ETHOS digital health

National review · Clinical risk management standards

DCB 0129 and DCB 0160: my response to the national review

Stuart Harrison · Founder & CEO, ETHOS Digital Health · Author of DCB 0129 and DCB 0160

NHS England is reviewing the two clinical risk management standards that are mandatory across health and social care: DCB 0129 and DCB 0160. I authored them. This is my full response to that review, published in the open.

These regulations were developed to ensure that all health information technologies undergo a clinical risk assessment before being used in patient care. DCB 0129 addresses technology manufacturers' needs, while DCB 0160 provides guidelines for Health & Social Care organisations on implementing the safe deployment and use of these systems. Since 2009, both documents have had an impact. The review will now determine whether these two standards apply to today's software development, releases, and changes.

My position on this matter is quite brief. Modernise the model that underpins the standards while keeping the same standards; i.e., keep the one-off safety case signed at go-live as a standard, but replace it with continuous, through-life assurance of the regulated software. Remove deployment from the release process. Generate evidence over time (i.e., as software matures). Plan to make changes ahead of time; this is what regulated software development is moving toward now, with the MHRA's Predetermined Change Control Plans, and this is how many organisations develop their products in a progressive delivery manner.
I am posting my complete response here so that responses about such important topics are visible and not left in a blank survey field. If you use DCB 0129 or DCB 0160, I recommend responding to this review also.

Response submitted to the NHS England public consultation. The consultation closes on 11 September 2026.

The standards

Scope

Q5. Should alternative risk management standards replace the current ones?

My answer: No.

The current standards should be retained and modernised, not replaced. Their underlying discipline: hazard identification, risk evaluation, risk control, and residual-risk acceptability is the same as that expressed in ISO 14971, the harmonised medical-device standards, and the revised standards, and the revised standards should explicitly state that alignment of principle. What DCB 0129 and DCB 0160 add, and what a device-centred standard does not reach, is the deployment, configuration, integration and use context, and the treatment of the clinical safety case as a sociotechnical artefact

rather than a device file. Replacing them would lose that context. Align in principle; retain the distinct scope.

Q6. Should DCB 0129 and DCB 0160 apply to medical devices?

My answer: Yes, in a defined and limited sense.

Where a medical device includes software deployed within an NHS or social care pathway, the act of deploying, configuring and using it in that pathway creates clinical risk that the device's own conformity file does not address integration with other systems, local configuration, workflow, and the conditions of use. DCB 0160 should apply to that deployment. The standards must not replicate nor replace the UK Medical Device Regulation (MDR) compliance pathway for the medical device itself. A proper relationship exists as a linkage: the standards ensure that the clinical safety of the medical device will be maintained when it has been appropriately deployed, and that they reference the device's regulatory status rather than evaluating/assessing the device again.

Q7. Should the standards include specific cybersecurity requirements?

My answer: Yes, as a clinical-safety interface, not a cyber framework.

The standards should require that cybersecurity risks with a genuine path to patient harm, loss of availability, integrity, or confidentiality that affects care, be identified and assessed within the clinical risk management process. The resulting clinical hazards can then be controlled and documented like any other. The technical controls will be developed within several established frameworks. These include the Data Security and Protection Toolkit (DSPT), Cyber Essentials, National Cyber Security Centre (NCSC) guidance and the Department of Health's Digital Technology Advisory Committee (DTAC) security criteria. The standards need to describe how they interface with these frameworks and incorporate the clinical safety implications of cyber risks into the safety case, rather than simply replicating them.

Terminology

Q8. Do specific terms need clarification or updating?

My answer: Yes, and one distinction should be added.

> **Deployment and release should be defined as distinct events.** Modern delivery decouples the two: software can be deployed to production behind a control and not yet released to users. The standards currently read as if deployment and release are a single moment, which does not align with progressive delivery and blocks the continuous-assurance model that the revised standards should support.
>
> **Hazard, hazardous situation and harm** should be aligned to the ISO 14971 vocabulary, so the shared discipline is visible in the language.
>
> **Clinical risk** should be defined as the risk of harm to patients or service users arising from the health IT system.
>
> **Health IT system** and the boundaries of a system should be defined to accommodate interconnected and third-party components.
>
> The **accountability terms** around the clinical safety officer role should distinguish accountable clinical sign-off from the supporting safety-engineering analysis.

Q9. Should the word ‘clinical’ remain or be replaced?

My answer: Remain.

The word defines the scope of the standards; the risk being managed is harm to patients and service users, not general IT, information, or project risk. Replacing it would blur that boundary and invite the standards to be read as generic technology-risk management. Tighten the definition of clinical risk so the word carries its scope precisely but keep the word.

