A JOINT WHITE PAPER

MedStride × Xybern

AI in Medical Devices:

How Founders Can Meet EU MDR and EU AI Act Requirements Without Getting Stuck

Regulatory clarity for founders building AI-enabled medical devices in the EU and UK

May 2026

Introduction

This white paper is written for founders and clinical entrepreneurs who are building AI-enabled medical devices and preparing to bring them to market in the EU or UK. It is written for people who are moving fast, who understand their technology, and who want a clear, honest account of what the regulatory environment actually requires of them right now.

The regulatory landscape for AI in medical devices has changed materially in the past two years. The EU Medical Device Regulation (EU MDR) has been the baseline requirement since 2021, but it was written before AI became the dominant technology it is today. The EU AI Act, which came into force in August 2024 and is now entering full applicability, adds a second, overlapping framework that specifically addresses AI systems. Together, they create a compliance challenge that is more complex than either framework addresses alone.

Notified bodies are already asking harder questions. Submissions that would have passed two years ago are being returned with requests for additional evidence about how AI systems are governed, validated, and monitored. Founders who have built excellent AI systems are finding that the quality of their technology is not, by itself, sufficient evidence of compliance.

This paper explains what both frameworks actually require, where the gaps are, and what founders need to do about them. It is written jointly by MedStride, a regulatory consultancy founded by a former Notified Body reviewer, and Xybern, which builds the authorisation infrastructure that AI-enabled medical devices need to demonstrate compliant governance of their AI systems, and which has taken that infrastructure one step further by offering Authorised Agents: pre-built, production-ready AI agents with governance embedded by design.

The goal is practical. By the end of this paper, you should know what documentation you need, what questions notified bodies are going to ask, and what technical and governance infrastructure you need to build, or deploy ready-made, to answer those questions with confidence.

Why now?

The EU AI Act's high-risk provisions are now in full effect. Notified bodies have updated their assessment criteria. Founders who delay building compliant AI governance into their technical files will face longer review cycles, more rework, and the real risk of certification being refused or delayed. Founders who build governance in from the start will move faster through certification, not slower.

SECTION 1 · AUTHORED BY MEDSTRIDE

The Opportunity and the Challenge

AI is transforming medical devices

Artificial intelligence is no longer an emerging feature in medical devices. It is increasingly the core of the product. Diagnostic imaging tools use AI to detect anomalies at accuracy levels that match or exceed experienced clinicians. Clinical decision support systems use AI to synthesise patient histories and present treatment recommendations in real time. Remote monitoring platforms use AI to identify deterioration before clinical signs are visible to a human observer.

The commercial opportunity is significant. The global AI in medical devices market is growing rapidly, and EU and UK markets represent substantial opportunity for founders who can navigate the regulatory environment. Clinical buyers, hospital procurement teams, and integrated care systems are actively looking for AI-enabled tools that can improve throughput, reduce diagnostic errors, and support overloaded clinical workforces.

But the path from a technically excellent product to a product in the hands of clinical users runs through regulatory certification. And that path is where founders are hitting walls.

Why founders hit walls at the regulatory stage

The wall is not usually technical. Founders building AI-enabled medical devices in 2025 and 2026 are typically building on mature machine learning infrastructure, using validated datasets, and testing against appropriate clinical endpoints. The problem is rarely that the AI does not work. The problem is that the regulatory framework requires founders to prove, with documented evidence, not just that the AI works but that it works safely, consistently, and in a governed way that can be audited.

This is a different kind of challenge. It requires regulatory expertise, careful technical documentation, and increasingly, specific infrastructure for AI governance that most founders have not built because it was not previously required.

The two frameworks every founder needs to understand

There are two regulatory frameworks that govern AI-enabled medical devices in the EU. Understanding both, and understanding how they interact, is the starting point for any serious regulatory strategy.

EU MDR (Regulation (EU) 2017/745) governs medical devices as such. It sets the requirements for safety, performance, clinical evaluation, post-market surveillance, and conformity assessment. Any software that meets the definition of a medical device, including most AI-powered diagnostic and clinical decision support tools, must comply with EU MDR before it can be placed on the EU market.

The EU AI Act (Regulation (EU) 2024/1689) governs AI systems specifically. It applies a risk-based classification to AI systems, with the highest requirements placed on high-risk AI systems. Medical AI sits firmly in the high-risk category. The EU AI Act adds requirements on top of EU MDR, with a particular focus on risk management, transparency, logging, and human oversight.

Together, they create a higher compliance bar than either does alone. A founder who understands EU MDR but not the EU AI Act will build compliant documentation for their device but will be missing the AI-specific governance evidence that notified bodies are now requiring. A founder who understands the EU AI Act but not EU MDR will understand the AI requirements but will lack the broader device compliance framework. This paper addresses both.

Key takeaways from Section 1

- AI is the core of many modern medical devices, not a peripheral feature.
- Regulatory certification requires documented proof of safe, governed AI behaviour, not just technical performance.
- EU MDR and the EU AI Act together create a combined compliance bar that is higher than either alone.
- Founders who understand both frameworks from the start will reach certification faster.

SECTION 2 · AUTHORED BY MEDSTRIDE

What EU MDR Actually Requires for AI

EU MDR does not have a dedicated chapter on artificial intelligence. It was written before AI became a mainstream component of medical devices. But its requirements, particularly Annex I and the post-market surveillance obligations, apply to AI software in full, and they place obligations on manufacturers that go significantly further than many founders realise.

