

EU AI Act vs FDA AI/ML Cybersecurity for Devices

Reviewed 2026-10-22 | How EU AI Act Article 15 obligations compare to the FDA's PCCP framework and Section 524B for AI/ML SaMD: model integrity, robustness, postmarket monitoring.

The EU AI Act governs AI/ML medical device cybersecurity mainly through Article 15, which demands robustness against data poisoning, model poisoning, and adversarial evasion, plus Article 10 data governance and Article 72 postmarket monitoring. The FDA addresses the same risks under Section 524B, its AI/ML SaMD framework, and Predetermined Change Control Plans that govern model updates. FDA evidence should be risk-based; the January 2025 AI lifecycle document is draft, not for implementation. AI use alone does not require a PCCP, adversarial-test format, or drift metric. Shared evidence may support both regimes, but each requires its own applicability and conformity assessment.

How EU AI Act Article 15 cybersecurity obligations for high-risk AI systems compare to FDA's PCCP framework and Section 524B applied to AI/ML SaMD - model integrity, adversarial robustness, postmarket monitoring.

AI and machine learning medical devices face a cybersecurity regime that is sharper than standard SaMD on both sides of the Atlantic. This guide compares EU AI Act Article 15 cybersecurity obligations against the FDA's AI/ML SaMD framework and Section 524B as applied to AI/ML devices.

Reviewed October 22, 2026 against the February 3, 2026 final cybersecurity guidance and August 18, 2025 final AI PCCP guidance. The January 2025 AI lifecycle document remains draft. Assess EU AI Act obligations under their applicable phase-in dates, not as universally operative for medical devices today.

Community project: ai-samd.com is a free reference resource we sponsor for AI-SaMD-specific topics - model lifecycle, PCCPs, EU AI Act conformance. Use this guide for the cybersecurity narrative; use ai-samd.com for AI-SaMD lifecycle depth.

At a glance - How the regimes apply to AI/ML SaMD

Dimension	EU AI Act	FDA
Classification	Most AI medical devices = "high-risk AI system" under Article 6(1) and Annex I, where third-party conformity assessment is required	Assess the Section 524B(c) definition and covered pathway; AI alone does not establish applicability
Anchor for cybersecurity	Article 15 - accuracy, robustness, cybersecurity	Section 524B + AI/ML SaMD framework + PCCP guidance
Model update governance	Assess substantial modification under the applicable rules	Applicable change requirements; an optional FDA-authorized PCCP may cover specified changes
Adversarial robustness	Explicit requirement (Article 15(5))	Risk-based testing recommended where relevant
Training data governance	Article 10 - data governance	Implicit via SPDF + 524B
Postmarket monitoring	Article 72 - post-market monitoring plan	524B(b)(1) vulnerability plan for covered submissions; device-specific performance monitoring, also discussed in the AI lifecycle draft
Incident reporting	Article 73 - serious incidents	Assess Part 803 MDR and Part 806 corrections/removals separately; CVD is not itself an MDR report
Standards anchor	ISO/IEC 42001, ISO/IEC 27090 (emerging), ENISA AI guidance	NIST AI RMF, IEC 81001-5-1, MITRE ATLAS

Why AI/ML SaMD is a different cybersecurity problem

A connected infusion pump and an AI-driven diagnostic share the standard threat surface - network interfaces, OTA updates, SBOM hygiene. AI/ML SaMD adds threats that don't exist for traditional software:

- **Model evasion** - adversarial inputs crafted to flip the model's output (an image perturbed to look benign to a tumor classifier).
- **Model inversion** - extracting training data by querying the model (a privacy attack against PHI used in training).
- **Membership inference** - determining whether a specific record was in the training set.
- **Model poisoning** - corrupting training data or weights to introduce a backdoor.
- **Model theft / extraction** - querying the model to reconstruct it.
- **Prompt injection** (for LLM-based devices) - instruction smuggling that bypasses safety guardrails.

- **Data drift** - statistical shift between training and deployment data that degrades performance without any malicious actor.

Assess relevant AI-specific threats in the actual device context. Neither AI use nor a generic threat list establishes sufficient evidence or predicts reviewer acceptance.