Risk-based and proportionate application

Q10. Should requirements differ by the risk profile of the technology?

My answer: Yes.

The rigour of the standards should scale with the technology's risk class and demonstrated safety maturity, rather than applying uniformly. A low-risk product and a high-risk clinical decision support system should not undergo the same set of activities. This is not a lowering of the standard; it is proportionate application of it, and it is already the direction of NHS England’s Clinical Safety Management System, under which a product’s clinical safety maturity determines the assurance route. The standards should provide a method that makes proportionality auditable, a risk-based determination of which requirements apply and a defined maturity progression against which assurance accrues, so that policy and process sit atop a method the standards themselves own.

Clinical safety officers

Q11. Should the requirement for a CSO remain?

My answer: Yes.

The determination that a system is safe to use in patient care is a clinical judgement and requires a registered professional to be accountable for it. Removing the role would remove the point of clinical accountability that gives the safety case its authority.

Q12. Does the CSO role need further definition?

My answer: Yes.

The role should be redefined around continuous, multi-disciplinary assurance rather than a point-in-time approval at go-live. As delivery becomes continuous and change becomes pre-agreed and incremental, the CSO’s function is to hold the safety argument and the evidence level for change, not to approve a single gateway once. The definition should also separate what only a registered clinician can do, accept residual clinical risk on behalf of patients, from the safety-engineering analysis that supports that decision.

Q13. Should the standards define a supporting role and clearer related responsibilities?

My answer: Yes.

The standards should formally recognise a safety-engineering role beneath the accountable CSO and name the responsibilities of the other contributors to clinical safety. The work of clinical safety-hazard identification, control specification, evidence assembly, and safety case development is a discipline that does not require clinical registration, while accountability for accepting residual risk does. Recognising that this approach ensures safety in practice increases the likelihood that engineers can complete their tasks within available capacity during a time of scarcity. This enables

accountability across the entire delivery team, rather than relying on the CSO alone to deliver all aspects of safety.

Documentation and lifecycle

Q14. Do the mandated safety deliverables need revision?

My answer: Yes.

The deliverables should support continuous assurance rather than a single pre-deployment safety case. The safety case should be able to declare in advance the scope of change it tolerates and the evidence level each change must meet before release, a pre-agreed change plan rather than a fresh case for each change. The hazard log and safety case should be living records that continuously accrue evidence, not documents finalised at go-live. A record of an auditable release should document each release of a product, based on the evidence available at that time. The MHRA's draft regulations for software as a medical device incorporate Predetermined Change Control Plans, which would be mandatory under UK law from 2026, but do not weaken the evidential requirements.

Q15. Should the standards accommodate agile and modern lifecycle practices?

My answer: Yes, this is the most important modernisation.

The standards should explicitly accommodate agile and progressive delivery, in which deployment is decoupled from release. Under progressive delivery, software is deployed to production behind a control (a feature flag) and released to users only when a pre-agreed evidence bar is met. Read through a clinical safety lens, the feature flag becomes a named safety control, the release decision becomes a standing, pre-agreed evidence bar, and safety evidence accrues continuously upstream rather than at a single assurance gateway.

The feature flag approach has already been successfully implemented nationally as part of U.K. Government Services and was documented by Hamill in 2019. Industry statistics also confirm that the use of flags to deliver features has become nearly ubiquitous, with elite delivery teams deploying many times per day and achieving very low failure rates for changes made via flags (LaunchDarkly 2024; DORA 2024). Additionally, the approach aligns with the body of work on dynamic-through-life-safety-cases (Denney, Pai and Habli 2015) and with the future regulatory environment for predetermined change control plans in the draft MHRA 2026 regulations. Therefore, the new standard should support an approach based on a series of deploy-and-approve rather than relying on a single event.

Q16. Are all phases of the product lifecycle sufficiently addressed?

My answer: No.

Development and initial deployment have been well addressed. The post-deployment process of making changes, continuously monitoring, and eventually decommissioning has not been; in fact, it appears that most of the safety work will be done by the time a system goes live. What is needed is a "through life" approach to managing systems after they are deployed; this would include pre-determined rules for how changes can occur after deployment, ongoing safety monitoring (with safety-related information being fed-back into the current version of the living Safety Case), and specific consideration to decommissioning and replacing a system (including the clinical/medical/safety and data risk associated with exporting data and during the cutover/cut over process) that are also relevant to Cyber Security.

Existing requirements

Q17. Which requirements should be amended?