Annex I: General Safety and Performance Requirements

Annex I of EU MDR sets out the general safety and performance requirements (GSPRs) that all medical devices must meet. For AI-enabled devices, the most relevant requirements are:

- Devices must be designed and manufactured to perform as intended when used under normal conditions of use. For AI, this means the system must perform as stated across the intended patient population, under the range of clinical conditions described in its intended purpose.
- The risk management process must address all known and foreseeable risks. For AI, this includes risks arising from model errors, distributional shift (where the data the AI encounters in clinical use differs from the data it was trained on), failure modes specific to the AI, and risks arising from the interaction between the AI output and clinical decision-making.
- Devices must be designed to be used safely with the clinical and technical knowledge relevant to their intended users. AI tools must account for the fact that clinicians may over-rely on AI recommendations, under-scrutinise unusual outputs, or misinterpret confidence scores.
- Software must be developed in accordance with the state of the art. For AI, this includes the use of appropriate development methodologies, validation approaches, and, increasingly, governance frameworks.

The practical implication of these requirements is that the technical file for an AI-enabled medical device must include more than a description of the algorithm. It must include documentation of the risk management process as it applies specifically to the AI, evidence of validation across the intended patient population, a description of how the AI behaves under failure conditions, and evidence that the intended users are equipped to use the AI safely.

Post-Market Surveillance and Live AI Systems

Post-market surveillance (PMS) is where many AI medical device submissions show their weakest documentation. The EU MDR requirement for PMS is clear: manufacturers must proactively collect and analyse data from devices in clinical use, and use that data to update their safety assessment on an ongoing basis.

For AI systems, this obligation is particularly demanding. An AI system in clinical use is not a static device. Its behaviour can change over time if it is retrained. Its performance can degrade if the clinical context in which it is deployed differs from its development environment. Its outputs may be influenced by updates to the clinical workflows or electronic health record systems it connects to.

A credible PMS plan for an AI-enabled medical device must specify:

- What performance metrics will be monitored in clinical use, and how.
- What thresholds will trigger a safety investigation.

- How model changes (including retraining) are controlled and validated before deployment.
- How the results of PMS feed back into the risk management file.

Notified bodies are paying close attention to PMS plans for AI devices. A PMS plan that describes monitoring in general terms, without specifying how AI-specific performance signals will be captured and acted upon, is likely to attract questions.

What Notified Bodies Look for in AI Technical Files

Based on direct experience reviewing submissions at BSI and supporting founders through the certification process, the areas where AI technical files most commonly fall short are:

- Insufficient specificity in the intended purpose statement. The intended purpose must be precise enough that a reviewer can assess whether the clinical evidence and performance data are sufficient to support it.
- Risk management that addresses general software risks but does not adequately address AI-specific risks, including model-specific failure modes, the impact of distributional shift, and the risks arising from AI output interpretation.
- Validation data that does not adequately represent the intended patient population, or that uses performance metrics that are not clinically meaningful.
- A lack of documentation showing how the AI was developed, including training data provenance, model selection, and the rationale for key design decisions.
- Weak or absent documentation of how the AI system is governed in clinical use, including who is authorised to change the system, how changes are controlled, and how decisions made by the AI are recorded.

The Most Common Mistakes Founders Make

The single most common mistake is treating the technical file as a documentation exercise to be completed after the product is built. Founders who approach their technical file this way are forced to reconstruct evidence that should have been generated during development, and often find that they cannot reconstruct it because the governance infrastructure was never in place.

The second most common mistake is underestimating the post-market obligations. Founders who achieve initial certification often find that the ongoing requirements for PMS, incident reporting, and periodic safety update reports (PSURs) are more demanding than anticipated, particularly when the AI system continues to evolve.

The third mistake, which has become more significant since the EU AI Act came into force, is failing to account for the AI-specific requirements that sit on top of the MDR obligations. This is covered in the next section.

What a notified body reviewer wants to see in your AI technical file

- A precise intended purpose that defines the AI's clinical function, intended user, and intended patient population.
- A risk management file that specifically addresses AI failure modes, distributional shift, and clinician over-reliance.
- Validation evidence that is representative of the intended use population and uses clinically meaningful metrics.
- A post-market surveillance plan with specific AI performance monitoring commitments.
- Documentation of how AI decisions are recorded, governed, and auditable.

SECTION 3 · AUTHORED BY MEDSTRIDE

What the EU AI Act Adds

The EU AI Act came into force in August 2024. For high-risk AI systems, which include the vast majority of medical AI, full compliance obligations apply from August 2026 onwards, though notified bodies are already incorporating EU AI Act alignment into their assessments. Founders who wait until the deadline to begin building compliance will find themselves in a difficult position.

Where Medical AI Sits in the EU AI Act Risk Classification

The EU AI Act uses a tiered risk classification. At the highest tier are prohibited AI practices. At the next tier are high-risk AI systems, which are subject to the most stringent requirements. Medical AI sits clearly in the high-risk category.

Annex III of the EU AI Act lists the categories of AI system that are automatically classified as high-risk. This includes AI systems intended to be used as safety components of products, or as products themselves, that are covered by EU harmonised legislation, which explicitly includes medical devices under EU MDR. In practice, any AI system that qualifies as a medical device under EU MDR will also be classified as a high-risk AI system under the EU AI Act.

This matters because the obligations on high-risk AI systems are substantial, and they add to, not replace, the EU MDR requirements.

Article 9: Risk Management

Article 9 of the EU AI Act requires providers of high-risk AI systems to establish, implement, document, and maintain a risk management system. This is, in some ways, familiar territory for founders who have already built a risk management file for EU MDR purposes. But the EU AI Act risk management requirements have a distinct focus.

