



Medicines & Healthcare products
Regulatory Agency

National Commission into the Regulation of AI in Healthcare

Government Response to the National Commission's recommendations

Published Tuesday 6th October 2026



© Crown copyright 2026

Open Government Licence



Produced by the Medicines and Healthcare products Regulatory Agency. www.gov.uk/mhra

You may re-use this information (excluding logos) free of charge in any format or medium, under the terms of the Open Government Licence. To view this licence, visit <http://www.nationalarchives.gov.uk/doc/open-government-licence> or email: psi@nationalarchives.gsi.gov.uk.

Where we have identified any third-party copyright material you will need to obtain permission from the copyright holders concerned.

The names, images and logos identifying the Medicines and Healthcare products Regulatory Agency are proprietary marks. All the agency's logos are registered trademarks and cannot be used without the agency's explicit permission.

Contents

Foreword	4
Introduction.....	6
Priority areas for delivery.....	8
Proportionate Lifecycle Regulation.....	8
System-wide Responsibility and Safe Management.....	9
Trust, Transparency and Predictability	10
Delivery and Governance.....	11
Action Plan.....	12
Key Principle: Ensure regulatory requirements are proportionate to a device's risk and benefits, tailored to evidentiary needs across the product lifecycle and support equitable access to safe and effective software and AI-enabled technologies	12
Key Principle: Consider regulation and safe management of software and AI health technologies as a system-wide function, with each part of the system delivering its responsibilities so these technologies can be used effectively and responsibly.	22
Key Principle: Design a framework that enhances trust, transparency and predictability for all stakeholders, to enable responsible development, adoption and use of safe and effective AI technologies	27
Annex A. List of Recommendations	31

Foreword

On behalf of the UK Government, I welcome the important recommendations made by the National Commission into the Regulation of AI in Healthcare.

I give my thanks to Professor Alastair Denniston and Professor Henrietta Hughes OBE, the Commissioners and Working Group members, and officials who have supported the National Commission's work. AI-enabled products, when deployed correctly and with proportionate safeguards, can support better outcomes for patients, free up professionals' time, and support the development and expansion of promising UK-based manufacturers, driving growth and job creation across the four nations of the UK. The National Commission's important work has therefore set high expectations for the regulation and use of AI in healthcare and has the potential to become an important, pioneering step towards establishing a national pathway for the safe, effective and scalable adoption of AI in healthcare.

We have therefore carefully considered the National Commission's recommendations and welcome the breadth of analysis that underpins them. We are accepting all 44 of their recommendations. This response outlines how the Government will use these recommendations to create a future ready regulatory framework for the development, regulation and adoption of AI in healthcare. But this publication does not mark the end of the process. It marks the next stage in an ongoing conversation about how the UK can maximise the benefits in this area, while managing the risks. The Government will continue to work closely with regulators, the NHS, industry, academia, patients and the public as we implement these recommendations and ensure that AI supports better, safer and more effective care for people across the country.

At the launch of the National Commission's recommendations, I said that innovation succeeds when it becomes implementation, and 'we dream' turns into 'we do'.

This Government agrees that the UK needs a new approach to AI product regulation in healthcare that is proportionate, risk and lifecycle based and takes a system wide approach to overseeing safe use. This is essential for shifting the narrative on AI adoption away from highlighting the potential benefits of AI in healthcare to demonstrating the benefits in real world settings, enabling successful innovation to be adopted at scale, and building the trust that is fundamental for the use of these pioneering products.

The actions set out in this response are designed not only to strengthen the regulatory framework, but also to deliver tangible benefits across the healthcare ecosystem. For patients and the public, we will create transparency and opportunities to inform policy decisions on the use of AI in healthcare. This includes better support to use devices that have been designed with their needs in mind.

For healthcare professionals, we will support greater device usability and AI literacy, so they have greater confidence to use AI products. For healthcare providers, we will create greater opportunities to participate in the testing of innovative medical devices that can directly benefit their patients. For medical device manufacturers, we will support innovation with regulation that is balanced between pre-market and post-market stages and clearer routes to market for promising devices.

This Government is committed to good growth in every postcode of the UK. By creating the conditions for safe, effective and trusted healthcare innovation, we can improve outcomes for patients while supporting the growth of a globally competitive life sciences and health technology sector all across the four nations of the UK.

The same collaborative spirit set by the Commission will now be taken by the Government when turning these recommendations into reality. We are committed to working with partners across the health system, industry, academia and the devolved authorities to deliver these commitments. We look forward to taking forward the recommendations of the National Commission and ensuring that the UK remains at the forefront of safe and responsible AI innovation in healthcare.

James Frith MP

**PARLIAMENTARY UNDER-SECRETARY OF STATE
FOR HEALTH INNOVATION**

Introduction

The government is fully committed to improving outcomes for patients and supporting sustainable, high-quality health and care systems. In doing so, the government recognises the transformative opportunity that AI technology promises the healthcare sector. AI products are already enabling earlier diagnosis of disease, automating repetitive clinical tasks, and reducing administrative burdens across the health system. For example, stroke imaging technology is being used across NHS stroke services to help clinicians rapidly assess brain scans and identify patients who may be eligible for life-changing interventions¹. By enabling faster, more consistent decision-making in time-critical situations, these tools are already helping patients access treatment sooner, improving patient outcomes and freeing up resource across the NHS.

We are only at the beginning of this transformation. The opportunity is significant, but so too is the responsibility to ensure that these technologies can be trusted by patients, healthcare professionals and the wider public. The [10 Year Health Plan for England](#)² and the [Life Sciences Sector Plan](#)³ set out an ambition for the NHS to become the most AI-enabled healthcare system in the world. This is not an end in itself, because getting the safe and effective use of AI right not only promises exciting upsides, it can help the NHS deliver the standard of care that the British public expects and deserves.

Earlier in 2026, the Government committed up to £10 billion over the next three years to invest in technology, digital, and data initiatives. The funding will support the transformation of the NHS into a fully digitally enabled service, operating 24/7 both online and offline, as laid out in the 10-Year Health Plan.

To realise the benefits of AI across health and care services, it is critical that we have a proportionate and agile regulatory system. Regulations provide the predictability and certainty needed for innovators, companies and investors to thrive while maintaining robust patient safety controls. How AI in healthcare is regulated depends on its intended purpose, functionality and use. Some AI-enabled products fall within the scope of the UK Medical Devices Regulations 2002 (as amended) (UK MDR 2002), while others are subject to data protection, professional, organisational or other sector-specific regulatory frameworks. This joint response brings together those organisations involved in developing and implementing such standards and regulatory frameworks.

¹ Nagarathnam et al (2025) [Artificial intelligence imaging decision support for acute stroke treatment in England: a prospective observational study](#).

² Department for Health and Social Care (2025) [10 Year Health Plan for England: fit for the future](#).

³ Department for Business and Trade (2025) [Life Sciences Sector Plan](#).

Because the use of AI has developed considerably since UK MDR 2002 was drafted, as well as the nature and pace of its development, the current medical device regulations are not fit for purpose.

The Medicines and Healthcare products Regulatory Agency (MHRA) launched the National AI Commission in September 2025 in recognition of the challenge that the healthcare system was never set up to fully realise the opportunities presented by AI. We believe that this is a pivotal moment to rethink how healthcare is delivered, to reaffirm the values that underpin our system, and create the incentives and conditions that will enable innovation while improving outcomes and safety for all patients across the UK.