The EU AI Act side

CRA Article 2(2)(a) and (b) exclude products regulated by MDR/IVDR, including medical AI products in that scope. CRA reporting of actively exploited vulnerabilities and severe incidents has applied since September 11, 2026, and main obligations apply December 11, 2027, for products within CRA scope, not excluded medical devices. See [Regulation \(EU\) 2024/2847](#) and [Regulation \(EU\) 2024/1689](#).

Classification

The EU AI Act (Regulation 2024/1689) creates four risk tiers. Under Article 6(1), AI systems that are products, or safety components of products, covered by Annex I legislation, including MDR/IVDR, are high-risk when those products require third-party conformity assessment. This is the medical device route, not Annex III.

Article 15 - accuracy, robustness, and cybersecurity

For Article 6(1) product-based high-risk systems, the corresponding obligations apply from August 2, 2027 under Article 113. The following describes those provisions, not blanket current medical-device obligations. Article 15 is the cybersecurity anchor. It requires that high-risk AI systems be designed and developed to achieve appropriate levels of accuracy, robustness, and cybersecurity throughout their lifecycle, and to perform consistently in those respects.

Article 15(5) is explicit on cybersecurity: high-risk AI systems must be resilient against attempts to alter their use, outputs, or performance by exploiting system vulnerabilities. Technical solutions must address AI-specific vulnerabilities, including:

- **Data poisoning** - attacks on training datasets.
- **Model poisoning** - attacks on pre-trained components.
- **Adversarial examples / model evasion** - attacks on inputs at inference time.
- **Confidentiality attacks** - model inversion, membership inference, model theft.
- **Flaws** in the AI model itself.

Article 10 - data governance

Article 10 requires training, validation, and testing datasets to be subject to appropriate data governance - including examination for biases, gaps, and data integrity. From a cybersecurity perspective, training data lineage becomes a controlled artifact you must protect, audit, and version.

Article 72 - postmarket monitoring

Providers of high-risk AI systems must establish and document a post-market monitoring system proportionate to the system's nature and risks. For medical AI, this overlaps with the MDR PMS plan; you build one system that feeds both.

Article 73 - serious incident reporting

Providers must report serious incidents to market surveillance authorities. For medical AI, this overlaps with MDR vigilance reporting with the applicable medical-device reporting rules.

Substantial modification

EU AI Act requires re-conformance if the system undergoes a "substantial modification" that affects compliance with Article 15 (or other Section 2 requirements). This is the EU analogue to the FDA's PCCP question.

The FDA side

Section 524B applies

If the AI/ML SaMD meets the Section 524B(c) cyber-device definition and is submitted on a covered pathway, the 524B duties apply (see our [FDA Premarket Cybersecurity Submission Checklist](#)). The February 2026 final guidance does not carve out AI/ML; it adds expectations.

AI/ML SaMD framework

The 2019 discussion paper introduced SaMD Pre-Specifications (SPS) and an Algorithm Change Protocol (ACP) as historical proposals, not the current final PCCP structure. The August 18, 2025 final AI PCCP guidance recommends:

- **Description of Modifications** - the planned changes.
- **Modification Protocol** - how changes will be developed, evaluated, and implemented.
- **Impact Assessment** - the assessment of the modifications' effects.

A PCCP is optional and must be authorized by the FDA before it supports covered changes without another marketing submission.

PCCP and cybersecurity

The PCCP is where AI/ML cybersecurity gets specific. A PCCP is optional. If proposed, assess relevant cybersecurity impacts and consider device-specific evidence such as:

- How retrained models are validated for adversarial robustness (not just accuracy).
- How the SBOM and model card update with each release.
- How signing, integrity protection, and rollback work for model artifacts.
- How drift monitoring triggers a retraining cycle vs. a postmarket safety notification.

Assess relevant cybersecurity impacts from the specified changes; no omission alone predicts a deficiency.

Pen testing scope for AI/ML

For AI/ML SaMD, penetration testing should address applicable AI-specific threats. The FDA recommends tester independence, which can be provided by independent internal testers or qualified third parties.

Risk-based scenarios include:

- Adversarial input crafting against the deployed model.
- Model extraction attempts via the inference API.
- Prompt injection (for LLM-based systems).
- Input validation under malformed or out-of-distribution data.
- Authentication and authorization on model update endpoints.

A pen test that only covers network/web/protocol surfaces and ignores the model is not a complete AI/ML pen test.