The EU AI Act risk management process must address risks arising specifically from the AI system's characteristics, including risks arising from data quality and representativeness, from the AI's opacity or lack of explainability, from unintended bias, and from the interaction between the AI output and human decision-making. These are not comprehensively addressed by the general risk management requirements in EU MDR, and founders who treat their EU MDR risk management file as sufficient for EU AI Act purposes are likely to find gaps.

Article 12: Logging Requirements

Article 12 of the EU AI Act requires high-risk AI systems to be capable of automatically generating logs of their operation. These logs must enable post-deployment monitoring of the system and must support the identification of risks and substantial modifications.

For medical AI founders, this is a significant requirement. It means that the AI system must be designed from the outset to generate logs that are sufficient for regulatory purposes, not

just operational monitoring. The logs must capture what the AI did, in what context, with what input, and what output it produced. They must be structured in a way that supports retrospective analysis and audit.

Many founders have logging in place for operational purposes, but operational logs and regulatory logs are not the same thing. Operational logs are designed to support debugging and performance optimisation. Regulatory logs must support audit, investigation, and accountability. The distinction matters, and founders who discover it for the first time during a technical file review will face significant rework.

Article 17: Quality Management

Article 17 requires providers of high-risk AI systems to implement a quality management system that covers the entire lifecycle of the AI system, from design and development through deployment and post-market monitoring. The QMS must be documented and must include specific provisions for data management, validation, monitoring, and corrective action.

For founders who already have a QMS in place for EU MDR purposes, the EU AI Act QMS requirements are largely additive rather than duplicative. The EU AI Act adds specific requirements for AI data management, including requirements around the provenance, quality, and representativeness of training data. It also adds requirements for the documentation of AI model performance across the intended deployment contexts.

How EU MDR and the EU AI Act Overlap, and Where the Gaps Are

There is significant conceptual overlap between EU MDR and the EU AI Act. Both require risk management, quality management, performance validation, and post-market monitoring. For founders, this overlap is a practical opportunity: a well-constructed EU MDR compliance programme will address many of the EU AI Act requirements, provided it is designed with the EU AI Act in mind.

The gaps are in the AI-specific provisions. EU MDR does not require the level of logging specificity that the EU AI Act demands. EU MDR does not address AI transparency and explainability requirements in the same way. And EU MDR does not address the governance of AI decisions, including who authorised each action, under what policy, and where the evidence of that authorisation is stored, in the way that the EU AI Act, read in conjunction with notified body expectations, now requires.

These gaps are the subject of Section 4.

UKCA Considerations

For founders targeting the UK market, the regulatory picture is distinct but similar in its demands. The UK MDR 2002 (as amended) is the operative framework for medical devices in Great Britain, and the MHRA has been clear that software as a medical device is within scope. The UK has not adopted the EU AI Act, but the MHRA has published guidance on AI as a medical device that incorporates many of the same principles, including requirements for transparency, validation, and ongoing performance monitoring.

Founders pursuing both EU and UK certification should note that, while the frameworks are not identical, a compliance programme that addresses EU MDR and the EU AI Act will provide a strong foundation for UKCA certification. The additional work required for UKCA is largely in the specifics of the UK regulatory pathway rather than in substantively different technical requirements.

EU AI Act obligations at a glance for medical AI founders

- Article 9: A risk management system that specifically addresses AI-specific risks, including bias, data quality, and opacity.
- Article 12: Automatic logging of AI operations in a format suitable for regulatory audit.
- Article 17: A quality management system with AI-specific provisions for data management and model validation.
- General: Transparency and explainability requirements for high-risk AI systems.
- General: Human oversight provisions ensuring that AI decisions can be reviewed, overridden, and corrected.

SECTION 4 · AUTHORED JOINTLY BY MEDSTRIDE AND XYBERN

The Governance Gap Nobody Talks About

There is a gap in the AI medical device compliance conversation that almost nobody is talking about directly. It sits between having a good AI system and being able to prove, with auditable evidence, that the AI system behaved correctly, was governed appropriately, and can be held accountable for every action it took.

This gap is not an obscure regulatory technicality. It is the difference between a submission that passes and one that does not. And it is the area where the most technically sophisticated founders are most likely to be caught out, because they have focused on building excellent AI and assumed that documentation will follow.

Building AI That Works vs Being Able to Prove It Behaved Correctly

An AI system that works is one that produces accurate outputs when given appropriate inputs, has been validated against a relevant clinical dataset, and performs within its intended use specifications. Most founders who reach the certification stage have built AI systems that work in this sense.

An AI system whose correct behaviour can be proven is something more. It is a system where every action the AI took can be traced back to a specific authorised instruction, where the context in which the action was taken is recorded, where the policy under which the action was authorised is versioned and auditable, and where the evidence of all of this is stored in a tamper-evident record that cannot be altered after the fact.

The gap between these two states is large, and most founders underestimate it. Building AI that works requires machine learning expertise, clinical validation, and good engineering practice. Being able to prove correct behaviour requires a governance infrastructure that most founders have not built, because it was not previously made explicit as a requirement.

What Auditable AI Actually Means

Auditable AI is a term that is used loosely in the industry, sometimes to mean explainability (understanding why the AI reached a particular conclusion), sometimes to mean logging (recording what the AI did), and sometimes to mean governance (controlling what the AI is

permitted to do). All three are relevant to EU MDR and EU AI Act compliance, but the governance dimension is the least well understood and the most likely to create submission problems.

