical Safety Case Report
Document version number: 1.0
Organisation: MindBay Technologies
Issue Date: Jul 25, 2025
Document Owner: Taha Ouertani

Document Version Control:

Title	Version	Date approved	Changes	Approver
Clinical Safety Case Report	1.2	Jul 25, 2025, 09:42	Updated with new product name	Paul Jewell
Clinical Safety Case Report	1.0.1	Jul 20, 2025, 09:56	Renamed product from WellnessOne to MindBay; no scope change.	Taha Ouertani
Clinical Safety Case Report	1.0.0	Jun 23, 2025, 17:47	First version	Paul Jewell

Contents

Clinical Safety Case Report

The Clinical Safety Case Report (CSCR) summarises and reviews the clinical safety activities performed to support the implementation, deployment and use of product(s), relative to their phase of the product lifecycle.

Introduction

The Clinical Safety Case Report (CSCR) summarises and reviews the clinical safety activities carried out to support the implementation, deployment and use of the MindBay Technologies (hereafter referred to as "we", "our", "us") product(s), relative to the phase of the product lifecycle. The CSCR is structured to be iterative as new functionality and changes are introduced.

The CSCR forms part of our Clinical Risk Management File and is part of a set of clinical safety documentation, which has been produced in order to meet requirements of the DCB0129 standard of clinical risk management and review and to address the requirements of DCB0160 clinical risk management. This report contains the software definition, clinical hazards and mitigation/acceptance, and supporting evidence to provide an assurance statement on the clinical safety of our product(s).

Purpose

The purpose of this document is to clearly define our Clinical Risk Management processes in support of the development, modification and use of our products, as well as identifying, assessing and managing clinical safety hazards that may arise from the deployment and use of our product(s).

Specifically, this current CSCR has been written to support the implementation of the following product(s) and the current phase of their product lifecycle(s):

Product	Version	Lifecycle phase
MindBay	1.0	Pilot

Subsequent and wider deployment of existing or new products will be supported by appropriate Clinical Risk Management (CRM) and an uplift to this CSCR.

Scope

This CSCR applies to MindBay (the app) v1.0. This policy also applies to any local customisations or specific configurations made to the Mindbay MindBay v1.0 IT system by . CSCR are software product, version and development stage specific.

This scope extends to all clinical risk management linked activities undertaken during the MindBay v1.0 life cycle. All clinical functions and use cases that have potential to cause harm to patients and/or system users are incorporated.

Medical Device Regulatory Assessment

MindBay v1.0 is classified as a class I medical device.

Service Overview

Mindbay Technologies builds and operates digital mental health solutions developed through collaboration between mental health professionals and AI experts.

Our product, MindBay, is an interactive, chat-based conversational agent that delivers eight sessions of Cognitive Behavioural Therapy (CBT), along with between-session CBT exercises, meditation, and problem-solving tool.

MindBay uses Large Language Models (LLMs) trained on extensive datasets, enabling the generation of nuanced, context-specific responses in real time. Unlike earlier chatbots that relied on rigid rules or scripts, our tool offers a more personalised, precise and engaging experience, addressing key limitations of earlier AI-driven mental health interventions.

Clinical context

Before MindBay, GPs faced only two main first-line tools for mild-to-moderate depression or anxiety: (1) start an antidepressant - taken by 8.7 million people in England in 2023/24 despite common side-effects and patchy adherence - or (2) refer the patient to NHS Talking Therapies, where the average gap between first and second appointments exceeded three months and 1.76 million yearly referrals often received little more than review slots or self-help leaflets while they waited.

MindBay, a **Class I medical device** that delivers AI-supported, CBT-based conversations, is designed to **supplement - not replace - these established pathways**. A GP can generate an activation code through the electronic prescribing system to accompany medication or a Talking Therapies referral; clinicians can likewise issue a code at assessment so that patients on wait-lists gain immediate, guided CBT practice. Usage and symptom-change dashboards feed back to the referrer, supporting stepped-care reviews and helping teams prioritise follow-up. In this way, MindBay functions as a scalable, cost-effective **adjunct** that offers timely psychological support while preserving - and never substituting for - the clinician's judgment, pharmacotherapy, or face-to-face therapy.