By subject:

> The safety case report and hazard log requirements, so both are living records that accrue evidence continuously and can express a pre-agreed change envelope, rather than documents completed before deployment.
>
> Deployment and go-live requirements, to distinguish deployment from release and to permit release against existing evidence levels.
>
> CSO requirement, to define the role around continuous assurance and separate accountable clinical approval from supporting safety-engineering analysis.
>
> The risk evaluation requirements, to make application proportionate to risk class and demonstrated maturity.

Q18. Are there requirements that should be removed?

My answer: No.

The gaps in the standards are matters of what is missing or needs modernising, not of redundant requirements; reform rather than deletion. Removing requirements would weaken the evidential floor at the point where the direction of travel calls for it to be extended through-life.

Q19. What new requirements should be added?

> **Through-life change control.** If a design for the product is expected to go through some form of change that will be done on a continuous or incremental basis, the safety case needs to define the "change window" or the acceptable parameters of allowable changes prior to them being implemented. In addition, it must also determine the level of supporting evidence required by each proposed change to allow its implementation before deployment. The records from the implementation of these changes must be auditable.
>
> **Transferable and cross-system hazard ownership.** Hazards are owned by a single, accountable entity that is responsible for each of these hazards across organisational boundaries.
>
> **Continuous safety monitoring.** Defined post-deployment monitoring that feeds evidence back into the living safety case.
>
> **Cyber-safety interface.** Cybersecurity risks with a credible pathway to patient harm are carried into the clinical risk management process.
>
> **Equity of safety.** Safety assessment considers whether the technology performs safely across the diverse patient groups it will serve.

Implementation and wider system

Q20. Aligning the standards to the broader patient safety effort.

The standards and their guidance should explicitly connect to the wider patient safety system rather than run alongside it. Digital clinical safety is currently treated as a separate workstream from patient safety teams' work, which weakens both. Clinical safety officers should have a defined relationship to the organisation’s patient safety function and its governance. The standards should reference the national patient safety architecture, the Patient Safety Incident Response Framework and the Learn from Patient Safety Events service, so that safety events arising from health IT are captured and learned from through the same channels as other patient safety events. And the CSO

role should be positioned as part of the organisation's safety culture, not an isolated compliance function.

Q21. Alternative pathways for non-registered professionals to become CSOs.

My answer: Retain registered clinical accountability for sign-off; formally recognise a safety-engineering role beneath it.

The determination that a system is safe to use in care and the acceptance of residual clinical risk on behalf of patients are clinical judgements that should remain with a registered health or care professional. The safety-related aspects of this judgment, including the development of a risk assessment (hazard analysis), description of controls, and collection of relevant evidence, are an engineering-based discipline in which a registered professional is not required. Therefore, the analytical aspects of this position should be made available, on a competency basis, to both registered and unregistered engineers with experience.

Q22. Minimum qualification requirements for CSOs.

The accountable clinical sign-off role should require current UK registration, patient safety training and demonstrated clinical safety experience. The supporting safety-engineering tier should require the technology, patient safety, and clinical safety competences, as well as continuing professional development, but not necessarily registration. Formal risk-management qualification is useful but should not be mandatory in its own right; demonstrated competence matters more than a specific certificate.

Q23. Support CSOs need to assure cybersecurity.

Consistent with the interface position at Q7, the CSO does not need to become a cybersecurity specialist. The CSO needs the means to bring cyber risk into the clinical safety assessment, including working with the people who own the technical controls, access to cyber expertise, engagement with the digital and cyber teams, and templates that make the clinical-safety consequences of cyber risk straightforward to capture. Clearer regulatory guidance on the boundary between clinical safety and cyber assurance would reduce duplicated or missed effort.

Q24. Other ongoing support CSOs need most.

The two support needs that most affect safety in practice are protected time and senior leadership backing. Where clinical safety is done in the margins of another role, and without visible leadership ownership, the analysis is thinner and the culture weaker. Peer networks and shared templates raise the floor across organisations, which varies widely today.

Q25. Extending scope to interoperability and third-party integration.

My answer: Yes.

Interoperability and integration with third-party systems are where a growing share of clinical risk now sits, and the standards should extend to cover the clinical safety of those interfaces. When systems exchange data or depend on one another, hazards arise that no single system's safety case captures, a change in one system creating a hazard in another, or data losing meaning as it crosses a boundary. Extending the scope of the integration itself, as a clinical-safety concern, closes a gap that currently exists between organisations and suppliers.

Q26. Extending scope to organisational accountability for transferable risk.

My answer: Yes, the requirement most clearly missing today.

