SMART CODE, SAFE CARE:

Navigating AI Acceleration in Regulated Medical Devices

IMed Consultancy

&

CONTENTS

Management Summary

Whether you are a university spinoff with an innovative algorithm that could improve the lives of patients, or an established corporate integrating a new AI component into an existing physical device to improve speed, accuracy and save clinician time, a number of considerations have to be made to ensure that the device can enter its target market (and stay there) safely. These considerations depend on how different digital components relate to the device's functionality and also, critically, how each market regulates Software as a Medical Device (SaMD) and AI-powered medical technology.

International regulators put patient safety at the top of the agenda but approach the issue in different ways. In addition to this, the compliance map keeps changing, keeping manufacturers on their toes. In addition to keeping up with a changing rulebook, digital medical device manufacturers must also integrate new procedures and skills into their team. From managing data quality and privacy to preventing cybersecurity issues, through to understanding which geographies require which ISO certifications and assessments, the digital medical device market is quite literally learning on the job.

As new professional roles that stand at the cross-roads between software development and medical device manufacture start to take shape, this white paper combines the knowledge of medical sector software development teams at Firefinch Software with medical regulatory specialist expertise from IMed Consultancy to support both established businesses and startups navigate this exciting new phase of the digital (r)evolution.

Introduction

Artificial intelligence is reshaping the global healthcare technology market at unprecedented speed. In 2024, the global AI in healthcare market size was valued at USD 29.01 billion: it is now projected to grow from USD 39.25 billion in 2025 to USD 504.17 billion by 2032 at a CAGR of 44.0%.[1] Not only are more and more medical device manufacturers integrating software or AI components in their solutions, but AI-enabled software now represents one of the most promising pathways to improving diagnostics, treatment accuracy, operational efficiency, and ultimately patient outcomes.

Computational capability and cloud infrastructure have matured to the point where complex models can run efficiently, securely, and cost-effectively. AI tools too have finally reached a level of real-world performance that makes clinical value demonstrable: faster radiology reads, more reliable decision-support processes, personalised treatment insights, and continuous patient monitoring can be reliably delivered through smart devices. All this couldn't come at a better time as the healthcare sector is under enormous pressure. Clinician shortages, backlogs and rising costs impede access to care for millions across the globe making automation, better decision support, and workflow optimisation an urgent requirement.

While technology has reached a good level of maturity, its application to the highly regulated world of healthcare is still in its infancy: the EU AI Act for example, is not solely focused on medical AI but covers healthcare among a range of uses of the technology. The act itself also only came into force on 1st August 2024, and even then, with a phased implementation.

Another legislation, the EU Data Act, in force from September 2025, similarly imposes on manufacturers the need to share control over data generated by smart devices, such as digital health trackers, with users and businesses and sets conditions to regulate data access that promote greater user control over personal information. The opportunity and the risk posed by AI and SaMD still need to be finely balanced: misdiagnoses, incorrect dosages or hacked drug delivery systems pose a tangible danger to patient health, so safety must come first. At the same time, the risk for businesses is also high: the wrong regulatory approach can delay market entry for months or years, inflate development costs, or even force a device withdrawal.

The variability of international regulation for AI and SaMDs adds another layer of complexity: while the UK, US and EU all agree on putting patient safety at the top of their concerns, their pathways, classifications, and evidence requirements diverge.

This environment places unprecedented demands on organisations developing digital medical devices. Teams must develop or acquire new cross-disciplinary skills that bridge software engineering, product development, clinical risk management, cybersecurity, data science, and regulatory compliance. They must approach from scratch new ways to manage data ethically, implement robust AI governance, navigate ISO requirements, validate systems that behave probabilistically, and adapt to shifting regulatory norms across different geographies.

Are you really dealing with a SaMD?

The first step to ensuring that your digital-native device or integration is compliant and poised for success, is determining whether its software or AI features actually make it a SaMD or not. Simple tools that add up a series of numbers as scores, for example, are often not considered to be medical devices at all. On the other hand, when academic spin outs and startups develop an algorithm, have had some access to patient data that initially validates it and want to use it more broadly, the device is typically classified as a SaMD.

Cases like the above are clean-cut, but most are far more murky. Take Ambient Voice Technology (AVT) for example, a simple transcription service that can quickly develop into a more sophisticated platform capable of summarisation, clinical decision support, and more. Depending on its degree of functionality and intended use, an AVT may well cross over into medical device territory.