Development lifecycle

The current phase of the product life cycle is the **initial feasibility and pilot stage.**

Existing systems

MindBay does not supersede any current IT solutions. During the feasibility and pilot phases it operates as a stand-alone product, with no interfaces to other IT systems.

Intended users

MindBay intended users are adult patients (18 years and older) with mild-to-moderate symptoms of depression and anxiety. GPs or NHS Talking Therapies will refer patients to MindBay.

For the feasibility and pilot stage, GP practices will send referral links via SMS, using an AccuRx template they create for this purpose. No additional clinical or administrative platforms are required; all patient interaction after referral is handled within the MindBay app.

Clinical Risk Management System

We have established a Clinical Risk Management System, with processes in place for both proactive and reactive clinical risk control to ensure that as many credible hazards and associated risks can be identified and anticipated as possible before they occur and any subsequent incidents that occur can be detected and resolved efficiently.

All our staff members work with clinical safety in mind, and think of potential clinical risks when requesting, designing and developing new changes to the software to ensure each new feature does not introduce new risk to the software.

Additionally, processes are in place so that all changes or modifications are reviewed prior to a release, in line with the **Clinical Safety Change Management Procedure**, to ensure no changes could result in patient harm, and that any bugs that pose a potential clinical risk can be identified and dealt with accordingly.

Our CRM activities cover the following:

- Risk Analysis
  - Scope Definition
  - Clinical Hazard Identification
  - Clinical Risk Estimation
- Risk Evaluation
  - Evaluation of initial level of risk of each identified hazard using pre-defined criteria
- Risk Control
  - Control Option Analysis
  - Clinical Risk Benefit Analysis
  - Control Measure Implementation
  - Completeness Evaluation

This comprehensive and repeatable clinical risk management process is applied throughout the lifecycle of the product.

Key personnel

The list below identifies key people responsible for clinical safety within our organisation, making up the Clinical Safety team.

Clinical Safety team

Name	Role	Responsibilities
Paul Jewell	*Clinical Safety Officer*	*Create clinical risk management processes* *Review clinical risk management documentation* *Approve clinical risk management documentation* *Lead hazard identification reviews* *Lead regular, episodic clinical risk management reviews*
Taha Ouertani	*Product Manager*	*Product development identification of risk* *Product development identification of controls for named risks* *Communicate product development with CSO* *Communicate user feedback with CSO* *Join hazard identification reviews* *Join regular, episodic clinical risk management reviews*

Clinical Safety Officer

Name	Professional qualifications	Professional Body	Competency/Skills/Experience	Clinical Safety Training
Paul Jewell	Medical doctor, BMBCh, MRCP	GMC: 7515784	Medical doctor with over 8 years experience, NHS Digital Clinical Safety Officer training and certification	NHS digital CSO Course

All documents created as part of the Clinical Safety Management System are maintained, version controlled and managed by the Clinical Safety Team and authorised by the Clinical Safety Officer. Documents are stored in the Clinical Risk Management File, which has appropriate access control measures in place.

Clinical safety review process

Clinical Safety change management procedure

We maintain a **Clinical Safety Change Management Procedure**. This ensures any product modifications or updates, or any new bugs that have been logged, will be assessed for any potential clinical risks that may be introduced, prior to release onto a live environment. These will be reviewed by the Clinical Safety team to ensure any change is either not released, or has an acceptable level of risk. The hazard log will be updated accordingly.

Clinical Safety incident management procedure

We maintain a **Clinical Safety Incident Management Procedure**, in which incidents that might impact patient safety are reported and managed appropriately, with requirement for CSO assessment if the severity dictates.

Clinical Risk Analysis

As is the nature of clinical software, hazards are always likely to be present. It is important to reduce the likelihood of these hazards occurring, as well as the potential impact of each hazard.

Clinical risk analysis involves identification of hazards, description of patient safety consequences, explanation of hazard causes and effects, identification of existing mitigating controls and estimation of clinical risk.