Over the past year, the MHRA, working with health system partners, have laid the foundations for this transformation by convening leading and diverse voices from healthcare, technology and medical devices industry, patient groups, and government representatives across the four UK nations to develop a shared vision for the future of AI-enabled healthcare. The National Commission, working with the MHRA and partners, has undertaken one of the most comprehensive engagement exercises on the role of AI in healthcare, ensuring the Commission's work was grounded in a rich, diverse and extensive evidence base. The leadership of the Patient Safety Commissioner ensured that patient needs, experiences and expectations remained at the heart of the process. Together, this work reflects our commitment to creating the conditions for AI to deliver meaningful improvements across the health system, in a way that is fast, trusted, and safe.

The National Commission concludes that fundamental shifts are needed for our regulatory regime to support this transformation so that patients can benefit earlier from AI technologies that are safe and trusted. The National Commission's report provides a clear framework for the next phase of regulatory reform. The recommendations align with what is set out in the 10-Year Health Plan for England, the Life Sciences Sector Plan and the AI Opportunities Action Plan⁴, translating those ambitions into a practical programme of action.

The government appreciates the depth and breadth of the work of the National Commission and intends to accept all of the recommendations it put forward. Delivering the full benefits of AI in healthcare should be viewed as a shared endeavour across the healthcare system. The Department of Health and Social Care (DHSC) and devolved health departments have an important leadership role in establishing the conditions for safe adoption, including supporting organisational readiness, workforce capabilities and effective governance. Regulators such as the MHRA have an important role to play as it will require a regulatory system that is agile, proportionate and responsive to rapidly evolving technologies. It will

⁴ Department for Science, Innovation, Technology (2025) [AI Opportunities Action Plan](#).

also require stronger regulatory capabilities, and a continued focus on equitable access, patient safety and public trust.

Although the MHRA is the competent regulatory authority for medicines and medical devices in the United Kingdom, the government also recognises that successful implementation will require partnership across the healthcare system and with the devolved authorities. The government is therefore committed to working collaboratively with relevant organisations to deliver these recommendations. These include, but is not limited to, the National Institute of Health and Care Excellence (NICE), the Health Research Authority (HRA), the Care Quality Commission (CQC), and others, alongside equivalent organisations in devolved authorities, such as the Scottish Intercollegiate Guidance Network (SIGN), Health Technology Wales (HTW), Scottish Health Technologies Group (SHTG), Healthcare Improvement Scotland (HIS), Healthcare Inspectorate Wales (HIW), and the Regulation and Quality Improvement Authority (Northern Ireland) (RQIA).

Priority areas for delivery

Over the next year, the government will focus on the following priority areas to begin delivering this transformation.

Proportionate Lifecycle Regulation

The National Commission's first theme focuses on proportionate lifecycle regulation considering how regulatory approaches should evolve to reflect the characteristics of AI-enabled technologies. The National Commission reflected on how regulation needs to be proportionate to risk and benefit, informed by evidence generated across the product lifecycle, and designed to support timely and equitable access to safe and effective technologies. Government priorities on this theme include:

Priority 1: Modernise the regulatory framework for AI-enabled medical devices

The MHRA will take steps to modernise the regulatory framework for software and AI-enabled medical devices to ensure it remains fit for rapidly evolving technologies through legislative reform and clear guidance. This includes providing clearer lines on when a product qualifies as a medical device, an improved classification approach and clarity on how intended purpose includes design and functionality considerations.

Priority 2: Delivering equitable access and outcomes for patients and members of the public, with a special focus on underserved populations

Partners across the health, regulatory and care system will work together to ensure that benefits of AI-enabled healthcare technologies are realised fairly across the population, with particular attention to underserved or underrepresented groups. Building on the National Commission's findings, government will promote approaches that consider equity throughout

the lifecycle of AI-enabled technologies, from development and evaluation through to deployment, monitoring and ongoing assurance.

Priority 3: Shift towards stronger post-market assurance

The MHRA will work to develop and implement a more proportionate, lifecycle-based approach to regulation that better balances pre-market evidence requirements with robust post-market assurance. This will include enhanced post-market surveillance mechanisms, clearer evidence expectations and improved mechanisms for monitoring and responding to quality signals, device safety and performance changes in real-world settings, enabling a conscious shift towards ongoing assurance rather than overreliance on pre-market assessments and passive adverse event reporting.

Priority 4: Enable safe iteration and earlier patient access

To support innovation and earlier patient access, the MHRA will develop robust and tailored approaches to bringing safe AI innovations to patients faster, including by exploring practical, forward-looking implementation of Predetermined Change Control Plans (PCCPs) to better manage changes, novel staged authorisation pathways that enable earlier access while generating real world evidence, and expanded use of regulatory sandboxes to test novel technologies and regulatory approaches. To achieve this, the MHRA will work with NICE, and equivalent organisations in devolved authorities, and other partners.

System-wide Responsibility and Safe Management

The National Commission's second theme examines how the safe adoption, management and oversight of AI-enabled technologies should be supported across the healthcare system. It considers how responsibilities should be shared across manufacturers, healthcare providers, professionals, regulators and other system partners, and how accountability, organisational readiness, workforce capability and governance can support their safe and effective use. Government priorities on this theme include:

Priority 5: Creating the assurance conditions necessary for safe, real-world deployment, including risk controls

The MHRA will develop guidance setting out how manufacturers should identify, describe and communicate the operational conditions necessary for the safe deployment and use of AI-enabled devices. This will also involve providing practical guidance and illustrative examples of risk controls.

Priority 6: Building system capability and readiness for widespread adoption of AI, including training of the workforce.

DHSC and NHS England will work with the devolved health departments to develop and maintain an AI Readiness Toolbox to support healthcare organisations in assessing and preparing for the safe adoption of AI-enabled technologies.

Professional regulators, education bodies and healthcare organisations will work together to strengthen AI literacy and ensure that healthcare professionals receive the training and support necessary to safely use AI-enabled technologies in clinical practice.

Priority 7: Strengthening accountability, including through a culture of safety and reporting

DHSC, MHRA, CQC and equivalent organisations in the devolved authorities, NHS Resolution (NHSR), the Professional Standards Authority (PSA) and professional regulators will form a working group to ensure that responsibilities are clear across manufacturers, healthcare providers, professionals and regulators.

DHSC, and the devolved health departments, working together with the MHRA, CQC, and equivalent organisations in devolved authorities will focus on strengthening mechanisms for reporting, monitoring and learning from the real-world use of AI technologies, enabling safety concerns and implementation challenges to be identified and addressed at an early stage. DHSC will work with the Patient Safety Commissioner and relevant organisations across the system to ensure patients have access to redress.

Trust, Transparency and Predictability

The National Commission's third theme explores how trust, transparency and predictability can be strengthened across the healthcare ecosystem. It considers how greater transparency, clearer accountability, meaningful stakeholder involvement and more navigable pathways can support the responsible development, adoption and use of safe and effective AI technologies. Government priorities on this theme include:

Priority 8: Strengthening transparency and public confidence

The MHRA and partners will increase patient and public engagement in regulatory decision-making and explore new approaches to communicating information about AI-enabled technologies and emerging safety concerns.

DHSC will work with healthcare partners across the UK to improve transparency around the use of AI-enabled technologies in healthcare and support a consistent, proportionate approach to informing patients about how these technologies are used in their care.