In the regulatory context, auditability means that a reviewer, whether a notified body assessor, an MHRA inspector, or an internal quality auditor, can examine the record of the AI system's behaviour and determine, for any action the AI took:

- What the AI was asked to do.
- Under what authorisation the action was permitted.
- What policy governed the decision to authorise or deny.
- What version of the policy was in force at the time.
- What the outcome was and when it occurred.
- Whether any human oversight was applied, and by whom.

This is a demanding standard. It is not met by operational logs. It is not met by explainability tools. It requires a specific kind of governance infrastructure that is designed for regulatory accountability, not just operational visibility.

The Question Notified Bodies Are Starting to Ask

The specific question that is beginning to appear in notified body assessments of AI medical devices is this:

> **The question**
> Who authorised each action your AI took, and where is the evidence of that authorisation?

This question does not have a comfortable answer for most founders. The typical response is some version of: the AI was trained on validated data and tested against clinical benchmarks, and it is operating within its intended use specification. That is an answer about performance. It is not an answer about authorisation.

The question is about governance. Did your system have a mechanism that determined, before each action, whether that action was within the scope of what the AI was permitted to do? Was that mechanism documented? Was it versioned? Was its operation recorded? Can you produce evidence that any specific action was authorised by a specific policy rule, at a specific time, under a specific authority?

For most founders today, the answer is no. The infrastructure for this kind of authorisation governance has not been built because the question has not been asked explicitly until recently.

Real Examples of Where the Gap Appears

Diagnostic imaging tools

A diagnostic imaging AI produces a risk score that is presented to a radiologist. The technical file documents the AI's accuracy on the validation dataset. But can the manufacturer demonstrate that the AI was only permitted to produce outputs within the intended use specification? That the output was generated under a versioned policy that specifies exactly what the AI is permitted to assess and report? That the output is cryptographically tied to the model version and policy version in force at the time? In most cases, the answer is no.

Patient data workflows

An AI system that processes patient data to generate a care pathway recommendation is touching sensitive personal data, making clinical assessments, and influencing clinical decisions. The EU AI Act's logging requirements apply in full. But logging that the AI processed data is not the same as demonstrating that the AI was authorised to access that data, under a specific policy, for a specific purpose, at a specific time. The distinction matters both for EU AI Act compliance and for GDPR compliance, and many founders have not built the infrastructure to evidence it.

Clinical decision support

A clinical decision support system that provides treatment recommendations to a prescribing clinician is operating in a context where the consequences of an unauthorised or incorrectly governed AI action could be serious. Notified body reviewers are asking whether the system has a mechanism that prevents the AI from operating outside its intended use specification, and whether that mechanism generates auditable evidence of its operation. In most implementations, it does not.

The governance gap in plain terms

Most medical AI founders have built AI systems that work. Far fewer have built the infrastructure to prove that those systems behaved correctly, were governed appropriately, and can be held accountable for every action they took. That infrastructure is what notified bodies are beginning to require, and what Sections 5 and 5b explain how to build, or deploy ready-made.

SECTION 5 · AUTHORED BY XYBERN

How to Close the Gap Technically

The governance gap described in Section 4 is real, but it is solvable. The solution is an authorisation layer, infrastructure that sits between your AI agents and your clinical systems, intercepts every action before it executes, checks it against a versioned policy, and records the outcome with cryptographic proof. This section explains what that infrastructure looks like, how it works, and how it maps onto the specific requirements of EU MDR and the EU AI Act.

What an Authorisation Layer Is

An authorisation layer is not monitoring. It is not guardrails. It is not explainability tooling. It is a mandatory decision point that every AI action must pass through before it executes. The question it answers is binary: is this action authorised, or is it not?

In a medical device workflow, the authorisation layer sits between the AI system and the clinical infrastructure it interacts with, including the electronic health record, the diagnostic output interface, the patient data store, and the clinical decision support interface. Every action the AI proposes is intercepted at this layer. The layer checks the action against the current authorisation policy, verifies the identity of the AI agent proposing the action, and

returns a decision: authorise or deny. If the action is authorised, it proceeds. If it is not, it is blocked before it reaches the clinical system.

This is fundamentally different from post-hoc monitoring or audit logging. Post-hoc monitoring tells you what happened after the fact. An authorisation layer determines what is permitted to happen before it occurs.

How Every AI Action Is Intercepted, Checked, and Logged

Xybern's implementation of the authorisation layer operates as a five-stage pipeline that processes every AI action:

- Intercept: The layer sits between the AI agents and the clinical infrastructure. No action reaches the clinical system without passing through the authorisation pipeline. There is no bypass.
- Identify: Every AI agent carries a cryptographic identity. The authorisation layer verifies exactly who is acting, including which agent, under which role, with what trust level, before any action is evaluated.
- Authorise: The proposed action is checked against the current version of the authorisation policy. The policy is expressed in code, versioned in source control, and defines precisely what each agent is permitted to do and under what conditions.
- Decide: The layer returns a binary decision, either authorise or deny. The decision is deterministic and traceable to the exact policy clause that governed it. There is no scoring system, no probabilistic threshold, and no ambiguity.
- Record: Every decision is written to the Provenance Vault with a cryptographic signature and a hash chain linking it to the previous record. The record is immutable from the moment it is written. It cannot be altered, deleted, or backdated.

This pipeline runs for every action, every time, without exception. The result is a complete, tamper-evident record of every AI action, the authorisation decision that governed it, and the policy version in force at the time.

What Cryptographic Provenance Means for Regulatory Submissions

Cryptographic provenance is the mechanism by which the authorisation records can be proven to be authentic and unaltered. Xybern implements this using a combination of HMAC-SHA256 signatures and SHA-256 hash chains.