Hazard Identification (HAZID) is performed using the Structured What If Technique (SWIFT) and/or Functional Failure Analysis (FFA). At each step of the process; a checklist of possible 'things that could go wrong' are documented following structured discussions, taking into account technical, human and organisational factors within processes and procedures that may affect safety and the potential safety consequences; the functionality is questioned in terms of what could happen, what a user could do, what a user will know and what the system does. This hazard assessment considers information security, operational security, information governance, privacy and confidentiality and human behaviour and usability / GUI hazards and threats that could potentially have an impact on patient safety.

Hazards are defined using the cause > effect > hazard > harm framework:

- Cause: Consider events or actions that occur in the care-pathway that causes a deviation in the intended care process. Have all the potential scenarios that could cause a hazard or threat to occur been considered?
- Effect: The deviation that occurs in the care-pathway as a result of the cause
- Hazard: The condition that is created in the care-pathway as a result of the effect that has the potential to cause harm
- Harm: The realisation of harm to the patient

Details of any completed hazard assessment workshops can be found in the Appendix.

Description of patient safety consequences

MindBay offers personalised CBT-based intervention, but its AI-driven nature means patient safety depends on the accuracy of content, reliability of crisis detection and robustness of data protection. Consequences fall into two groups.

Direct consequences: software or model limitations may miss subtle references to self-harm, so the suicide-prevention workflow is never triggered, or may trigger without actually connecting the user to a live hotline. Incorrect or outdated helpline details can likewise leave high-risk users unsupported. Some patients or clinicians may assume the app is a stand-alone treatment and postpone professional therapy or medication, allowing symptoms to escalate. Finally, unauthorised access, whether a cyber-attack or someone picking up an unlocked phone, could expose highly sensitive session data, causing distress and loss of trust.

Indirect consequences: missed or ineffective crisis detection and over-reliance on self-help can push patients to present later and sicker, increasing the severity of episodes. Perceived failures can reduce confidence in digital therapeutics and in the referring NHS services, which may discourage future use of beneficial tools.

Clinical Risk Evaluation & Control

Each identified hazard is evaluated for the initial level of risk using pre-defined criteria.

The risk matrix grade is based on:

- Likelihood: How likely each identified consequence is to actually occur
- Consequence: For each combination of cause and hazard, were the hazard to be realised, what potential outcomes for patient safety are there
- Severity: The seriousness of each consequence is considered, in terms of individual patient harm
- Grade: Each risk is graded by consensus opinion (combination of consequence severity and likelihood), using the Department of Health Informatics Directorate (DHID) Risk Matrix

For each risk we consider the following as part of our evaluation:

- Actions: Actions taken or proposed to prevent or reduce the safety risk where possible.
- Acceptability: What top management consider to be acceptable risks and why
- Review: How often each residual risk/hazard is reassessed and actions evaluated

Qualitative probability

The defined levels for the probability of harm are as follows:

Likelihood Classification	Interpretation
Very high	Certain or almost certain; highly likely to occur
High	Not certain but very possible; reasonably expected to occur in the majority of cases
Medium	Possible
Low	Could occur but in the great majority of occasions will not
Very low	Negligible or nearly negligible possibility of occurring

Qualitative severity

The defined levels for the severity of harm are as follows:

Severity Classification	Interpretation	Number of Patients Affected
Catastrophic	Death	Multiple
	Permanent life-changing incapacity and any condition for which the prognosis is death or permanent life-changing incapacity; severe injury or severe incapacity from which recovery is not expected in the short term	Multiple
Major	Death	Single
	Permanent life-changing incapacity and any condition for which the prognosis is death or permanent life-changing incapacity; severe injury or severe incapacity from which recovery is not expected in the short term	Single
	Severe injury or severe incapacity from which recovery is expected in the short term	Multiple
	Severe psychological trauma	Multiple
Considerable	Severe injury or severe incapacity from which recovery is expected in the short term	Single
	Severe psychological trauma	Single
	Minor injury or injuries from which recovery is not expected in the short term.	Multiple
	Significant psychological trauma.	Multiple
Significant	Minor injury or injuries from which recovery is not expected in the short term.	Single
	Significant psychological trauma	Single
	Minor injury from which recovery is expected in the short term	Multiple
	Minor psychological upset; inconvenience	Multiple
Minor	Minor injury from which recovery is expected in the short term; minor psychological upset; inconvenience; any negligible severity	Single