To help provide some clarity, the UK Medicines and Healthcare products Regulatory Agency (MHRA) which regulates medicines, medical devices and blood components for transfusion in the UK, has provided a new critical piece of guidance that offers an illustrative starting point for determining whether a device is a SaMD. Although designed for the UK market, its clear examples and definitions offer a solid foundation for devices entering any geography.

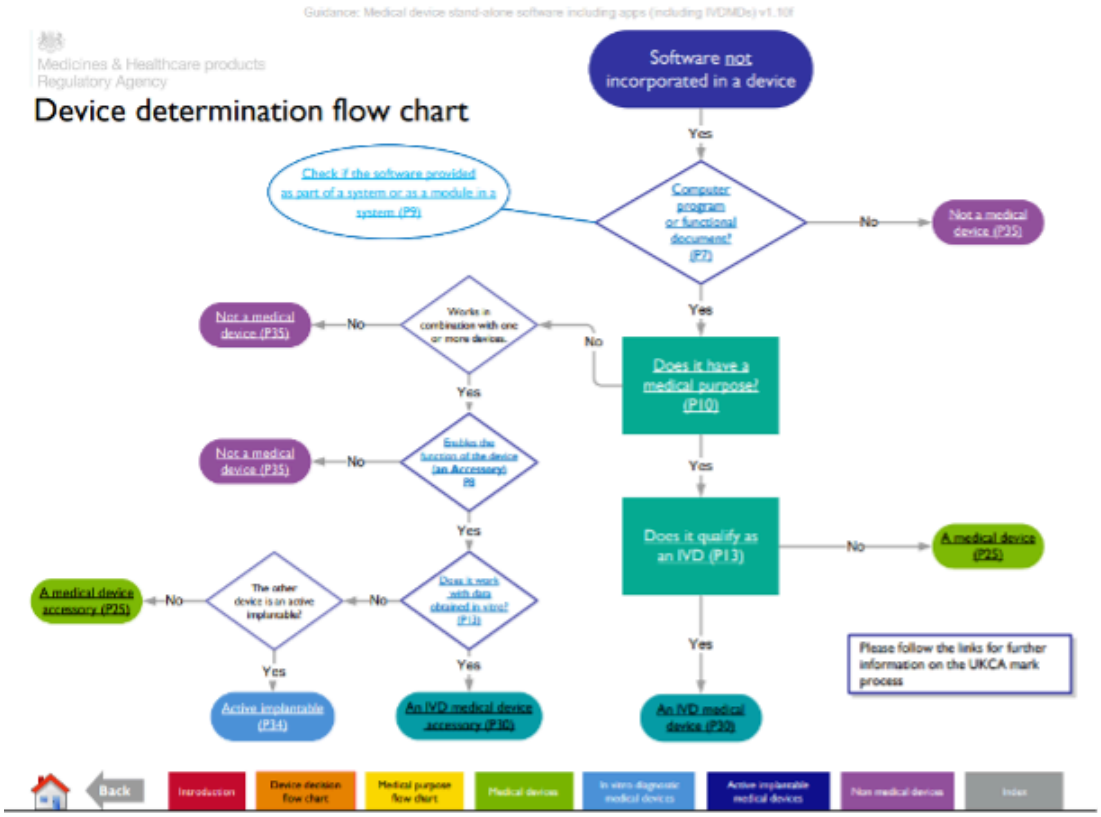

Figure 1: Device Determination Flow Chart, MHRA, Guidance Medical device stand-alone software including apps (including IVDMDs) v1.10f

The MHRA's new framework for determining whether a Digital Mental Health Technology (DMHT) qualifies as a medical device considers intended purpose and level of functionality. The first step in qualification is assessing what the manufacturer claims the tool is designed to do. Labelling, inclusive of instructions for use and promotional materials, as well as technical documentation will need to be assessed. If the software explicitly states that it diagnoses, treats, prevents, or monitors a medical condition, it is more likely to fall under medical device regulations. Despite this, a digital device may have a medical purpose but still be excluded from regulation if its functional impact is low. This ensures that only those tools providing a clinical effect or influencing patient care decisions are classified as medical devices.