Priority 9: Improving regulatory clarity and navigability

The MHRA and partners will provide clearer guidance, resources and opportunities for early engagement for developers and manufacturers, including user-centred design and regulatory expectations across the product lifecycle. These measures will make regulatory and assurance pathways easier to navigate and enable responsible innovation by improving predictability and confidence in the regulatory system.

Priority 10: Strengthening regulatory capability and the evidence base for AI in healthcare

The MHRA and NICE, alongside equivalent organisations in the devolved authorities, will work together to ensure that regulatory, evidence generation and assurance expectations are well-aligned across the lifecycle of AI-enabled health technologies, creating a more coherent and joined-up experience for developers and healthcare providers – and ensuring that the best technologies are adopted effectively.

The HRA and equivalent organisations in the devolved authorities, will build on its guidance and decision-support tools to improve clarity on when AI development, deployment, evaluation and monitoring activities constitute research and require research governance, ensuring that oversight arrangements remain proportionate and aligned with wider system guidance.

Delivery and Governance

To ensure effective delivery of this programme, the government will establish a cross-system Programme Board with representation from key partner organisations across the health, regulatory, research system and representatives from the devolved authorities. The Board will convene quarterly to oversee progress, monitor delivery against agreed objectives, identify and address emerging risks, and ensure coordinated action across participating organisations. The Board will also be tasked with identifying how to foster collaboration and mechanisms to facilitate coordination between the four nations on the delivery of shared recommendations and goals that have come out of the National Commission. This governance structure will support shared accountability for delivering the programme's ambitions and maintaining a focus on outcomes for patients, healthcare professionals and the wider health system.

An implementation plan and roadmap, published by Spring 2027, will provide further detail on roles, responsibilities, timelines and key deliverables, with annual progress updates thereafter to maintain transparency, accountability and predictability across the ecosystem.

The Board will provide progress updates to the Life Sciences Delivery Board (LSDB) on a quarterly basis, ensuring ongoing senior oversight and alignment with wider government priorities.

Action Plan

Key Principle: Ensure regulatory requirements are proportionate to a device's risk and benefits, tailored to evidentiary needs across the product lifecycle and support equitable access to safe and effective software and AI-enabled technologies

Recommendation	Government Response	Timeline
Recommendation 1: Modernising the Regulation	(Accept) The MHRA will launch a consultation on a new, proportionate approach to classification and qualification for software and AI-enabled devices. The MHRA will then use secondary legislation to introduce: • An updated definition of a medical device that builds upon the current definition by providing clarity on when certain types of software and AI-enabled products are not classed as medical devices (such as certain administrative, general wellbeing and decision support software). This may include carve-outs for certain software and AI-enabled products that would not be regulated as medical devices. • A classification system and approach which better reflects a proportionate approach to the benefits and risks of software and AI-enabled devices.	MHRA intends to launch the consultation by Spring 2027

	An updated definition of intended purpose to ensure factors such as device design, functionality and context of use are adequately considered when evaluating an intended purpose for device qualification and classification.	
Recommendation 2: Improving Regulatory Guidance	(Accept) As a starting point, the MHRA will review and identify opportunities to provide further clarity on device qualification, classification and intended purpose through guidance under the existing regulatory framework while legislative reform is considered (Rec 1). For example, the MHRA intends to provide an update to the Guidance: Medical device stand-alone software including apps (also covering IVDMDs). This update will also provide additional clarity on the types of software and AI-enabled products that do not qualify as medical devices. Following any updates to the UK Medical Device Regulations through secondary legislation, the MHRA will issue new guidance for manufacturers. This guidance will, among other things: - Detail the rules for qualification of AI-enabled medical devices. - Specify how risk-based classification systems will be applied to AI-enabled medical devices. - Provide details on the interpretation of intended purpose for these devices.	Further details to be shared by Spring 2027
Recommendation 3: Proportionate Regulatory Oversight	(Accept) The MHRA will consider how the regulatory framework can best support proportionate, flexible and tailored regulatory oversight over time. This will be in response to the best available evidence about the performance, capability, benefits and risks of AI-enabled medical	Further details to be shared by Spring 2027

	devices. The MHRA will consider how to clearly signal and consistently apply our allocation of resources and oversight to best assure the public health. For example, through issuing draft guidance on the application of enforcement discretion of sufficiently low risk software and AI-enabled medical devices.	
Recommendation 4: Function Based Regulation	(Accept) The MHRA will consider options, including secondary legislation, for how to introduce a function-based approach to regulation that focuses on the functionality with a medical intended purpose in products where there may be multiple purposes. The introduction of function-based regulation will be supported with clear guidance for manufacturers.	Further details to be shared by Spring 2027
Recommendation 5: Lifecycle-based evidence and oversight	(Accept) The MHRA will work with partners, to explore options to establish a more lifecycle-based regulatory framework. Under this approach, evidence generation and oversight should increasingly take place across the product lifecycle, supported by robust, ongoing post-market oversight and less reliance on point-in-time pre-market assessments. This may include seeking potential licensing powers, legislative reform, and new guidance. The MHRA would seek to work with NICE, and equivalent organisations in the devolved authorities, to align on evidence expectations. To strengthen clarity for manufacturers on how a revised lifecycle-based approach to regulation works in practice, the MHRA will issue guidance or other resources to address how this approach would be applied to emerging technologies.	Further details to be shared by Spring 2027

Recommendation 6: Predetermined Change Control Plans	(Accept) The MHRA will issue draft guidance on PCCPs for public comment. This guidance will describe a proposed approach to PCCPs that allows for manufacturers to establish boundaries to define the scope of allowable change (rather than only allowing specific modifications to be prespecified), including clear approaches for adaptive AI-enabled medical devices, while maintaining a device's intended use.	MHRA intends to issue draft guidance by December 2026
Recommendation 7: Master Files for general-purpose models	(Accept) The MHRA will engage with experts to explore options for proportionate information on general purpose AI models that would be valuable, practical and supportive of establishing device safety and performance. This will be supported by the development of guidance for manufacturers. MHRA's work may also include initiation of a Master File pilot focused on general-purpose models or platforms.	Further details to be shared by Spring 2027
Recommendation 8: Reporting product dependencies on underlying general- purpose models	(Accept) The MHRA will provide guidance on the types of information, including provenance, nature, dependencies, related risks and their controls, that such software and AI-enabled devices will need to provide as part of a regulatory submission. The MHRA will work closely with DHSC and the devolved health departments to ensure expectations to provide this information are built into procurement processes and contract terms for software and AI-enabled products.	Further details to be shared by Spring 2027
Recommendation 9: Transparent information for users and the public	(Accept) The MHRA will consider how best to support manufacturers in communicating relevant information about how these devices operate, inform clinical decision-making and may affect care. This will include providing appropriate guidance on relevant approaches, such as model cards, and accessible user-interface design and dynamic labelling.	Further details to be shared by Spring 2027