Clinical Risk Management Risk Matrix

Likelihood	Very High	3	4	4	5	5
	High	2	3	3	4	5
	Medium	2	2	3	3	4
	Low	1	2	2	3	4
	Very Low	1	1	2	2	3
		Minor	Significant	Considerable	Major	Catastrophic
		Severity				

Risk Matrix key - Acceptability

5 4	Unacceptable level of risk. Mandatory elimination or control to reduce risk to an acceptable level
3	Undesirable level of risk Attempts should be made to eliminate or control to reduce risk to an acceptable level. Shall only be acceptable when further risk reduction is impractical.
2	Acceptable where cost of further reduction outweighs benefits gained.
1	Acceptable, no further action required

The subsequent approach for review and re-evaluation of these hazards will be the responsibility of the clinical safety team which will include:

- Consider the existing safeguards which would prevent or reduce the safety risk.
- Given those safeguards, grade each risk using the combination of consequence severity and likelihood.
- Identify recommendations for mitigation or controlling hazards to reduce risk.
- Given those mitigations, re-grade each risk using the DHID risk matrix.

Identified hazards

Summary of identified hazards for MindBay First Version.

Clinical Risk Category	Initial Risk (count)	Residual Risk (count)
Very High	-	-
High	-	-
Significant	-	-
Moderate	11	9
Low	-	2

Hazard initial and residual risk summary table for MindBay.

List of identified hazards

Hazard ID	Label	Initial Risk	Residual Risk
HAZ-001	Incorrect and potentially harmful advice	2	2
HAZ-002	No detection of severe distress, self-harm intention and suicidal ideation	2	2
HAZ-003	Safety feature activation without providing the adequate support	2	2
HAZ-004	Patient or clinicians overestimates app's capability, viewing it as a direct substitute for current standard of care	2	2
HAZ-005	Inaccessible language	2	2
HAZ-006	Unauthorized local user access	2	1
HAZ-007	App unavailability or technical issue	2	1
HAZ-008	Cyberattack or data breach compromises patient information	2	2
HAZ-009	Incorrect helpline signposting	2	2
HAZ-010	Forseeable misuse	2	2
HAZ-011	Poor engagement with the app	2	2

Discussion of relevant hazards

The most important risk is that patients treat MindBay as **a replacement for** consulting their GP, taking their medication, seeing a mental health professional or referring themselves to NHS Talking Therapies. We address this with clear disclaimers at sign-up, pop-up reminders inside the app, and plain-language manuals that spell out what the eight AI-driven CBT sessions can and cannot do. Together, these measures remind users that the app is a supplement, not a substitute.

The second-ranked risk is that MindBay might miss signs of severe distress, self-harm or suicidal thinking. Software bugs, obscure wording and the occasional referral of users with more severe mental-health illness may raise the risk. To mitigate the risk, we run layered unit, integration and system tests; push every change through an automated CI/CD pipeline; and review code line by line. The detection protocol itself is trained on open suicide-ideation datasets and is manually and regularly tested against many edge-case phrases. Referring clinicians rely on their clinical judgement and their existing diagnostic pathways to identify high-risk patients and route them to more appropriate services.

A third hazard is that the safety feature may trigger yet still fail to guide the user to the right help. Out-of-date crisis numbers or missing local resources could fail to provide users promptly with the correct helplines. The same robust testing, code reviews and rapid patching protect the underlying logic, while a scheduled content review checks every helpline and NHS signpost for accuracy and regional fit. Additional audits track how quickly and correctly the routing works in practice, allowing prompt fixes when performance slips.

Hazard log

The full **Hazard Log** can be found separately. It is provided in the NHS Digital format, and provides evidence to support compliance with DCB0129

Hazard Workshops

Date: 5th June 2025

Attendance: Paul Jewell(CSO and facilitator), Rosie Taylor, Taha Ouertani, Vaidotas Gulbinas, Mouafak Dakhlaoui

Minutes: Hazard Workshop Minutes

Test issues

No outstanding test issues identified at time of writing.