Evaluating team compliance skills

After determining whether the functionality provided makes the tool a SaMD, any size of business will need to make a thorough gap analysis assessing whether they have the right set of in-house skills to ensure their device is compliant.

Many of the internal software and compliance skills that do exist will not easily apply to SaMDs. For example, it's typical with a Machine Learning (ML) software product to update the ML model once it is out on the market with new data from the field to improve accuracy. In software development, the data are added continuously, in real time, while in the healthcare sector there are strict safety protocols to follow for every additional data set. A software engineering team may not be aware of this and may not know that post market surveillance needs to be built in prior to launch.

Specifically, unlike a physical device, an AI model can always be improved just by adding more data, as additional information could improve accuracy or reduce risk. The data the model is trained on effectively is a part of the device, so an additional data set could mean that the model needs to go back through validation and verification, with the technical file also requiring updates. In this instance, early support from specialists in medical device software compliance would ensure that Predetermined Change Control Plans are included in the initial files to speed up inclusion of future data. Finally, from a compliance perspective, as regulations are evolving, it is very difficult for an internal team to always be up to date.

Evaluating team software development skills

Determining whether the in-house skills to develop the device or integration are available in-house is also critical. Startups and spin-offs will find that there is pressure from investors to prove that they have all the skills required within their teams. Outsourcing to partners that can offer software roadmaps, help hire lead engineers and advise on regulatory issues such as post market surveillance processes are key to building a strong business case.

On the other hand, larger, more established manufacturers that want to develop a whole product will typically start with a hardware element and often find that they lack the in-house skills to build a software stack inclusive of interfaces. They would typically use a product design company, a cybersecurity expert to run their vulnerability analysis and penetration test and then would need some ongoing internal capacity devoted to maintenance once the solution is up and running. Additionally, in these cases it's critical to find a trusted partner that can help train internal teams so they are up to date on their skills.

Typically, engineers that have already worked with ML will need to be part of the team, alongside professionals with experience developing Convolutional Neural Networks (CNNs), a very different skill set to querying transformer based models such as LLM.

In addition to this, skills required for software engineering and AI are distinct, so engineers might not have had exposure to how to train an AI model. On top of this, core skills such as risk assessment, data management and ML Ops become essential because they ensure full traceability across the system. Developing AI-enabled medical software requires a broader and more sophisticated skill set than traditional software development, making it difficult to hire and retain the right staff.

Finding recent graduates with the right skills is usually fairly easy, but recruiting more senior roles, such as team leads, is far more difficult. CompTIA's 2024 Tech Workforce Report [2] shows that 87% of organisations struggle to hire AI developers, taking on average 142 days to fill senior roles, highlighting a talent gap that is widening according to McKinsey Global Institute.[3] Another often overlooked aspect is the need for engineers that have user experience and UI design skills. These also come with their own set of specific compliance requirements. Verification and Validation (V&V) skills are equally critical, as is IEC 62304, which expands once AI is introduced.

AI: skills and regulations

The increasing use of AI brings about a series of complexities, many tied to the fact this technology is still nascent and at the earliest stages of regulation. The concerns are real: protecting patient safety and privacy, but without stifling much needed innovation.

The EU has taken its own approach with its EU AI Act, effective 1st August 2024, which introduces a comprehensive regulatory framework for AI with classifications. In the EU, any AI system that acts as a safety-critical component or functions as a product in its own right can be classed as "high-risk" if it falls under specific harmonised legislation. AI-driven software designed for medical purposes must respond to added expectations around managing risks, using reliable and well-curated data, providing clear user guidance and ensuring that humans remain in control of key decisions. Under the EU AI Act, systems are grouped by risk level: minimal, limited, high or unacceptable. Within the medical device world, most SaMD products fall into Class IIa or higher, reflecting their clinical importance. Only the simplest digital tools, such as fertility-tracking apps or basic prognostic calculators, typically remain in Class I.

The FDA has taken a similar stance with its guidance in its AI/ML-Based Software as a Medical Device (SaMD) Action Plan, where the agency openly acknowledges that: "The FDA's traditional paradigm of medical device regulation was not designed for adaptive artificial intelligence and machine learning technologies. Many changes to artificial intelligence and machine learning-driven devices may need a premarket review." This means regulators will often need to evaluate these products individually rather than applying a one-size-fits-all approach.