Recommendation 10: Cyber Security guidance and resources	(Accept) The MHRA will issue draft guidance and educational resources setting out cyber security expectations for software and AI-enabled medical devices, working with partners as appropriate. This guidance will outline practices that should be considered throughout the lifecycle to ensure ongoing safety and performance. The MHRA will also continue to demonstrate international leadership by leveraging its role as co-chair of the IMDRF cyber security working group to promote greater international alignment and share learning with like-minded international regulatory partners.	Further details to be shared by Spring 2027
Recommendation 11: Direct to consumer devices	(Accept) The MHRA will issue guidance on the regulatory considerations and expectations for direct-to-consumer applications and wearables with medical device functionalities. This guidance will outline how to carefully consider the intersection of various aspects of the proportionate, lifecycle-based framework in the context of direct-to-consumer devices, such as transparency, user-centred design and post-market monitoring.	Further details to be shared by Spring 2027
Recommendation 12: Human factors and usability expectations	(Accept) The MHRA will issue guidance for how human factors and usability expectations should be considered across the product lifecycle, as it is important that the needs of different groups be considered as part of the design, development and maintenance of software and AI-enabled devices.	Further details to be shared by Spring 2027
Recommendation 13: Recognising the diversity of the intended population	(Accept) The MHRA will work with partners to issue guidance to support manufacturers in demonstrating that their devices are safe, perform as intended and do not create or exacerbate disparities for relevant population groups, including children, women and underrepresented or underserved population groups, including ethnic	Further details to be shared by Spring 2027

	minorities. This guidance will address expectations for representative data, evaluation of performance across relevant groups and ongoing monitoring and evidence collection throughout the product lifecycle.	
Recommendation 14: Staged authorisations	(Accept) The MHRA will explore options to introduce staged authorisation pathways, including through proposed licensing powers, regulatory sandboxes, and support for proposed legislative reforms to create statutory mechanisms for regulatory sandboxes. The MHRA will work with NICE, and equivalent organisations in the devolved authorities, to consider how evidence that is generated through any staged authorisation processes can be used to support economic assessment and appraisal of these devices. The MHRA will also consider how these approaches will impact the ecosystem. The MHRA commits to engaging with relevant stakeholders to ensure a coordinated and safe approach to device deployment, adoption, and scaling through such pathways.	Further details to be shared Spring 2027
Recommendation 15: Regulatory sandboxes	(Accept) The MHRA, in collaboration with system partners, will take forward a number of actions to ensure regulatory sandboxes support the development, testing and deployment of innovative software and AI-enabled medical devices and corresponding regulatory approaches. This includes: • Delivering phase 3 of the MHRA's AI Airlock programme • Working with the Department for Business, Innovation, Science and Trade (BIST) to support sandbox provisions into their upcoming Regulating for Growth Bill, enabling the ongoing	Further details to be shared Spring 2027

	development of sandboxes for software and AI-enabled medical devices. • Working with groups across the UK to consider opportunities to engage with MHRA on sandboxes which reflect their unique needs. • Working with NICE, and equivalent organisations in the devolved authorities, to consider how evidence generated through regulatory sandboxes can be used to support the economic assessment and appraisal of these devices.	
Recommendation 16: Recognition and reliance pathways and international harmonisation	(Accept) The MHRA will continue efforts to maintain and enhance international harmonisation, where appropriate to do so. For example, this will come through the MHRA's membership of the International Medical Device Regulators Forum (IMDRF), and the trilateral and bilateral relationships with international partners such as the Food and Drug Administration (FDA), Health Canada and the Therapeutic Goods Administration (TGA), among others. The MHRA will also take forward proposals set out in its 2026 economic growth goals⁵, by working with recognised international regulators to establish specific recognition and reliance arrangements for AI-enabled medical devices. These pathways can support the growth of the UK market for medical technologies, while enabling patients to benefit from safe and innovative technologies developed around the world.	Further details to be shared by Spring 2027

⁵ Medicines and Healthcare products Regulatory Agency (2026) [Medicines and Healthcare products Regulatory Agency \(MHRA\) economic growth goals](#).

Recommendation 17: Tailored post-market surveillance	(Accept) The MHRA will build on its existing post-market surveillance requirements and introduce an enhanced toolkit of mechanisms that could be required individually or in combination to enable a more tailored approach to post-market oversight. This will include consideration of any necessary legislative reform and issuance of guidance. The MHRA will also utilise the third phase of the AI Airlock to investigate post-market surveillance challenges/practices as a key focus area, building upon the findings from the initial phases.	Further details to be shared by Spring 2027
Recommendation 18: Reporting and information sharing across the lifecycle	(Accept) The MHRA will work with manufacturers, other regulators, healthcare providers and professionals, and the public to introduce strengthened mechanisms for reporting and information sharing, including best practice, across the lifecycle. This will include: • Monitoring and reporting processes and requirements to be included in PCCPs to support continued safety and performance while a device iterates • Working with manufacturers to consider mechanisms to allow product users to more seamlessly share feedback during use and providing clarity as part of relevant guidance, where appropriate • Exploring ways to enhance partnerships and data sharing to enable information to be shared in a more transparent and timely way • Delivering improvements to the Yellow Card scheme. This could include additional adverse incident reporting mechanisms for software and AI-enabled medical devices and strengthened transparency on how Yellow Card data is being used as an integral part of a lifecycle-based approach to keep patients safe.	Further details to be shared by Spring 2027

Recommendation 19: Sharing information on adverse incidents with the public	(Accept) The MHRA will work to introduce a public facing database/tool where members of the public can access and search for information on reports of adverse incidents for medical devices, including those that are software and AI-enabled. This database will build on the success of the Interactive Drug Analysis Profiles (iDAPs) ⁶ . MHRA will also explore potential legislative reform needed to enable such data sharing.	Further details to be shared by Spring 2027
Recommendation 20: Traceability and Monitoring	(Accept) The MHRA will publish guidance on the mandatory use of unique device identifiers (UDI) for medical devices as required in the pre-market statutory instrument. This guidance will include provisions for software and AI-enabled medical devices and will outline practical requirements for UDI and version control across updates. The MHRA will also work closely with DHSC and the devolved health departments to ensure such version control traceability is captured in the electronic patient record.	Further details to be shared by Spring 2027
Recommendation 21: Automating low burden reporting	(Accept) The MHRA will explore opportunities to automate low burden reporting requirements. This will include working with healthcare providers to explore how automated reporting can be embedded into integrated systems such as electronic patient records and with manufacturers to embed reporting into software and AI-enabled medical devices. The MHRA will also consider providing guidance on the types of reporting that can be automated, which will be supported with use case examples of what automated reporting looks like in practice.	Further details to be shared by Spring 2027

⁶ GOV UK [Check drug analysis profiles \(iDAPs\)](#).

Recommendation 22: Enhanced enforcement mechanisms	(Accept) The MHRA will explore additional approaches to further support communication and reporting of emerging issues and safety and quality signals in a timely, clear and accessible manner with partners across the health system. The MHRA will also explore options to strengthen its enforcement approach for software and AI-enabled medical devices. This will include, implementing the civil sanctions regime by the Medicines and Medical Devices Act, enabling the issuance of financial penalties where appropriate.	Further details to be shared by Spring 2027
Recommendation 23: Compliance pilots	(Accept) The MHRA will work with stakeholders to test how more responsive approaches to engagement and compliance could work in practice. The MHRA will also consider, when appropriate, how to best share information from these efforts with partners across the health system.	Further details to be shared by Spring 2027

Key Principle: Consider regulation and safe management of software and AI health technologies as a system-wide function, with each part of the system delivering its responsibilities so these technologies can be used effectively and responsibly.