If hazards can be transferred from one system to another (and so on), if they are created in one and realised in another, if nobody has responsibility for them now and they fall outside of the current

definitions of what Safety Cases include. The Standards should demand that all transferable and cross-boundary risks be formally identified and given a name by which an accountable owner(s) can be designated within and beyond their organisations and suppliers. This would represent the only completely "new" element required for this digital era: assigning hazard ownership rather than liability based on a physical system boundary. While it will take some care to ensure it is implementable, the principle of assigning accountability in practice (not just stating) should be established.

Q27. How legacy systems should be managed.

My answer: Use a risk-based approach to determine which standards apply.

Legacy systems should be subject to risk-based determinations rather than treated the same as new systems or exempted. The right level of requirement is determined by the clinical risk the system poses and the evidence of its performance. Long, safe in-use history is itself evidence and should be usable to justify a proportionate rather than full re-assessment, a proven-in-use argument, made explicitly and recorded, rather than an unexamined assumption that an old system is safe because it has not yet failed. This is the same proportionate, maturity-based logic as Q10, applied to systems already in service.

Q28. Extending implementation guidance to AI-enabled technology.

My answer: Yes, but not as a separate silo.

The defining challenge of AI in this context, systems that change after deployment, is a specific instance of the general problem the revised standards should already solve through pre-agreed, through-life change control. An AI model update is change within a defined envelope against a standing evidence bar, the same mechanism as any other post-deployment change, with AI-specific evidence requirements. Building the general continuous-assurance model first and treating AI as a case of it is more durable than writing AI-specific rules that date as the technology evolves. This is consistent with the MHRA's approach, where Predetermined Change Control Plans apply to software as a medical device generally, not only to AI.

Q29. Integrating with existing cybersecurity frameworks.

The integration needs to occur by reference and proportion, rather than by requiring a single universal certification. If cyber risk is to be included in the clinical safety assessment and aligned with established frameworks (the NCSC, the DTAC security criteria, DSPT status), this will incorporate the clinical/safety implications of cyber risk into that process but will not duplicate the cyber regulatory regime. However, if mandatory certification (such as ISO 27001 or Cyber Essentials Plus) is required for all products, regardless of risk level, it will disproportionately affect small suppliers, who may gain an unequal measure of additional safety from those costs; therefore, certification should scale with risk.

Q30. Which cybersecurity risks require clinical assurance?

Loss of availability will disrupt health services; loss of data integrity will provide clinicians with incorrect or misleading information; loss of confidentiality will prevent patients from presenting for treatment or seeking medical help. In addition to extended periods of unavailability (e.g., more than two hours), there are disruptions caused by supply chain disruptions and third-party integrations within care pathways. Clinical assurance is required if a cyber risk has a reasonable chance of causing patient harm through an identifiable pathway. Clinical assurance is tied to the identified pathway of potential patient harm, rather than to a specific vulnerability in technology.

Q31. Assuring cybersecurity across the full lifecycle.

The same "through-life" approach to data security should be applied to cybersecurity as it is to clinical safety. That means ensuring the data remains secure throughout the entire lifecycle, from import when first made live until export at decommissioning. Ongoing assessment of suppliers and systems supporting the application for assurance; planned downtime for patches; planned decommissioning before a vendor withdraws support; and proper access management will ensure the system remains safe throughout its use.

Q32. Enabling innovation, particularly for smaller companies.

The proportionate, maturity-based approach at Q10 is in fact the primary response; a smaller company creating a low-risk product will need to meet a proportional subset of requirements, not the complete suite designed for higher-risk products. The cost reduction associated with meeting clinical safety regulatory requirements will be considerable by reusing similar patterns (blueprints) rather than repeating everything each time. Smaller organisations have the added benefit of releasing their products at a very fast rate. In addition, sharing clinical safety templates and work samples, as well as clinical safety subject-matter experts across organisations, could help smaller organisations become compliant without requiring them to develop those competencies internally.

Q33. Sector-specific implementation challenges.

Each sector faces the same realities when implementing solutions that generic guidelines miss. In general, primary care and community services implement supplier-configured systems with little or no local safety capacity; therefore, the generic guidance will need to support and rely on suppliers' assurances for use by those without specialist knowledge. In secondary care, there is an estate of highly interconnected entities that are most severely affected by the transferable hazard problem. The social care sector has the least developed digital safety infrastructure and needs access through the most easily available route. In ambulance and emergency services, availability and failover are the two most significant hazards when operating in time-critical situations. A series of sector-specific implementation guides, using the same methodology, will serve each of these areas better than a single generic implementation guide.

Q34. Supporting smaller organisations with limited resources.

Smaller organisations want the process to be available in the same way as larger ones; but they don't want the 'entry' requirements of the process reduced. They would like to have shared (or templated) example(s) of how to do things, with appropriate entry points based on their level of risk. They would also like access to other people's pooled or peer-based clinical safety expertise (to avoid duplication of effort) and, when possible, use the suppliers' assurance processes (if they are using supplier-configured systems).