Generative AI adds yet another layer of complexity. Because these models can create new content on their own, and are therefore more susceptible to bias and "hallucinations". Studies such as Anthropics' Alignment faking suggest they may even choose to "hide" the fact that they are generating fabricated information, raising their risk profile significantly.[4] This unpredictability makes it harder to define or control the device's intended use. The foundation models underpinning GenAI are often trained on massive, loosely curated datasets and are not always transparent or tailored for medical applications, factors that further complicate their safe and reliable use in regulated healthcare settings.

In the UK, the MHRA has also provided a Software and AI Change Programme Roadmap which builds upon wider reforms for medical devices and provides guidance for programmes to ensure regulatory requirements for software and AI are clear and patients are protected. MHRA guidelines on identification of a SaMD can prove to be a useful starting point for businesses regardless of their target market as they provide a framework for interpretation and validation.

AI Governance

In conventional software, traceability focuses on who developed a feature, who reviewed it and which requirements it links back to. With AI, however, the data itself becomes part of the device, meaning teams must track where the data originated and how it was quality-checked, and maintain comprehensive logs of each step.

Explainability is even more important in the medical sector, as even complex neural networks can now be interrogated to understand why a model behaves a certain way, enabling teams to detect bias and assess failure modes. Ensuring upstream traceability for data is fundamentally different from code, yet validation is still required to demonstrate that the system performs as intended. In addition to traditional DevOps practices such as cloud-based version control, MLOps requires systematic review of not only the code but also the data pipelines as data quality is fundamental for AI. This includes sanitising and anonymising patient data so that no individual is identifiable and verifying that every file listed in the manifest is present, correct and appropriately handled. Part of the risk for startups is not being able to predict the volume of data their model will become reliable with and how much annotation will be required.

When it comes to data governance specifically, the Association for the Advancement of Medical Instrumentation (AAMI) Technical Information Report (TIR) 45 2023 recommends that organisations designate 4 separate roles for:

1. Product owner
2. Cybersecurity
3. Usability
4. Risk management

Update meetings would thus give voice to the positions and concerns of four different stakeholders. The designated cybersecurity expert might, for example, contribute to a decision as to where specific data sets need to be stored, depending on sensitivity. Similarly, data that will be placed on a hard drive should be encrypted. Similarly, they may advise on encryption at rest, justification for using or storing specific data sets or make a selection to minimise risk, preserve anonymisation and optimise integration with NHS or other third-party patient data management systems. In addition to this, the cybersecurity expert will need to advise and plan on a course of action in case of a breach.

Data and Privacy

Protecting patient data and privacy is of key importance when working with AI technologies, but even more risk-laden in the healthcare sector where information is highly sensitive.

For example, normally an anonymous patient identifier would be stored next to their data for any storage outside of third-party systems, such as NHS databases. This allows the business to easily link it back to the patient, but their identifiable data never leaves the secure third-party platform. This is often particularly pertinent with already available banks of data, where an agreement on how to handle data has already been made.

Understanding SOUP Compliance

Putting an effective Quality Management System (QMS) in place is essential when working with Software of Unknown Provenance (SOUP), ISO standards and cybersecurity requirements. QMS must have effective procedures in place to handle Software of Unknown Provenance (SOUP) in order to meet regulatory standards and cybersecurity requirements. Teams must maintain a clear record of every external software component they use, why it was chosen and how it links back to specific software requirements. Keeping your own copy of each SOUP element is equally important, as it gives you true ownership of the component and ensures you can meet your obligations under ISO 62304. This level of documentation and control becomes the backbone of demonstrating safe and compliant software development.

In most industries, using SOUP is simply practical: there's no need to rebuild something that already exists, and leveraging proven components saves time and reduces costs. In the medical sector, however, this convenience comes with added responsibility. Every piece of SOUP must be assessed more thoroughly to understand its risks, its expected behaviour and its long-term reliability. You also need a clear plan for how bugs, vulnerabilities and quality issues will be managed throughout the product's lifecycle. Once the device is on the market, ongoing surveillance is essential, as manufacturers remain responsible for monitoring third-party components and ensuring they continue to meet safety and performance requirements.