Recommendation	Government Response	Delivery Timeline
Recommendation 24: Clear allocation of responsibility	(Accept) DHSC, MHRA, CQC and equivalent organisations in the devolved authorities, NHSR, PSA and the professional regulators will form a working group to ensure that responsibility for different actors across the lifecycle is clear and appropriate to their respective roles, while recognising each organisation’s statutory remit. This will include building on work led by the PSA and the Centres of Excellence for Regulatory Science and Innovation in AI (CERSI-AI) with professional regulators and Accredited Registers to clarify health and care professional responsibilities associated with the use of AI in healthcare. Their work involves a development of a shared set of high-level principles for regulators, to support the development of policy, guidance and standards. These build on existing principles and focus on clarifying new responsibilities professionals may be less familiar with, such as those arising from increased focus on post-market surveillance. The shared development of case studies may also contribute to this effort. The group will not determine or redistribute responsibilities. NHS England and NHSR will continue to progress guidance on indemnity to provide reassurance on how indemnity applies in this area.	Further details to be shared by Spring 2027

	DHSC and devolved health departments are actively considering governance best practices for AI tools used in healthcare to clarify accountability, including tools with increasing degrees of autonomy.	
Recommendation 25: Patient communication and redress	(Accept) DHSC will convene a working group with representatives across relevant healthcare organisations from England and the devolved authorities, as well as the Patient Safety Commissioner, to ensure that patients have access to redress if they receive a standard of care below what is reasonably expected. Relevant organisations who invite feedback on care, including CQC and equivalent organisations in the devolved authorities, will communicate with patients and service users to define what high quality care looks like in an AI-enabled system; how to report concerns and the most appropriate routes to redress, as software and AI-enabled medical products evolve. This could include considering what a shared model for complaints could look like.	Further details to be shared by Spring 2027
Recommendation 26: Operational conditions for safe deployment	(Accept) The MHRA will review and develop guidance setting out how manufacturers should identify and describe operational conditions necessary for safe use and consider how these should be reflected in future regulatory requirements.	Further details to be shared by Spring 2027
Recommendation 27: Sharing best practice examples of risk controls	(Accept) The MHRA will develop and include illustrative examples of best practice risk controls where appropriate, as part of planned guidance on operational conditions, which will form a key part of risk management practices for safe deployment and use.	Further details to be shared by Spring 2027

Recommendation 28: Contractual allocation of responsibilities	(Accept) DHSC and devolved health departments commit to reviewing and identifying opportunities to more clearly allow responsibility allocation between manufacturer and providers. This will include making sure providers have access to the necessary information to understand where responsibilities should practically sit. DHSC, working with the devolved health departments, will also consider how providers can ensure NHS contractual terms and risk management standards are consistently applied across the system.	Further details to be shared by Spring 2027
Recommendation 29: AI readiness Toolbox	(Accept) DHSC, working with the devolved health departments, will continue working to build on existing AI readiness work, including the AI Readiness Checklist delivered by CERSI AI and overseen by NHS England. This checklist will complement the Digital Technology Assessment Criteria, which currently functions as a national assurance gateway for AI and broader digital technologies. As part of this AI readiness work, DHSC and devolved health departments will also explore mechanisms to collect and share evolving practices across the system to provide access to learnings and resources to promote best practice.	Further details to be shared by Spring 2027
Recommendation 30: National leadership for AI readiness and adoption	(Accept) DHSC, working with the devolved health departments, will provide national direction to mitigate regional disparity in AI readiness, adoption and governance adequacy, including by bringing together existing networks such as the AI Ambassador's Network and Communities of Practice to disseminate national learnings and best practice.	Further details to be shared by Spring 2027

Recommendation 31: National best practice on lifecycle management	(Accept) DHSC, working with the devolved health departments, will explore and identify opportunities to provide national best practices on lifecycle management of software and AI-enabled technologies. For example, this will include strengthening guidance under DCB0129/0160, the clinical risk management standards that apply through a product's life and are currently under review. DHSC, in collaboration with MHRA, will also publish technology-specific guidance where additional clarity is needed to support safe adoption, using existing work on Ambient Voice Technology as a model for national guidance for other AI categories as they mature. This could align with relevant AI-specific guidance such as the statutory code of practice on AI and automated decision-making being developed by the Information Commissioners Office (ICO).	Further details to be shared by Spring 2027
Recommendation 32: Coordinated approach to training professionals and wider staff	(Accept) DHSC, working with the devolved health departments, will support a coordinated approach to improving AI literacy across the ecosystem, including working with healthcare professionals, procurement teams, professional regulators, those responsible for setting and delivering training and the wider public, building on existing activities including the NHS Digital Academy and AI Ambassador Network. This will aim to establish a shared baseline of general AI capabilities, with learning tailored to the roles, responsibilities, and contexts of different groups. For healthcare professionals and students, DHSC, working with the devolved health departments, will work with the relevant organisations to ensure that capability expectations are embedded within existing professional education.	Further details to be shared by Spring 2027

Recommendation 33: Provision of training by healthcare providers	(Accept) DHSC, working with the devolved health departments, will ensure that providers support their staff to meet existing training obligations set out by professional regulators and/or required by risk management standards. This will build on existing bespoke training material including that of the NHS Digital Academy, AI Ambassador Network, wider government resource and NHS-specific training courses, among others.	Further details to be shared by Spring 2027
Recommendation 34: Culture of reporting product experience	(Accept) DHSC and devolved health departments will work with the MHRA, CQC and equivalent organisations in the devolved authorities to strengthen the culture of reporting AI product experience across the lifecycle, for medical devices as well as relevant non-device AI tools. This should include working with partners to provide robust education, training and continuous development on reporting mechanisms for healthcare professionals. DHSC will work with the MHRA to ensure the Yellow Card Reporting scheme is fit for purpose for the implementation of AI at scale, including by increasing awareness for the tool through public campaigns and educational material and exploring how suppliers could also facilitate reporting requirements (e.g. standardised reporting templates).	Further details to be shared by Spring 2027

Key Principle: Design a framework that enhances trust, transparency and predictability for all stakeholders, to enable responsible development, adoption and use of safe and effective AI technologies

Recommendation	Government Response	Delivery Timeline
Recommendation 35: Patient transparency and trust	(Accept) DHSC, working with the devolved health departments, will consider how existing Information Governance requirements could be complemented to clarify transparency expectations where AI is used in patient care, what choices patients have, and how this applies to different functionalities. They will work together with NHSR where appropriate. This work will also consider when opt-out arrangements may be feasible and appropriate, particularly where AI functionality forms an integral part of a clinically approved medical device or care pathway, and how any such arrangements should be balanced against clinical effectiveness, equitable access, and patient safety considerations.	Further details to be shared by Spring 2027
Recommendation 36: Patient and public collaboration	(Accept) The MHRA will work with relevant partners to undertake a programme of patient and public engagement on topics relevant to software and AI-enabled devices. This programme will support public understanding and confidence in these technologies, ensure patient and public perspectives inform regulatory decision-making, and draw on professional regulators' insights where relevant.	Further details to be shared by Spring 2027