Each record in the Provenance Vault is signed with a cryptographic signature that allows any reviewer to verify that the record was written by Xybern's authorisation layer and has not been altered since. The records are linked in a hash chain: each record includes a cryptographic hash of the previous record, so any alteration to any record in the chain causes the chain to break and the alteration to become immediately detectable.

The practical implication for regulatory submissions is significant. When a notified body asks for evidence that a specific AI action was authorised, or that a specific action was correctly denied, the manufacturer can produce a Provenance Vault record that is cryptographically verifiable as authentic. The record shows the action, the agent identity, the policy version, the decision, and the timestamp. A reviewer can independently verify that the record has not been tampered with.

This is the kind of evidence that the EU AI Act's Article 12 logging requirements are designed to support. It is also the kind of evidence that a notified body asking who authorised this action, and where is the proof, is looking for.

How Delegation Records Work in Multi-Agent Workflows

Many AI-enabled medical devices are not built around a single AI model. They involve multiple AI agents working in sequence or in parallel, with one agent to process incoming data, another to generate a preliminary assessment, another to aggregate evidence, and another to produce the output that reaches the clinical user. In these multi-agent architectures, the governance question becomes more complex: not only must each individual agent action be authorised, but the delegation of authority from one agent to another must also be governed and recorded.

Xybern addresses this through delegated authority chains. When Agent A delegates a task to Agent B, the delegation is itself an action that must be authorised. The scope of the delegation, meaning what Agent B is permitted to do on behalf of Agent A, is explicitly defined and bounded. Agent B cannot act beyond the scope of what was delegated to it. And the delegation record is stored in the Provenance Vault with the same cryptographic integrity as any other authorisation record.

This means that for any output produced by a multi-agent AI system, a reviewer can trace the complete chain of authorisations: which agent produced the output, what it was delegated to do, who delegated it, under what authority, and what each step in the chain was permitted to do. This is the kind of accountability chain that auditors and regulators require, and that most multi-agent medical AI systems cannot currently produce.

How This Maps onto the MDR Technical File and EU AI Act Logging Requirements

The Xybern authorisation layer is designed to generate evidence that maps directly onto the specific requirements of EU MDR and the EU AI Act.

- EU MDR Annex I risk management: The authorisation layer implements the control that ensures the AI cannot act outside its intended use specification. The Provenance Vault provides the evidence that this control operated correctly.
- EU MDR post-market surveillance: The Provenance Vault records provide the data source for PMS monitoring of AI behaviour in clinical use. Deviations from expected behaviour patterns can be detected from authorisation records without requiring additional clinical logging infrastructure.
- EU AI Act Article 12 logging: The Provenance Vault is the logging infrastructure that Article 12 requires. The records are automatically generated, cryptographically secured, and structured for regulatory review.
- EU AI Act Article 9 risk management: The authorisation policy is the documented, versioned control that manages the risk of AI acting outside its permitted scope. The policy-as-code approach means that the risk control is directly auditable.
- EU AI Act human oversight requirements: The Xybern escalation mechanism provides the human review queue for actions that require operator approval before proceeding. This is the implementation of the human oversight obligation that the EU AI Act requires for high-risk AI systems.

What a Notified Body Actually Receives as Evidence

When a manufacturer using Xybern's authorisation layer prepares their technical file, the AI governance section of the file can be supported by the following evidence:

- The authorisation policy, expressed in code, versioned in source control, and linked to the risk management file as a documented control.

- Provenance Vault records demonstrating that the authorisation pipeline operated during the validation and clinical testing phases, with cryptographic proof of record integrity.
- Delegation chain records for any multi-agent workflows, demonstrating that agent-to-agent delegation was governed and bounded.
- Evidence of the shadow mode testing process, showing that new policy versions were tested against live traffic before deployment.
- Escalation records demonstrating that the human oversight mechanism functioned as described.

This is evidence that a notified body can examine, verify, and rely on. It answers the question of who authorised each AI action, under what policy, and where the proof is, with cryptographic certainty.

Founders who need to build their own agents on top of this infrastructure can do so using Xybern's Policy-as-Code SDK and deployment model. But for founders who want to move faster, Section 5b describes a second option: Authorised Agents, production-ready AI agents that are built on the authorisation layer from day one.

What the Xybern authorisation layer provides for your technical file

- Cryptographically verifiable records of every AI action and the authorisation decision that governed it.
- Versioned, auditable authorisation policies expressed in code and linked to the risk management file.
- Delegation chain records for multi-agent workflows.
- A human oversight escalation queue with full logging.
- A Provenance Vault that meets EU AI Act Article 12 logging requirements.
- Shadow mode testing evidence for policy changes.

SECTION 5B · AUTHORED BY XYBERN

Authorised Agents: Governed AI, Ready to Deploy in Healthcare

Building the authorisation layer and then building AI agents on top of it takes time and engineering resource. For many healthcare founders, the faster path is to deploy agents that already have governance built in by design. Xybern has taken the authorisation layer it built for enterprise AI and used it as the foundation for a suite of production-ready AI agents, called Authorised Agents.

Every other agent platform on the market adds governance after the fact: monitoring layers, guardrails, and post-hoc audit logs that are retrofitted onto agents that were built without governance in mind. Authorised Agents are architecturally different. They are built on the authorisation layer from day one. Every skill, every automation, every connector runs under

the authorisation pipeline before it touches anything. Governance is not a feature. It is the foundation.

For founders building AI-enabled medical devices, this distinction has direct regulatory significance. The authorisation evidence that Section 5 describes, the Provenance Vault records, the versioned policies, the delegation chains, is generated automatically by Authorised Agents as a native output of how they operate. There is no integration work required to produce the evidence that notified body reviewers are asking for. It is already there.