Postmarket: drift monitoring

Consider performance drift monitoring where relevant to the device's risks. The January 2025 AI lifecycle draft proposes performance monitoring; Section 524B(b)(1) requires a cybersecurity vulnerability plan, not a universal AI drift metric. Drift can be benign (population change) or adversarial (poisoning at inference time). Your monitoring must distinguish.

Side-by-side: AI/ML cybersecurity deliverables

Deliverable	EU AI Act	FDA AI/ML SaMD
Threat model with AI-specific threats	Risk-based evidence supporting Article 15(5), not a prescribed document format	Guidance-recommended, risk-based evidence
Adversarial robustness testing	Appropriate technical measures under Article 15(5), not a universal test suite	Risk-based testing recommended where relevant
Training data governance	Required (Article 10)	Implicit (SPDF + design controls)
Model signing / integrity	Appropriate integrity measures under Article 15; signing is an implementation option	Risk-based integrity controls; 524B(b)(2) requires cybersecurity processes and updates
Drift monitoring	Required (Article 72)	524B(b)(1) vulnerability plan; AI drift monitoring is device-specific, discussed in lifecycle draft
Change-control governance	Assess substantial modification	Applicable change requirements, with an optional authorized PCCP
Cybersecurity in technical documentation	Annex IV technical documentation	Applicable cybersecurity evidence; PCCP only if proposed
Incident reporting clock	Applicable Article 73 and MDR/IVDR vigilance provisions	Separate Part 803 MDR and Part 806 assessment; CVD is not itself an MDR report

Build sequence for dual-market AI/ML SaMD

- **Build the threat model with AI-specific threats from day one.** STRIDE plus ML-specific extensions (LINDDUN for privacy, MITRE ATLAS for AI-system attacks).
- **Lock training data lineage as a controlled artifact.** Dataset versioning, integrity hashes, provenance documentation. Satisfies Article 10 and FDA design controls.
- **Design the PCCP to satisfy both regimes.** An optional FDA-authorized PCCP may provide useful change documentation, but does not determine EU substantial modification or conformity outcomes. An out-of-PCCP change does not automatically require a new 510(k), and FDA authorization does not determine EU conformity.
- **Scope the pen test to include AI-specific scenarios.** Adversarial inputs, extraction attempts, prompt injection, model update endpoint testing.
- **Run drift monitoring in production with two thresholds:** statistical drift (retrain trigger) and integrity drift (security incident trigger). The latter triggers assessment under applicable medical-device incident reporting rules.
- **Maintain a model card** alongside the SBOM. Model card covers intended use, training data summary, performance characteristics, known limitations; SBOM covers software components. Both update on every release.

Common AI/ML cybersecurity failure patterns

1. **Treating the model as a black box in the threat model.** Threat models stop at the API and don't model attacks on the model itself. Both regulators now expect ML-specific threats addressed.
2. **PCCP without cybersecurity.** PCCP describes the model retraining cadence and validation but ignores the cybersecurity implications of releasing new model weights. PCCP is a security artifact, not just a performance artifact.
3. **Drift = retrain only.** Investigate drift using device-specific evidence. A signal may concern performance, data integrity, or security; it does not automatically establish a reportable incident. Assess clinical risk and reporting independently.
4. **Training data not under change control.** Training data lives in a data scientist's notebook, not the QMS. Article 10 and FDA design controls both expect documented lineage and integrity.
5. **Generic web-app pen test.** Pen test report covers the inference API as a web service but does no adversarial input testing or model extraction attempts. Not a complete AI/ML pen test.

FAQ

Does the EU AI Act apply on top of MDR?

Yes. For medical AI, both regulations apply. The AI Act layers additional obligations (Article 10 data governance, Article 15 robustness/cybersecurity, Article 72 monitoring, Article 73 reporting) on top of MDR/IVDR. The conformity assessment under MDR is the gateway; AI Act requirements are folded into the technical documentation.

Do I need a PCCP for the EU?

The EU doesn't have a PCCP construct, but the underlying need is the same: define which model changes are pre-authorized vs. which trigger re-conformance ("substantial modification"). A PCCP may support EU change-control documentation, but it does not automatically satisfy AI Act requirements.

What's MITRE ATLAS and should I use it?