Recommendation 37: Public communication on safety concerns	(Accept) The MHRA will identify opportunities to engage with patient and public groups and other relevant partners to explore new mechanisms and approaches for providing public-facing safety information that are more tailored to meet their needs. The MHRA will also consider additional mechanisms to provide clear public facing information on AI-enabled medical devices. This will include working to publish a regularly updated list of authorised AI-enabled medical devices in the UK. As part of broader legislative reform for licensing, the MHRA will carefully consider how to ensure adequate information is summarised and available to the public on authorised software and AI-enabled medical devices.	Further details to be shared Spring 2027
Recommendation 38: AI sentiment analysis	(Accept) DHSC, working with the devolved health departments, will explore additional mechanisms to enhance its work to use patient, public and workforce insight to inform AI policy, including by considering whether existing survey methods could be expanded or complemented to collect wider sentiment around AI in healthcare.	Further details to be shared Spring 2027
Recommendation 39: User centred design and transparency	(Accept) The MHRA, in collaboration with relevant partners, will develop guidance on user-centred design for AI-enabled medical devices. This guidance will set out expectations for how developers and manufacturers should consider the needs and contexts of intended users when designing and deploying these products.	Further details to be shared Spring 2027
Recommendation 40: Comparative evidence to support adoption decisions	(Accept) NICE, and equivalent organisations in the devolved authorities, will continue to support adoption of the best AI technologies via setting out the evidence required for effective use	Further details to be shared Spring 2027

	along the access pathway, proportionate to the technology lifecycle and its intended use. DHSC, working with the devolved health departments, will complement this work by reducing duplication and improving predictability across the system's adoption pathways through the National HealthTech Access Programme, Innovator Passport, Accelerated National Innovation Adoption (ANIA) pathway and value-based procurement methodology.	
Recommendation 41: Lifecycle journey map for developers and manufacturers	(Accept) The MHRA will work with CQC, NICE, HRA equivalent organisations in devolved authorities, and the ICO to explore and consider options for providing more user-friendly and interactive public facing educational resources, that help articulate the regulatory and assurance journey across the product lifecycle. These resources could help signpost relevant tools and best practices, guidance and templates at each stage of the product lifecycle and help manufacturers identify the specific requirements they need to comply with. These resources could also help address commonly asked questions and signpost to early opportunities for innovators to engage with regulators.	Further details to be shared by Spring 2027
Recommendation 42: Early regulatory advice and pathway confirmation	(Accept) The MHRA will explore options to provide manufacturers of software and AI-enabled medical devices with earlier, more predictable regulatory advice and greater clarity on the applicable regulatory pathway. This will include considering structured mechanisms through which manufacturers can submit specific questions and receive written feedback within indicative timescales. For example, this could include	Further details to be shared by Spring 2027

	expansion of the existing Integrated Advice Service for medicines to medical devices ⁷ , where advice is provided from the MHRA and NICE.	
Recommendation 43: Formal classification confirmation service	(Accept) The MHRA will consider options for establishing a formal classification or pathway confirmation service for developers and manufacturers. This work will draw upon learnings from other international models and consider new regulatory powers as these are introduced.	Further details to be shared by Spring 2027
Recommendation 44: Post-market evaluation clarity	(Accept) The HRA will build on its existing guidance, including the “<i>Is my study research?</i>” decision tool and defining research table, to help organisations determine when AI development, deployment, evaluation and monitoring activities constitute research and require research governance, and when alternative governance arrangements apply. The HRA will review and update relevant guidance to address emerging evaluation approaches, including silent local evaluation and in silico methods. Additionally, it will seek to provide greater clarity on the boundary between non-research and research uses of AI, including when real-world evaluation or use of an AI tool may constitute research and require research governance. The guidance will recognise that classification depends on an activity's purpose, design and context rather than its label, and will signpost relevant data protection, confidentiality and information governance requirements. This work will align with MHRA and wider system guidance.	Further details to be shared by Spring 2027.

⁷ Medicines and Healthcare products Regulatory Agency (2026) [Request Integrated Scientific Advice: step by step.](#)

Annex A. List of Recommendations

Key Principle: Ensure regulatory requirements are proportionate to a device's risks and benefits, tailored to the evidentiary needs across the product lifecycle and support equitable access to safe and effective software and AI-enabled technologies.

Clear and Consistent Device Qualification and Classification

Recommendation 1. The MHRA should systematically review and update, as necessary, the existing UK Medical Devices Regulations to allow for a more tailored approach to software and AI-enabled medical devices. This should include:

- a. updates to the definition of a medical device to provide more clarity on when certain software and AI-enabled products are regulated as medical devices. For example, such updates should explicitly clarify when software intended for administrative purposes in a healthcare setting, general wellbeing purposes, and certain decision support purposes are not medical devices.
- b. a classification system and approach that considers the clinical risks, patient benefits, device lifecycle, and international harmonisation to provide a clear and resilient approach to classification of software and AI-enabled devices. This classification system should:
 - i. address the known limitations of self-declared Class I devices in the current framework.
 - ii. appropriately apply requirements and/or oversight for devices where there is sufficient risk to patients.
 - iii. employ a proportionate approach to low-risk devices.
- c. updates to the definition of intended purpose to ensure that device design and functionality are appropriately considered in addition to a manufacturer's claims and promotional materials.

Recommendation 2. The MHRA should systematically review and update, as necessary, guidance and other resources to provide clarity and predictability for the ecosystem. The MHRA should develop and implement an ongoing process for reviewing and updating guidance and other resources to ensure supporting policies are updated and available in a timely fashion. This should include:

- a. clarification of which types of software and AI-enabled products are, and are not, medical devices.
- b. clarification of the application of the risk-based classification to relevant software and AI-enabled devices.
- c. clarification of how intended purpose is interpreted for software and AI-enabled devices, including practical examples relevant to both understanding device qualification and post-market compliance.

Tailored and Risk Proportionate Regulatory Approaches

Recommendation 3. The MHRA should be empowered to improve the flexibility of its application of the regulatory framework, allowing for tailored approaches over time to allocate its resources and oversight to best assure the public health. For example:

- a. adding safety-related requirements for devices where new/available evidence shows a need for this beyond which is covered by existing classification requirements to support patient safety.
- b. reducing or limiting requirements and/or oversight for devices where the best available evidence shows that the risk is sufficiently low (for example, applying concepts such as enforcement discretion, where appropriate). Where such actions are proposed, it is essential that this be signalled clearly, applied consistently, and aligned to risk proportionality and patient benefit.

Recommendation 4. The MHRA should consider a function-based approach to the regulation of software and AI-enabled medical devices, recognising increasing product complexity and supporting proportionate regulation. It is critical that regulatory oversight be tailored to and focused on the medical device functionality, rather than any non-medical device functionalities present in products where there may be multiple functionalities.

Recommendation 5. The MHRA and relevant partners should consider how to appropriately balance oversight and evidence expectations across the product lifecycle. In particular, this should include clear identification of sufficient pre-market data and post-market surveillance for software and AI-enabled devices. This should be in line with the expectations and best interests of the public, clinical risks, technological characteristics, and the intended purpose of a given device. For example, the MHRA should provide clarity regarding what pre-market evidence is needed to support the safety and effectiveness of an adaptive AI-enabled device while considering the real-world data monitoring and reporting that will be needed to reduce overall uncertainty.