What Authorised Agents Are

Authorised Agents is Xybern's agent platform, providing a complete environment for defining, deploying, and governing AI agents in regulated environments. The platform has three core components that map directly onto the needs of healthcare AI founders:

Skills Custom capabilities that define exactly what an agent can do. Build once, reuse across agents. Every skill executes under the authorisation pipeline before it touches clinical data, EHR systems, or any connected infrastructure. Skills are versioned, auditable, and linkable to the risk management file as documented capabilities with defined scope.	**Automations** Scheduled and triggered workflows that orchestrate agent activity. Automations pause at high-stakes steps and wait for an authorisation verdict before the next action executes. This is the EU AI Act's human oversight requirement implemented by design, not by configuration. Escalation to a human reviewer is built into the workflow architecture.
Connectors Integrations with clinical and enterprise systems, including Google Drive, Gmail, SharePoint, and other tools relevant to healthcare workflows. Each connector carries a declared permission scope that is enforced by the authorisation layer before any data is accessed. A connector cannot access data outside its declared scope, and the evidence of every access decision is in the Provenance Vault.	**Provenance Vault** Every action taken by an Authorised Agent is written to the Provenance Vault with HMAC-SHA256 signatures and SHA-256 hash chains. The records are tamper-evident from the moment they are written. Merkle proofs allow individual decisions to be disclosed selectively to regulators, auditors, or notified body reviewers without exposing the full audit trail.

Why This Matters for Healthcare AI Specifically

Healthcare is one of the sectors where AI errors are most consequential and least tolerable. A patient data access that falls outside an agent's authorised scope is not just a compliance event. It is a potential GDPR breach, a potential clinical safety incident, and a potential EU AI Act violation, all at once. An AI agent that takes a clinical action it was not authorised to take, even if that action is clinically correct, creates a governance exposure that cannot be resolved by demonstrating that the outcome was good.

The healthcare regulatory environment demands that AI systems operate within explicitly defined, auditable boundaries. Authorised Agents provide those boundaries by design. The permission scope of every connector, the authorisation policy governing every skill, and the escalation rules governing every automation are defined before the agent is deployed, versioned in source control, and enforced before any action executes. The Provenance Vault records that result are exactly the kind of Article 12 logging evidence that EU AI Act compliance for high-risk AI systems requires.

Authorised Agents and the EU MDR Technical File

For founders who choose to deploy Authorised Agents as part of their medical device, the AI governance section of the technical file can draw directly on the Authorised Agents architecture as documented evidence of compliant AI governance. Specifically:

- The skills catalogue provides a documented, versioned record of every capability the AI is permitted to exercise, equivalent to the risk control documentation that Annex I requires.
- The automation configurations provide evidence of how the AI workflow is structured, where human oversight is embedded, and how escalation operates, directly addressing the EU AI Act's human oversight requirements.
- The connector permission scopes provide evidence of how patient data access is controlled and bounded, directly addressing GDPR and EU AI Act data governance requirements.
- The Provenance Vault records from the development, testing, and validation phases provide the Article 12 logging evidence and the PMS data source that the technical file requires.

MedStride works with founders using Authorised Agents to translate this architecture into the technical file documentation that notified body reviewers expect. The result is a technical file that can answer the governance questions described in Section 4 with evidence that is native to the agent platform, not reconstructed after the fact.

The Two Paths: Build vs Deploy

Founders have two options when it comes to AI governance infrastructure on the Xybern platform. The first is to build their own AI agents using Xybern's authorisation layer, Policy-as-Code SDK, and Provenance Vault, integrating the governance layer into their existing or custom-built agent architecture. This is the right approach for founders who need deep customisation, have specific agent frameworks they are committed to, or are building highly specialised clinical AI that requires bespoke agent design.

The second is to deploy Authorised Agents, using Xybern's pre-built, pre-governed agent platform as the foundation for their clinical AI workflows. This is the right approach for founders who want to move faster, who need production-ready agents with governance built in, and who want the technical file evidence to be generated automatically rather than engineered from scratch.

In both cases, the regulatory outcome is the same: a governed AI system with cryptographic provenance, operating under a versioned, auditable authorisation policy, with Provenance Vault records that map directly onto the EU MDR and EU AI Act evidence requirements. The choice between the two paths is a product and engineering decision, not a regulatory one.

What Authorised Agents provide for healthcare founders

- AI agents built on the authorisation layer from day one, not retrofitted with governance after the fact.
- Skills, automations, and connectors that run under pre-execution authorisation before touching clinical data or systems.
- Native Article 12 logging through the Provenance Vault, with HMAC-SHA256 signatures and SHA-256 hash chains.
- Human oversight built into automation workflows, meeting EU AI Act human oversight requirements by design.
- Connector permission scopes that enforce data access boundaries before any patient

data is accessed.
- Technical file evidence that is native to the platform architecture, not reconstructed for submission.
- Deployable in under one week per workflow, without changes to existing clinical infrastructure.

SECTION 6 · AUTHORED JOINTLY BY MEDSTRIDE AND XYBERN

A Practical Roadmap for Founders

The requirements described in this paper are significant, but they are navigable. This section sets out a practical roadmap, covering what to do at each stage of the certification journey, how MedStride and Xybern work together in practice, and what to prioritise depending on where you are in the process.