Testing and Validation

Testing and validation are crucial steps when preparing AI-driven software for medical use. Techniques like quantization, which reduces the precision of model weights and activations to make inference faster and more efficient, can be extremely useful for developing lightweight yet high-performance models. However, once these algorithms are intended for a regulated environment, every stage of the pipeline requires a higher level of scrutiny. For example, if your workflow includes image preprocessing followed by a basic landmark-recognition step, transforming this into medical-grade software means verifying the code, validating each operation and ensuring that all outputs are directly tied to a defined requirement.

A structured approach including a Traceability Matrix is essential for understanding how and where things might go wrong. A thorough testing process often starts with assessing data-quality checks before the model even performs inference, helping to catch input issues that could compromise results. It also involves deciding how the system should respond when the model is uncertain or fails, by defining clear boundaries for when to report an inference failure. Thoroughly evaluating these scenarios can help teams build AI systems that behave predictably, safely and in line with regulatory expectations.

The Global Outlook

Geographical considerations are key to selecting a launch market for an AI-driven device or SaMD as incorrect navigation of, or lack of awareness of the complex regulations can mean delays and higher overall costs. For example, it might be easier and faster to place a device in the US with minor changes, if your device meets certain criteria, than to enter the EU market first. Early conversations with regulatory experts can help you identify areas where your SaMD/AI product may be adapted to target a particular area or prevent nasty surprises later in the development process.

The US for example currently provides significant predictability over timelines and costs compared to other markets with this information being publicly available. Additionally, FDA guidance exists on the positioning of products as general wellness devices and clinical decision support software which may mean that a product need not be regulated as a medical device. Although manufacturers should be aware that enforcement discretion can be overturned by the FDA, the go to market process could be greatly simplified if they can successfully justify that their product falls within the requirements of these guidance documents. Small adaptations to a product in its early stages may help compliantly navigate down one of these routes.

In the EU, by contrast, Notified Body (NB) fee structures and defined timelines have not been set historically making it hard to plan a market launch for the EU or other markets in which the CE mark is accepted. On the 12th December 2025 the European Commission published a draft implementing regulation outlining the need to set specific quotation and timeline requirements that notified bodies must meet to ensure greater regulatory consistency and predictability.

The Commission has proposed maximum timelines for notified bodies to follow, including 30 days to review applications and sign contracts, 120 days to conduct quality management system audits, 90 days to conduct and verify an assessment of the technical documentation of the device or the representative device, and 15 days to issue a final decision and certification after conducting a final review. The draft, available for feedback until the 26th January 2026, intends to bring greater predictability, but it may put additional pressure on notified bodies.

Also on the 16th December 2025, the Commission published COM(2025)1023: a proposal to amend the MDR, the IVDR and related legislation to simplify rules on medical devices and IVDs to alleviate some of the pressures on NBs and manufacturers. One of the many proposals is to amend the classification rule 11 which has meant that many lower risk SaMD/AI devices were Class IIa. If the proposal succeeds, the challenges for some manufacturers may lessen with more SaMD falling within Class I and not requiring the involvement of an NB thereby making the EU route more attractive. Finally, the Ai Act defines SaMDs including AI as high risk and will come into force in August 2027 with further updates in this area expected. Each of these changes within the EU may affect a manufacturer's regulatory strategy if they come into force.

In the UK, which is currently following classification rules defined at Brexit, more SaMD devices fall under class I than in the EU. New requirements will be introduced in 2026 and there has been increased scrutiny around devices such as ambient scribes with latest guidance from the MHRA, indicating that this opportunity is closing. The UK thus introduces the requirements to have a holistic regulation of SaMD devices but is less specific that the EU and US in terms of cybersecurity requirements which are planned to be integrated into future UK medical device regulations.

In short, depending on your SaMD, the UK and US currently offer simpler routes to market but the EU is taking steps to lessen the burden of the EU route in the future. Regulatory strategies will need to adapt to changing requirements but nevertheless, compliance should not be an afterthought and creating documentation as part of a

consistent and ongoing compliance management system is key to ensuring the business is prepared whatever changes take place in the regulatory framework of the medical devices' chosen entry market.

The Importance of Partnerships

In an evolving technological and regulatory environment, it's important for developers and manufacturers to always have the latest information and skills available. This level of preparedness ensures the business is able to stay competitive in a rapidly changing technological environment but also remain compliant in their key markets to avoid expensive delays, bottlenecks and recalls. This is a tall achievement for any business, never mind bootstrapped startups. Even large corporates will find that gathering all the necessary skills in-house and ensuring they continue to respond to innovation and regulator demand over time, is overly ambitious.