Q35. Applying the standards consistently across systems.

Where the same technology is deployed across many organisations, the safety work should be done once as a reusable pattern and instantiated locally, rather than repeated from scratch each time. A supplier or lead organisation can establish the core safety case and hazard analysis for the product; each deploying organisation then ensures the local deployment context, configuration, integration, and workflow align with that baseline. This is both more efficient and safer, because the shared core improves with each deployment. The standards should explicitly recognise and enable this pattern-based, deploy-many model.

Q36. Should safety assessments include accessibility considerations?

My answer: Yes, as mandatory requirements.

Equity of safety is a clinical safety consideration, not a separate concern. A technology that performs less safely for some patient groups than others. The pulse oximeter example is a well-known case where it is less safe, full stop, and that belongs inside the safety assessment rather than alongside it. Making it a mandatory consideration, rather than a matter of guidance, ensures it is assessed. The requirement should be to consider and evidence safe performance across the diverse patient groups the technology will serve, proportionate to the risk and to the relevance of group differences to the specific technology.

Q37. Informing and engaging patients and the public.

Patients and the public should be able to understand, in accessible terms, that digital health technologies are subject to clinical safety assurance and what that provides, without being asked to engage with the technical detail of the standards. Transparency about the existence and purpose of the assurance builds warranted trust. Meaningful engagement is better placed at the level of specific services and their safety, and through existing patient and public involvement structures, than in the technical standards themselves. Publishing plain-language information on what clinical safety assurance means and involving patient representatives in the governance of the standards would be proportionate and effective.

Q38. Support to help organisations achieve compliance.

Shared templates and standardised tools do the most to improve consistency, because much of the observed variation comes from organisations building the same artefacts from scratch in different ways. Training, peer mentoring and sector-specific guides address the capability and context gaps. Certification of products or suppliers can help, but should be approached carefully and proportionately, as it carries costs that fall unevenly on smaller suppliers.

Q39. Applying the standards more consistently and strengthening oversight.

The change that would most improve consistent application is a shared, method-level definition of how clinical safety is done, not just what must be produced, but a reusable pattern for producing it, supported by common templates and worked examples. Consistency comes from a shared method more than from added oversight. On oversight, a maturity-based approach, of the kind the Clinical Safety Management System is introducing, is the right direction: assess where an organisation or product sits and target support and scrutiny accordingly, rather than applying uniform checking. Over time, published maturity and clearer accountability for transferable risk would strengthen compliance without heavy-handed enforcement.

Q40. Regular review and updates.

The standards should be reviewed on a defined, regular cycle. They should also include a mechanism to track regulatory and technological developments between full reviews, so they do not fall behind. The pace of change in AI, progressive delivery, and the surrounding regulation, including the MHRA draft 2026 regulations, means a static standard quickly becomes outdated. A standing review mechanism, with a lighter-touch update process for keeping alignment current between major reviews, would keep the standards live rather than periodically obsolete.

Q41. What else could ensure patient safety in the digital age?

The organising theme of the digital era is through-life continuous assurance: the safety case as a continuing, evolving justification for the scope of permissible change, while continually gathering evidence to support it, rather than a one-time document created before "go live". All other items follow directly from this, separating deployment and release, using the feature flag as a safety-

related control and the release determination as a minimum level of standing evidence; applying proportionalism and maturity-based approaches to the applications; using AI as a means to agree upon change, as opposed to AI being used as a separate silo; and, ultimately, owning the hazards that can be transferred, or shared among systems. This approach will fit into the overall U.K. Regulatory Trajectory, as identified in both the Draft Regulations, to be implemented by 2026, and in the MHRA's predetermined change control plans. This is also supported by current research on through-life safety cases (Denney et al., 2015), and by those currently practised in the field of progressive delivery (DORA, 2024; Hamill, 2019).

The new standards provide an opportunity to specify the method, i.e. the "how", by which the Clinical Safety Management System, as policy and process, will evaluate compliance. Incorporating that method into the standards themselves (i.e. defining commonality in practice) is likely the most beneficial singular contribution this review may make.

That said, there is another, much tougher question that lies beneath the surface of this review. That is a question that should be asked openly and honestly. When we consider that clinical safety management standards are mandated throughout healthcare and social care, and when those standards carry actual patient safety implications, such standards deserve to be owned/stewarded independently, with significant amounts of current expert knowledge/practice experience drawn from practitioners (clinicians, suppliers, etc.) in collaboration with others within the NHS.