Pre-Submission: Building the Foundation

The decisions made before a technical file is assembled determine how smooth the certification process will be. Founders who are at the pre-submission stage should focus on the following:

Define your intended purpose with precision

The intended purpose of your AI system is the foundation of your entire technical file. It must define exactly what clinical function the AI performs, who the intended users are, what patient population it is intended for, and what clinical environment it is intended to operate in. Vague intended purpose statements are the single most common cause of delays in early technical file reviews. Get this right first.

Classify your device and your AI system

Device classification under EU MDR and AI risk classification under the EU AI Act are not the same exercise, but they are related. Most AI-enabled diagnostic and clinical decision support tools will be Class IIa or Class IIb under EU MDR and high-risk under the EU AI Act. Understanding your classification early determines which conformity assessment route you need and which notified body requirements apply.

Build your risk management infrastructure

Start your risk management process early and design it to address both EU MDR and EU AI Act requirements from the outset. This means including AI-specific hazards, such as distributional shift, model-specific failure modes, and clinician over-reliance, in the risk assessment. It also means linking your risk management process to your authorisation policy: the policy is a risk control, and it needs to be documented and evidenced as such.

Implement authorisation governance from the start

The single most costly mistake founders make is deferring the implementation of AI governance infrastructure until the documentation phase. By then, the evidence that the

governance infrastructure was operating during development and testing cannot be generated retrospectively. Whether you are building on Xybern's authorisation layer directly or deploying Authorised Agents, implement governance during development, run it through your validation testing, and generate Provenance Vault records that your technical file can rely on.

During Technical Documentation

The technical file review is where the quality of your pre-submission preparation becomes visible. The areas that attract the most attention from notified body reviewers, and that MedStride works with founders on most intensively, are:

- Clinical evaluation: Does your clinical evidence support the intended purpose? Is the evidence representative of the intended patient population? Are the performance metrics clinically meaningful?
- Risk management: Is the risk management process comprehensive, AI-specific, and linked to documented controls? Are the residual risks acceptable?
- AI governance: Can you demonstrate who authorised each AI action during clinical testing? Is there a versioned, auditable authorisation policy? Is there evidence of the authorisation infrastructure operating as described?
- Post-market surveillance: Is the PMS plan specific, realistic, and designed to capture AI-specific performance signals?

MedStride provides direct support at each of these stages, drawing on experience as a former notified body reviewer to anticipate and address the questions that reviewers are likely to ask.

Post-Market: Maintaining Compliance

Certification is not the end of the compliance journey. The post-market obligations for AI-enabled medical devices are ongoing and demanding. Founders need to:

- Operate the PMS programme as described in the technical file, collecting and analysing the performance data specified.
- Maintain the Provenance Vault as an ongoing record of AI system governance, providing the data source for PMS monitoring.
- Manage any model changes, including retraining, as design changes, with appropriate validation and technical file updates.
- Submit Periodic Safety Update Reports (PSURs) at the frequency required by device classification.
- Update the risk management file and authorisation policy as new risks are identified or clinical context changes.

How MedStride and Xybern Work Together

MedStride provides the regulatory expertise. Xybern provides the technical governance infrastructure, including both the authorisation layer for founders building their own agents and the Authorised Agents platform for founders who want production-ready, pre-governed AI. In practice, the collaboration works as follows:

MedStride leads on regulatory strategy, device and AI classification, technical file structure, clinical evaluation, risk management methodology, and notified body engagement. Xybern leads on the implementation of the authorisation layer or the deployment of Authorised Agents, the design of the authorisation policy, the configuration of the Provenance Vault, and the translation of authorisation records into the format required for the technical file.

The critical integration point is the risk management file. MedStride works with the founder to identify the AI-specific risk controls required. Xybern implements those controls as authorisation policies, either through the SDK or through the Authorised Agents skills and connector architecture, and generates the evidence that they operated correctly. MedStride incorporates that evidence into the technical file in the format that notified body reviewers expect.

This means that founders working with both MedStride and Xybern get a fully integrated regulatory and technical governance programme, not two independent workstreams that need to be reconciled at the documentation stage.

What to Prioritise at Each Stage

Early stage (pre-product or pre-clinical validation)

- Define intended purpose precisely before building technical infrastructure around it.
- Decide early whether to build on Xybern's authorisation layer directly or to deploy Authorised Agents. Both paths generate the same regulatory evidence; the choice is an engineering and product decision.
- Implement governance in your development environment so that Provenance Vault records are generated from the start.
- Engage a regulatory consultant early to shape your development process, not just to document it afterwards.
- Understand your classification and conformity assessment route before you design your clinical validation programme.

Approaching certification

- Conduct a gap assessment against Annex I GSPRs and the EU AI Act high-risk obligations.
- Ensure that your Provenance Vault records cover the entire development and validation period.
- Review your post-market surveillance plan against the specific AI performance signals your system generates.
- Engage with your chosen notified body early to understand their current expectations for AI technical files.

Questions to Ask Your Regulatory Consultant and AI Governance Provider

As you build your compliance programme, the following questions will help you assess whether your regulatory consultant and AI governance provider are equipped to address the combined EU MDR and EU AI Act requirements:

- How do you approach the EU AI Act requirements for high-risk AI systems, and how do they integrate with the EU MDR technical file?
- What evidence of AI governance have notified bodies been requesting in recent submissions, and how are you addressing it?
- How do you document the risk management controls specific to AI, including the authorisation policy?
- What format does the Provenance Vault evidence take in the technical file, and have notified body reviewers accepted it?
- How do you manage model changes and retraining under the post-market obligations?

- What does your PMS programme look like for an AI system that is continuing to evolve post-certification?
- If we deploy Authorised Agents, how does the skills and connector architecture map onto the technical file requirements?