Teaming up with specialist partners that cover different areas of expertise, particularly those areas that are most subject to change

such as AI development and SaMD regulation, can provide businesses with the reassurance that they are developing modern, effective solutions that adhere to regulatory requirements and protect patient safety and privacy at all times. For businesses of all sizes, the key is to balance these assurances with efficient working practices and time management, so that teams are not overstretched and expertise is guaranteed without necessarily adding headcount. From cybersecurity experts to regulatory consultants, to UI specialists, these professionals have made it their job to stay on top of change in their particular field of expertise and can provide much needed support to medical device manufacturers that want to keep their teams focused on what they do best.

Image recognition and AI: some caveats

One company set out to upgrade an existing prototype, originally built in MATLAB and unsuitable for cloud deployment, by rewriting its modules and demonstrating that the new version produced equivalent results.

The tool compared pairs of medical images taken at different time points and highlighted clinically meaningful changes. However, once the team began rewriting and testing the system, it became clear that the prototype had not been working as expected; it had simply performed well on the limited data it had been initially fed. While image recognition is a mature field, recognising meaningful changes over time is far more complex. The model was flagging many irrelevant differences because it had never been properly trained to distinguish clinically significant changes.

This discovery forced the team to accelerate validation, expand testing and commit to additional R&D. More AI specialist man-hours and a larger training dataset were necessary to support the level of rigour required.

To reach the required level of discrimination, segmentation had been introduced so the system could learn exactly where important changes were likely to appear. Human expertise was thus essential: annotations often begin with a trained researcher, such as a PhD student, and are then reviewed by a clinician to ensure pathological accuracy. The company had budgeted for basic annotation, but not for the far more intensive and

costly engagement of clinician reviewers required. Training data creation proved both expensive and labour-heavy, especially because some of the changes the model needed to detect were invisible to the human eye without 3D visualisation. Ultimately, the model required validation both as part of the software development process and again during clinical verification. An earlier, more experienced assessment from both an ML development and a compliance perspective could have saved this company from having to face significant costs later on.

Change Type	UK	EU	US
New intended use	Re-classify, new evidence	Re-classify, new conformity assessment	Often new 510(k) / PMA
Broadened user/population	May raise risk class	May raise risk class	Likely new submission
New clinical claim	Additional clinical evidence	Additional clinical/performance evidence	New clinical data in submission
Software algorithm change	Review; possibly new submission if changes safety	Notified Body if "significant change"	PCCP or new 510(k) if outside plan
Change to risk profile	Higher class triggers new process	Higher class triggers new process	New submission if impacts safety