Recommendation 6. The MHRA should consider robust approaches to change management, including providing clarity on the scope and implementation of PCCPs for software and AI-enabled devices. This should include ensuring PCCPs provide a clear path and support for:

- a. adaptive and continuously learning AI-enabled devices by moving toward an approach that allows for clear boundaries and guardrails to define the allowable scope of change rather than requiring specific, planned changes.
- b. site or sub-population specific tuning.
- c. devices intended to operate with broad intended purposes or goal-versus-task defined intended purposes (for example, broad medical domain applicability, regulated AI agents).
- d. evidence-driven use expansion (for example, detecting/characterising additional radiology image findings to support clinician diagnosis or validating additional genes as they are characterised for an AI-driven diagnostic) that does not alter the device's intended purpose.

Recommendation 7. The MHRA should consider establishing an opt-in 'Master File' approach for general-purpose platforms, foundation models or underlying technologies used by medical device developers. These 'Master Files' can be used to support the assessment of medical devices leveraging such platforms or models. This should include:

- a. engagement with technical experts to explore the kinds of proportionate information that would provide long term value within a 'Master File', such as benchmarks, model cards, internal guardrails and firewalled mechanisms of providing additional information as available and necessary. The information should aim to support assurance that a given device leveraging the underlying technology will continue to function as intended.
- b. development of MHRA voluntary guidance setting out the process for a 'Master File' and the types of information a 'Master File' should contain for such underlying technologies to adequately support device developers and regulatory decision making.

Recommendation 8. The MHRA, DHSC and the devolved health departments should consider how to obtain the necessary information when a medical product has an underlying dependency on a general-purpose model. Specifically, the manufacturer should be expected to transparently report this dependency, any related risks and mitigations (including continuity plans). This should be an expectation as part of regulatory submissions (for AI-enabled medical devices) and through procurement and contract terms (for AI-enabled products).

Recommendation 9. The MHRA should consider issuing guidance to support developers in providing clear, transparent information to users and the public, such as through the use of model cards, careful user-centred interface design and dynamic labelling to support the safe use of software and AI-enabled devices.

Recommendation 10. The MHRA should consider providing clear guidance and educational resources on cybersecurity expectations across the product lifecycle for software and AI-enabled devices, which are increasingly connected across networks and with other devices. Cybersecurity is a critical factor of device safety as cybersecurity threats can directly compromise the safety or effectiveness of devices, including the potential to cause rapid harm to many people.

Recommendation 11. The MHRA should consider providing clarity through guidance on regulatory expectations for direct-to-consumer applications and wearables to fully harness their potential for the UK such as by supporting a shift towards preventative care. It is important that careful consideration be given to how to apply a proportionate approach to empower people with access to information about their health while understanding the unique benefits and risks of direct-to-consumer use.

Recommendation 12. The MHRA should consider providing clarity through guidance on human factors, usability considerations and expectations to support safe use as part of the design, development and testing of software and AI-enabled medical devices.

Recommendation 13. The MHRA should consider working with relevant partners to provide guidance for manufacturers on health equity expectations for software and AI-enabled medical devices across the product lifecycle, including consideration of the diversity of the intended population, with particular regard to underserved population groups.

Clear and Timely Paths to Market (With sufficient risk controls)

Recommendation 14. The MHRA should enable staged authorisations, supported by appropriate deployment pathways that support evidence generation and infrastructure to enable sustainable, long-term patient access to beneficial devices. These authorisation pathways should be:

- a. sufficiently flexible to allow tailoring of the authorisations and pathways to support those devices, where a staged authorisation is appropriate. This should include consideration of technological maturity, clinical needs and necessary risk controls to be put in place to ensure patient safety;
- b. transparently defined and reported to ensure patient awareness and support public confidence;

- c. designed to address the recognised limitations of single-timepoint assessment of software and AI-enabled devices, including demonstrating real world performance, managing uncertainty, and providing lifecycle assurance of device safety and benefit; and
- d. clearly defined processes and pathways from staged to full market authorisation to enable predictable, uninterrupted delivery of beneficial devices to patients.

Recommendation 15. The MHRA, in collaboration with system partners, should continue using regulatory sandboxes and extend their remit to support the development and testing of regulatory approaches, best practice and safe deployment pathways for novel technologies that provide a clear path to market authorisation.

Recommendation 16. The MHRA should continue its efforts to advance international harmonisation of the regulation of software and AI-enabled devices by:

- sharing lessons learned and best practices for proportionate, lifecycle-based regulation to align where possible and support responsible innovation for patients;
- introducing recognition or reliance pathways for relevant devices with trusted international regulators, where it is in the best interests of patients.

Proportionate post-market assurance

Recommendation 17. The MHRA should continue to tailor post-market surveillance requirements to be proportionate with a device's risks, benefits and related uncertainty to ensure effective, resource-efficient oversight. This should include a menu or set of mechanisms that manufacturers may be required to implement, including:

- a. post-market surveillance plans and other real-world data collection / monitoring plans assessed at the time of device authorisation to reduce uncertainty.
- b. real-world post-market studies where needed.
- c. more regular, ongoing performance reporting made available publicly and/or to regulators, as appropriate, to provide continued assurance of device safety and effectiveness.
- d. established processes for escalation where performance degradation is detected, even if a reportable incident has not occurred.

Recommendation 18. The MHRA, manufacturers and relevant healthcare organisations from the devolved authorities should work with partners across the healthcare ecosystem to review and refine mechanisms for reporting and sharing information (including best practice) across the ecosystem to ensure that they are clear, accessible, and, where appropriate,

standardised to enable more robust insight into ongoing performance and safety. This should support timely sharing of information from deployment sites and greater awareness of available reporting routes.

For example, the MHRA should consider how the Yellow Card scheme can be improved for AI and software-enabled devices, including for different users (such as patients) who may need to interact with the system. Likewise, manufacturers could support this process by focusing on user-centred design that enables real-time feedback, including the reporting of adverse incidents and malfunctions.

Recommendation 19. The MHRA should consider developing a database or tool that enables members of the public to easily search for adverse incidents related to a specific software or AI-enabled medical device. This should build on initiatives such as MHRA's interactive Drug Analysis Profiles (iDAPs) and include functionality that enables members of the public to search and filter by manufacturer, device name and timeframe, among other variables.

Infrastructure, monitoring and enforcement

Recommendation 20. The MHRA should introduce mechanisms to identify AI-enabled medical devices once deployed, including appropriate version control, to provide clear traceability throughout a device's lifecycle. For example, a unique device identifier (UDI) should be explored as a potential mechanism for traceability within the patient record, enabling healthcare professionals and providers to identify and audit the outcomes of patients whose care involved a particular device. The MHRA should also work closely with DHSC and the devolved health departments to support implementation, ensuring that device traceability information, including version, is routinely recorded as part of the patient record.

Recommendation 21. The MHRA should identify opportunities for manufacturers to automate low-burden reporting requirements, including by integrating AI into existing reporting mechanisms.

Recommendation 22. The MHRA should develop and introduce enhanced enforcement mechanisms that enable responsive and timely engagement with manufacturers of software and AI-enabled devices to address emerging issues and support patient access to safe, effective and high-quality devices. For example, this could include mechanisms such as:

- a. more timely and transparent sharing of early quality and safety signals for a given device or device type with providers, healthcare professionals and the public. It is important that these stakeholders have early awareness of potential device safety concerns that may alter the benefit-risk profile.

- b. the issuance of financial penalties or fines to penalise manufacturers that fail to comply with legal requirements and place patients at risk.