Roadmap summary: the three things to do first

- Define your intended purpose with clinical and regulatory precision, as this is the foundation of everything else.
- Implement authorisation governance from the start of development, not at the documentation stage. Choose between building on Xybern's layer or deploying Authorised Agents based on your product architecture.
- Engage regulatory and AI governance expertise early, and make sure they are working as an integrated programme.

Conclusion

The regulatory environment for AI-enabled medical devices has changed materially, and it continues to change. The EU AI Act is now in full effect for high-risk AI systems. Notified bodies are incorporating AI governance requirements into their assessment criteria. The bar for certification has risen, and it will not come back down.

For founders, this is not a reason for pessimism. It is a reason to act now rather than later. The founders who build compliant AI governance into their devices from the start, who treat the authorisation layer, the Provenance Vault, and the risk management integration as foundational infrastructure rather than documentation afterthoughts, will move faster through certification, not slower. They will generate the evidence that notified body reviewers are asking for because they built systems that produce that evidence as a natural outcome of their operation.

The founders who delay will face a different experience. They will arrive at the documentation stage and discover that the evidence they need cannot be reconstructed. They will face requests for information they cannot provide. They will face rework cycles that extend timelines and consume resources. And they will face a regulatory environment that is only becoming more demanding, not less.

Xybern now offers two paths to the same compliant outcome. Founders who want to build their own agents can do so on Xybern's authorisation layer, using the Policy-as-Code SDK and Provenance Vault to generate the governance evidence their technical file requires. Founders who want to move faster can deploy Authorised Agents, with governance built in from day one, and clinical AI capabilities that are production-ready from the moment they are deployed. Both paths lead to the same place: AI that works, and proof that it behaved correctly.

The window to get ahead of this is now. The EU AI Act compliance clock is running. Notified bodies are updating their criteria. The questions are already being asked in submissions. Founders who treat regulatory compliance as a late-stage task will find themselves behind; founders who treat it as a design requirement will find themselves ahead.

Key takeaways

- EU MDR and the EU AI Act together create a compliance bar that is higher than either framework alone.
- The governance gap, meaning the inability to prove who authorised each AI action and provide cryptographic evidence, is the most common and most costly mistake in AI medical device submissions.
- Auditable AI requires authorisation infrastructure, not just logging or explainability tools.
- The authorisation layer must be built into development from the start; it cannot be retrofitted at the documentation stage.
- Authorised Agents are pre-built AI agents with governance embedded by design, providing the fastest path to compliant AI deployment in healthcare.
- Founders who build governance in from the start will move faster through certification, not slower.
- MedStride and Xybern provide an integrated programme that covers both regulatory expertise and technical governance infrastructure, including Authorised Agents for healthcare.

About MedStride and Xybern

MedStride

MedStride is a regulatory consultancy specialising in EU MDR and UKCA certification for medical devices, including AI-enabled software. MedStride was founded by Dr Benyamin Rahmani, a former Reviewer and Scheme Manager at the British Standards Institution (BSI), one of the world's leading notified bodies for medical devices.

MedStride brings direct notified body experience to every engagement. This means that the regulatory advice MedStride provides is grounded in an understanding of what notified body reviewers actually look for, not just what the regulations say. MedStride supports founders across the full certification journey, from regulatory strategy and device classification through technical documentation, clinical evaluation, and post-market compliance.

MedStride is a member of ABHI (the Association of British HealthTech Industries), a member of RAPS (the Regulatory Affairs Professionals Society), and holds BSI Standards Maker status.

MedStride works with medical device manufacturers and start-ups worldwide seeking entry into the EU and UK markets, with particular expertise in cardiovascular technologies and AI-enabled software as a medical device.

Contact MedStride

Website: medstride.co.uk
Email: contact@medstride.co.uk

Phone: +44 20 3432 3698
Address: 60 Tottenham Court Road, Office 413, London, W1T 2EW

Xybern

Xybern builds the authorisation layer for AI agents and, uniquely, builds agents on top of it. In regulated environments, including medical devices, financial services, and legal services, AI agents trigger real actions with real consequences. Xybern intercepts every action before it executes, returns a clear authorise-or-deny verdict, and records the outcome with cryptographic proof in the Provenance Vault.

Xybern is not a monitoring platform, an observability layer, or a guardrails tool. It is a mandatory authorisation pipeline: every agent action, authorised or denied before it runs. The Provenance Vault generates the tamper-evident, cryptographically verifiable records that EU AI Act Article 12, EU MDR post-market surveillance, and notified body AI governance requirements demand.

Xybern offers two products for healthcare AI founders. The authorisation layer, available via the Policy-as-Code SDK, allows founders to govern their own custom AI agents with pre-execution authorisation and Provenance Vault logging. Authorised Agents is Xybern's pre-built, pre-governed agent platform, providing skills, automations, and connectors that run under the authorisation layer from day one, deployable in under one week per workflow, with governance evidence generated natively for regulatory submissions.

Xybern is framework-agnostic and model-agnostic, working with CrewAI, AutoGen, LangGraph, and any custom multi-agent architecture. Xybern is backed by NVIDIA Inception and AWS for Startups.

Contact Xybern

Website: xybern.com
Email: info@xybern.com
Authorised Agents: xybern.com/authorised-agents
Pilot enquiries: xybern.com/enterprise

Ready to speak to our team?

If you are a founder building an AI-enabled medical device and you want to understand what EU MDR and EU AI Act compliance means for your product, or if you want to explore deploying Authorised Agents in your clinical workflows, speak to MedStride and Xybern. We work with early-stage founders and companies approaching certification. The earlier you engage, the more we can help.