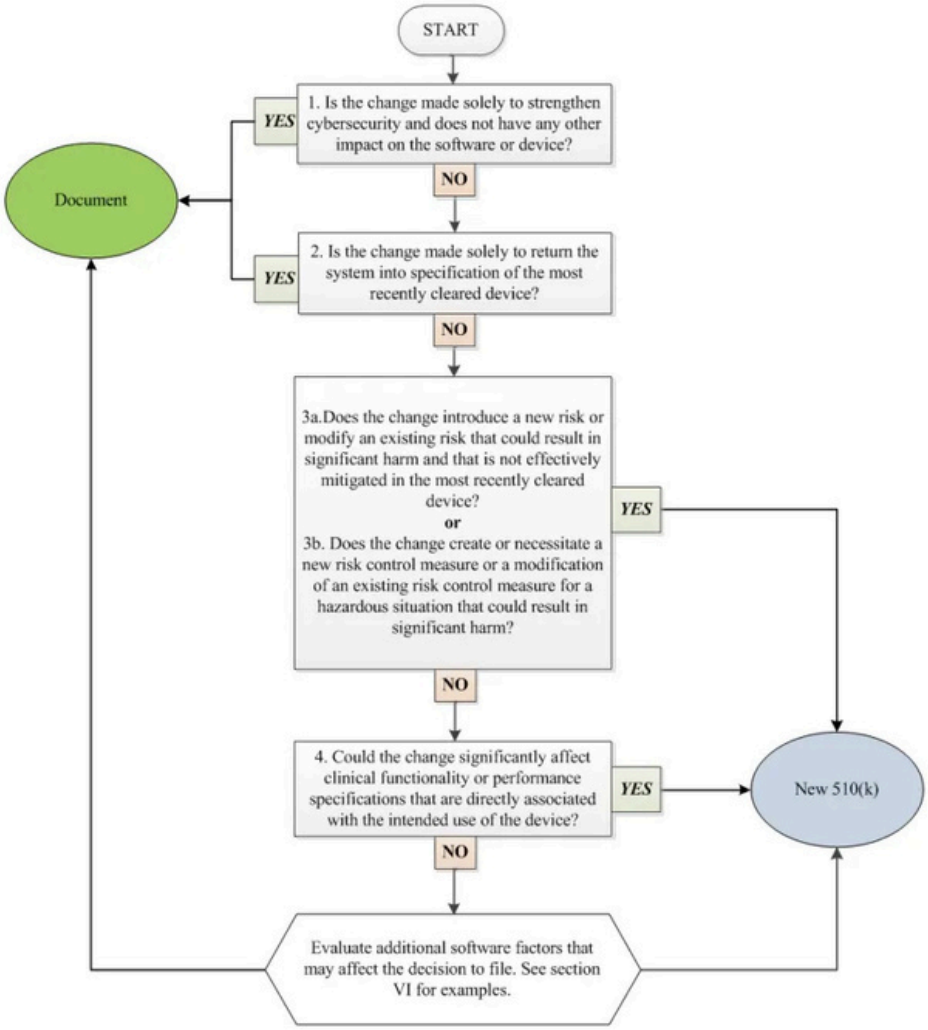

Source: FDA, Deciding When to Submit a 510(k) for a Change to an Existing Device

References

[1] Fortune Business Insights, AI in Healthcare Market Size, Share & Industry Analysis, By Platform (Solutions and Services), By Application (Robot-Assisted Surgery, Virtual Nursing Assistant, Administrative Workflow Assistance, Clinical Trials, Diagnostics, and Others), By End-user (Hospitals & Clinics, Pharmaceutical & Biotechnology Companies, Contract Research Organization (CRO), and Others), and Regional Forecasts, 2026-2034

[2] CompTIA, Turmoil weighs on tech jobs market, CompTIA reporting confirms, 2nd May 2025

[3] McKinsey, The state of AI in early 2024: Gen AI adoption spikes and starts to generate value, 30th May 2025

[4] Anthropic, Alignment faking in large language models

IMed Consultancy

Founded in 2012, IMed Consultancy offers a wide range of expert services to the global medical and health technology industry. We support medical device and in vitro medical device manufacturers to drive innovation, improve patient care and outcomes worldwide, providing assistance through all stages of the product lifecycle from concept and design through clinical studies and post-market surveillance.

IMed Consultancy's team of highly skilled and experienced medical regulatory professionals offers an outstanding yet accessible global regulatory service. We are committed to our team, to innovation, to our client's growth and success, to the health tech community, and patients. We build trusted relationships by offering agile solutions and education to advance healthcare innovations and adapting flexibly to client needs as we strive to stay ahead of industry trends, investing in developing our knowledge and seeking new opportunities to deliver value to the sector. With over 70 years of combined hands-on problem-solving expertise, our remit is truly global, ensuring that client devices are successfully launched and maintained in total compliance in the UK, EU and internationally.

www.imedconsultancy.com
hello@imedconsultancy.com,
+44 (0)1295724286

firefinch
software development

Firefinch Software develops custom software for life science companies. Our expert-led development team combines deep scientific domain knowledge with regulatory compliance and strategic insight to help you build the right product, ready to scale.

Strong domain expertise. *Firefinch deliver solutions in medtech, biotech and medical devices and with companies ranging from start-ups to large multi-nationals. We take pride in delivering excellent client outcomes while strengthening clients' core capabilities.*

Your software team for hire. *Firefinch can provide an entire cross-functional product team or can embed senior consultants with existing teams to strengthen their delivery capability. All projects are supported directly by a senior consultant and offer access to the skills of the entire team no matter the size of the project.*

Built for long-term success. *The Firefinch team adapt to the client's needs. From the code to control a device through to cloud platforms for interacting with your data, we are experts at end-to-end solutions aimed at both external customers and internal business users.*

https://firefinch.io/