Recommendation 23. The MHRA should work with a range of stakeholders, such as manufacturers and Approved Bodies, to identify, test and refine how more responsive approaches to engagement will support a more holistic, lifecycle-based approach, driving and incentivising regulatory compliance, while assuring patients have access to safe, effective and high-quality devices.

Key Principle: Consider regulation and safe management of software and AI-enabled health technologies as a system-wide function, with each part of the system delivering its responsibilities so these technologies can be used effectively and responsibly

Clarity of roles and responsibilities

Recommendation 24. DHSC and relevant healthcare organisations from England and the devolved authorities should work together to ensure that distribution of responsibility at each stage of the AI product lifecycle is clear and appropriate to their respective roles.

Recommendation 25. DHSC should convene relevant healthcare organisations from England and the devolved authorities to work together to ensure that, as AI technology is adopted, patients have access to redress if they receive a standard of care that falls below what is reasonably expected.

Relevant organisations that invite feedback on care, including the CQC and equivalent organisations in the devolved authorities, should work with patients (or other service users) to define what high-quality care looks like in an AI-enabled health system; and clearly communicate how to report concerns and access redress if they believe the standard of care has fallen below what is reasonably expected.

System-wide commitment to delivering on risk controls with shared understanding of roles

Recommendation 26. Building on existing risk management expectations, the MHRA should require manufacturers to specify, in pre-market submissions and risk management files, the operational conditions needed for safe deployment and use. This includes the controls that must exist in the user environment – healthcare provider, health professional, or consumer context, as applicable – such as those related to cybersecurity, user training, and institutional AI readiness.

Recommendation 27. To support manufacturers, the MHRA should publish guidance, including redacted illustrative examples of risk controls that demonstrate good practice and support more consistent adoption of this mechanism.

Recommendation 28. The contract between the manufacturer and healthcare provider should include an explicit allocation of responsibilities for delivering required risk controls, including those specified in the risk management file, and for meeting relevant regulatory commitments.

A healthcare system that is equipped and empowered to safely deploy AI technologies

Recommendation 29. DHSC and the devolved health departments should co-develop an adaptable 'AI readiness' toolbox that can be tailored to the implementation and ongoing use of specific medical products. This would involve providers sharing their evolving approaches to meeting the operational requirements and risk controls necessary for the safe deployment and use. In return, they would benefit from access to a repository of peers' emerging practices. Such a resource should also help providers recognise their obligations and capabilities with regard to the safe deployment of AI technologies, including:

- a. an ability to self-assess their 'readiness' to deploy a particular AI product.
- b. an ability to deliver the risk controls necessary for a particular AI product.
- c. an ability to articulate the governance structures in place to safely deploy AI products.

Recommendation 30. DHSC and the devolved health departments should provide the national direction to minimise local disparities in AI readiness, adoption processes and governance approaches. To this end, they should:

- a. define national best practice and expectations for organisational capabilities in AI readiness, including any baseline governance expectations, to ensure providers are equipped to enhance their AI readiness consistently across the system.
- b. establish a national learning function (or observatory) to collect learning from across the system, ensuring that implementation successes and challenges are shared across trusts to minimise siloed working. This could include centralising and maintaining the AI readiness toolbox for providers to access and contribute to.

Recommendation 31. DHSC and the devolved health departments should provide national best practice on the lifecycle management of emerging AI-enabled technologies (for example, agentic systems), including procurement, implementation, monitoring, reporting and end-of-life management, to support providers across the ecosystem.

Recommendation 32. DHSC and the devolved health departments, working with professional regulators, those responsible for setting and delivering education and training for the healthcare professions (including the medical Royal Colleges), and healthcare education providers, should establish a coordinated approach to building professional capability in the safe and effective use of AI in healthcare. This should span initial education, postgraduate training (where relevant) and continuing professional development, and should complement setting-specific training on technologies. Expectations should be expressed as defined capabilities so that they remain applicable across the four nations and adaptable to the differing needs of professional groups and providers.

Recommendation 33. Healthcare providers must ensure that professionals and their wider staff are well-equipped to use these technologies by undertaking appropriate training, including technology-specific training aligned with the AI technologies (including those with automated decision-making functionality) used in their practice.

Recommendation 34. DHSC and the devolved health departments should promote a strong culture of reporting product experience, including adverse outcomes and product malfunctions, throughout the product lifecycle to support continuous improvement and ongoing awareness of product performance in real-world settings. This should include robust education, training and continuing professional development on reporting mechanisms and expectations for healthcare professionals as part of their duty of care.

Key Principle: Design a framework that enhances trust, transparency and predictability for all stakeholders to enable responsible development, adoption and use of safe and effective AI technologies.

Transparency and engagement as a foundation for public trust

Recommendation 35. Relevant healthcare organisations from England and the devolved authorities should adopt a proportionate, system-level approach to transparency about the use of AI-enabled products in patient care, including addressing patients' reasonable expectation of being informed of the use of such technologies in their care and having the ability to opt out where possible or appropriate.

Recommendation 36. The MHRA and relevant healthcare organisations from England and the devolved authorities should enhance mechanisms for ongoing inclusion of patient and public perspectives to support a learning regulatory system. These mechanisms include a Patient Engagement Advisory Committee, workshops and deliberative events.

Recommendation 37. The MHRA should work with patient and public groups and other relevant partners to develop robust mechanisms for providing public-facing safety information that is clear, transparent and tailored to the needs of the public.

Recommendation 38. DHSC and the devolved health departments should consider implementing a periodic ‘AI sentiment census’ that includes both patients and professionals, so ongoing policymaking is informed by both patient and workforce experience and perspectives.

Recommendation 39. The MHRA should work with relevant partners to develop guidance for developers on user-centred design considerations that support transparency for device users, including patients, professionals and the public, and enable simple and intuitive reporting of user feedback on the device experience.

Recommendation 40. DHSC and the devolved health departments should work with NICE and equivalent organisations in devolved authorities, to consider ways to provide adopting organisations with additional information regarding clinical effectiveness and value to support decisions on whether to adopt an AI health tool and where appropriate, which product best fits their needs.

Regulatory clarity and predictability to support the innovation journey

Recommendation 41. The MHRA, CQC, NICE, HRA, ICO, and equivalent organisations in devolved authorities, and relevant organisations from the devolved authorities should develop clear, accessible, public-facing educational resources and tools within their scope of activities, with the overall intent of making it easier for developers and manufacturers to navigate the regulatory and assurance journey for AI health technologies. These resources should demystify the journey and support new developers by reducing ambiguity across the lifecycle.

Recommendation 42. The MHRA should introduce clearer, earlier and more predictable engagement mechanisms for developers and manufacturers, in collaboration with relevant partners. These mechanisms should include a structured route that provides written feedback (and, where necessary, meetings) on topics such as device qualification, classification, appropriate path to market, evidence requirements and post-market expectations.

Recommendation 43. The MHRA should also consider offering a formal ‘classification confirmation’ service, enabling manufacturers to submit a short package and receive a written determination on a product’s regulatory status, reducing avoidable delay and rework.

Recommendation 44. As AI-enabled technologies increasingly rely on post-market monitoring and evaluation, the HRA, and equivalent organisations in devolved authorities, should build on its existing guidance and decision-support tools to help organisations determine when development, deployment, evaluation and monitoring activities constitute research and require research governance, and where they may sit outside of research, requiring governance and oversight through other mechanisms. This work should be aligned

with MHRA and wider system guidance as the regulatory framework develops.