MITRE ATLAS is the adversary tactics and techniques knowledge base for AI systems - the AI analogue of MITRE ATT&CK. It's the most actionable framework for building an AI-specific threat model. Neither regulator mandates it; both accept it.

Is ISO/IEC 42001 mandatory?

Not yet, but it's becoming the reference AI management system standard. Expect Notified Bodies to look for it (or NIST AI RMF, its U.S. analogue) as evidence of an AI governance program.

How do I test for adversarial robustness?

For image models: project gradient descent (PGD), FGSM, AutoAttack against the deployed model. For text/LLMs: prompt injection benchmarks, adversarial suffixes. For tabular: feature perturbation under model-specific constraints. The pen test report should include scope, methodology, and results - not just "we tested adversarial robustness."

What's the FDA postmarket reporting trigger for AI drift?

The FDA specifies no universal drift threshold that automatically triggers reporting. Assess whether available information meets the death, serious-injury, or malfunction criteria in 21 CFR 803.50 and any Part 806 correction/removal duties. The January 2025 lifecycle draft proposes performance monitoring but establishes no reporting exemption or automatic trigger.

Where this fits in our content

- [SaMD Cybersecurity Requirements](#) - the base SaMD layer.
- [FDA Premarket Cybersecurity Submission Checklist](#) - the 524B baseline.
- [EU MDR vs FDA Medical Device Cybersecurity](#) - the non-AI dual-market crosswalk.
- [STRIDE Threat Modeling for Medical Devices](#) - extend with MITRE ATLAS for AI.
- [ai-samd.com](#) - community-sponsored AI-SaMD lifecycle reference.
- [GMLP Controls Crosswalk for Medical Devices](#) - each GMLP principle mapped to engineering controls.
- [AAMI CR515:2025 and ML device cybersecurity](#) - the ML threat catalog behind both regimes.

How Blue Goat Cyber helps

We build cybersecurity programs for AI/ML SaMD teams filing in the U.S., the EU, or both, and the dual-market work is mostly about generating shared evidence with separate applicability, conformity, and change assessments.

- **Threat models that include the model.** Training data poisoning, model extraction through the inference API, adversarial input crafting, and prompt or feature manipulation are enumerated as threats with owners, not mentioned as an AI caveat.
- **Adversarial and conventional testing in one scope.** The pen test covers the application and infrastructure; a separate work package probes model-specific behavior, including input perturbation and inference-boundary abuse, with results written into the same threat table.
- **PCCP authoring that survives both regimes.** The FDA wants bounded modifications with a validation protocol; the AI Act wants documented post-market monitoring of substantial modification. We prepare shared cybersecurity evidence while keeping FDA authorization and EU conformity assessments separate.
- **Drift-aware postmarket monitoring.** Performance drift and security monitoring share one intake, because a degradation report and an exploitation report can look identical for the first 48 hours.

See [FDA premarket cybersecurity services](#) and [EU MDR cybersecurity services](#).

Sources & primary references

- Regulation (EU) 2024/1689 - Artificial Intelligence Act
- FDA, Cybersecurity in Medical Devices: Quality Management System Considerations and Content of Premarket Submissions (final guidance, February 2026)
- FDA, Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions (final AI PCCP guidance, updated August 18, 2025)
- FDA, Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan (2021)
- Section 524B of the Federal Food, Drug, and Cosmetic Act
- MDCG 2019-16 Rev.1 - Guidance on Cybersecurity for Medical Devices
- NIST AI 100-1 - Artificial Intelligence Risk Management Framework (AI RMF 1.0)
- MITRE ATLAS - Adversarial Threat Landscape for Artificial-Intelligence Systems
- ISO/IEC 42001:2023 - Information technology - Artificial intelligence - Management system

Free tools for this

- [AI/ML Submission Timeline](#) - sequence the AI/ML cybersecurity work from premarket through postmarket.
- [AAMI CR515 Crosswalk](#) - map CR515 clauses to GMLP, PCCP, AIBOM, and Section 524B evidence.

Current web guide: <https://bluegoatcyber.com/guides/eu-ai-act-vs-fda-ai-ml-cybersecurity>

Informational material. Confirm requirements for your device and market. The manufacturer owns risk decisions, QMS approvals, and regulatory submissions. No regulatory outcome is promised.