Assessment of third party products

Third party product	Clinical safety consideration
Microsoft Azure NHS DSPT ISO/IEC 27001 (Information Security) ISO/IEC 27017 (Cloud Security) ISO/IEC 27018 (Cloud Privacy) ISO/IEC 27701 (Privacy Information) ISO 22301 (Business Continuity) ISO 9001 (Quality Management)	We host MindBay and all patient data exclusively in Microsoft Azure's UK data centres. Azure guarantees AES-256 encryption at rest, TLS encryption in transit, role-based access controls, detailed audit logging and automated backups. It holds NHS Data Security and Protection Toolkit accreditation and complies with GDPR, delivering ≥99.9 % uptime. This infrastructure stability and data integrity underpins safe, uninterrupted access for both patients and clinicians.
Microsoft Azure AI Services ISO/IEC 27001 (Information Security) ISO/IEC 27017 (Cloud Security) ISO/IEC 27018 (Cloud Privacy) ISO/IEC 27701 (Privacy Information) ISO 22301 (Business Continuity) ISO 9001 (Quality Management)	All language-processing and risk-detection models run within Azure AI Services under Microsoft's Responsible AI framework. Each model is version-controlled, validated against established clinical data sets and subjected to continuous performance monitoring. Crucially, the AI only flags potential risk and never autonomously issues clinical advice, ensuring our clinical governance remains robust.

Summary Safety Statement

As the appointed Clinical Safety Officer for Mindbay, I have overseen the implementation of the clinical risk management system and the completion of clinical risk analysis. This report summarises the clinical hazards identified through multidisciplinary clinical safety meeting and workshops with key stakeholders in Mindbay.

It is written in the context of the intended use within the scope of a UK pilot. All identified hazards have been evaluated using a structured approach to determine their likelihood and severity. Subsequently, where necessary, appropriate risk controls have been implemented to ensure clinical risk is minimised and the clinical safety of patients is upheld.

To summarise, this process has identified nine hazards, with none deemed to have high risk. All moderate and low risk hazards identified have been mitigated and controls put in place to ensure that the residual risk is as low as reasonably practicable.

In summary, the Mindbay MindBay application is a very valuable adjunct to existing care. Before MindBay, GPs faced only two main first-line tools for mild-to-moderate depression or anxiety: (1) start an antidepressant - taken by 8.7 million people in England in 2023/24 despite common side-effects and patchy adherence - or (2) refer the patient to NHS Talking Therapies, where the average gap between first and second appointments exceeded three months and 1.76 million yearly referrals often received little more than review slots or self-help leaflets while they waited.

Any risks identified are minimal and appropriately mitigated, and the benefit of using Mindbay greatly outweighs any of the identified risks. Any further product developments, deployment phases or safety incidents that arise following this report will follow the same rigorous risk management process to ensure ongoing safety. Key deliverable documents, including the Hazard Log and this Clinical Safety Case Report, will be updated to reflect this.

I confirm that the risks identified and mitigations implemented provide a sufficient basis to consider MindBay clinical safe for its first NHS pilot, and has my full confidence and support.

Dr. Paul Jewell, Clinical Safety Officer

Quality Assurance and Document Approval

This CSCR has been developed, reviewed, and approved in accordance with the arrangement described within the **Clinical Risk Management Plan**.

References

Reference 1

DCB0129: Clinical Risk Management: its Application in the Manufacture of Health IT Systems, NHS Digital, NHS England, https://digital.nhs.uk/data-and-information/information-standards/information-standards-and-data-collections-including-extractions/publications-and-notifications/standards-and-collections/dcb0129-clinical-risk-management-its-application-in-the-manufacture-of-health-it-systems

Reference 2

DCB0160: Clinical Risk Management: its Application in the Deployment and Use of Health IT Systems, NHS Digital, NHS England https://digital.nhs.uk/data-and-information/information-standards/information-standards-and-data-collections-including-extractions/publications-and-notifications/standards-and-collections/dcb0160-clinical-risk-management-its-application-in-the-deployment-and-use-of-health-it-systems

Reference 3

Guidance: Medical device stand-alone software including apps (including IVDMDs), MHRA (Medicines & Healthcare products Regulatory Agency) https://assets.publishing.service.gov.uk/media/64a7d22d7a4c230013bba33c/Medical_device_stand-alone_software_including_apps__including_IVDMDs_.pdf