

ODA3 INSTITUTE

SECTOR APPLICATION NOTE

Healthcare

Applying AI-IRF™ v1.0 to hospitals, diagnostics, medical devices and health data platforms

ODA3-2026-06-RPT-SCS-025 | Public

Document Control

Doc ID	ODA3-2026-06-RPT-SCS-025
Doc Type	Sector Application Note (RPT)
Framework Tag	AI-IRF™ v1.0 (79 controls across 19 domains)
Sector	Healthcare — hospitals, medical devices, diagnostics, health data platforms
Regulatory Crosswalk	FDA medical device rules and AI guidance (US); HIPAA (US); EU MDR/IVDR; EU AI Act; CDSCO Medical Device Rules 2017 (India)
Companion Documents	ODA3-2026-06-TCR-HAI-001 (Technical Report); ODA3-2026-06-EXB-HAI-002 (Executive Brief)
Evidence Basis	Implementation guidance. Regulatory references checked against primary texts; control references checked against ODA3-2026-06-TCR-HAI-001
Publish Date / Review Cycle	Issued 1 June 2026; updated 04 October 2026; review cycle: annual
Classification	Public

Sector Application Note — Healthcare

Applying AI-IRF™ v1.0 to hospitals, diagnostics, medical devices and health data platforms

Purpose: This note shows how AI-IRF™ v1.0 incident response controls relate to the obligations healthcare organisations and medical device manufacturers already work under. It is implementation guidance, not regulatory, legal or clinical advice. Organisations should consult qualified regulatory counsel before making any compliance claim.

Licence notice — GEL v1.0 | Public. This document is part of the GAISSE Ecosystem published by ODA3 Institute (ODA3 Pvt Ltd) under the GAISSE Ecosystem Licence (GEL) v1.0 (effective 1 May 2026); it was adopted under GEL v1.0 on 1 June 2026. GAISSE™, UAIF™, AI-IRF™, PAI-SF™ and ODA3™ are trademarks of ODA3 Pvt Ltd; registration applications are pending or the marks are common-law marks (GEL §9.1). Commercial use — including consulting, advisory work, regulatory submissions prepared for a fee, or training delivery — requires a separate written commercial licence (GEL §§3.4, 17). AI training data use is prohibited without a licence (GEL §3.6). Provided "AS IS" without warranties (GEL §12). Governing law: Republic of India; DIAC arbitration, New Delhi seat (GEL §14). Attribution (GEL §6): GAISSE Ecosystem developed by ODA3 Institute, licensed under GEL v1.0.

1. Why AI Incident Response Matters in Healthcare

Healthcare AI — diagnostic imaging analysis, clinical decision support, patient risk stratification and administrative automation — operates where errors can directly harm patients. When an AI system fails, healthcare organisations may face several overlapping reporting regimes at once: device adverse-event reporting, data-breach notification and, in the EU, AI-specific incident reporting.

AI-IRF™ v1.0 provides incident response controls — detection, containment, rollback, evidence preservation and regulatory routing — for those failures. This note shows where those controls fit

alongside the medical device, data protection and AI regimes healthcare organisations already follow.

2. Regulatory Landscape

2.1 United States — FDA

- **Device reporting (21 CFR Part 803):** manufacturers must report device-related deaths, serious injuries and certain malfunctions, generally within 30 calendar days, and within 5 work days where remedial action is needed to prevent an unreasonable risk of substantial harm. AI functions in regulated devices are within scope.
- **Predetermined Change Control Plans (PCCPs):** FDA's final guidance on PCCPs for AI-enabled device software functions (December 2024) lets manufacturers obtain authorisation in advance for specified future model modifications. A PCCP is optional, not mandatory.
- **Good Machine Learning Practice (GMLP):** ten guiding principles published by FDA, Health Canada and the UK MHRA in 2021, covering data management, model development, testing and monitoring.
- **Cybersecurity (FD&C Act section 524B):** manufacturers of "cyber devices" must submit a software bill of materials and plans to monitor and address post-market cybersecurity vulnerabilities.
- **Premarket pathways:** 510(k) submissions demonstrate substantial equivalence to a predicate device; De Novo requests are for novel low-to-moderate-risk devices without a predicate.

2.2 United States — HIPAA

- **Security Rule technical safeguards (45 CFR 164.312):** access control, audit controls, integrity, person or entity authentication and transmission security apply to AI systems that create, receive, maintain or transmit electronic protected health information (PHI).
- **Minimum necessary standard:** uses and disclosures of PHI, including for AI purposes, are limited to the minimum necessary, subject to the Rule's exceptions.
- **Business associate agreements:** required where AI platforms or service providers handle PHI on a covered entity's behalf.
- **Breach Notification Rule:** breaches of unsecured PHI — including through AI-specific attacks such as training-data extraction — require notification to affected individuals without unreasonable delay and within 60 days of discovery, with HHS and media notification where 500 or more individuals are affected.

2.3 European Union — MDR/IVDR and the AI Act

- **EU MDR (2017/745) and IVDR (2017/746):** software, including AI, can be a medical device. Manufacturers must maintain post-market surveillance (MDR Article 83) and report serious incidents (MDR Article 87) — within 15 days generally, 2 days for a serious public health threat, and 10 days for death or an unanticipated serious deterioration in health.
- **EU AI Act classification:** AI systems that are medical devices, or safety components of them, under the MDR or IVDR are high-risk through Article 6(1) and Annex I where third-party conformity assessment applies — not through Annex III. Annex III separately covers, among others, AI used to evaluate and classify emergency calls, dispatch emergency services and triage patients in emergency healthcare.

- Requirements and reporting:** high-risk systems must meet Articles 9–15 (risk management, data governance, documentation, record-keeping, transparency, human oversight, accuracy and robustness), with deployer obligations in Article 26. For MDR/IVDR devices, Article 73(10) limits AI Act serious-incident reporting to fundamental-rights infringements; other serious incidents are reported under the MDR/IVDR. Application dates depend on the operative AI Act timetable.

2.4 India — CDSCO

Software with a medical purpose is regulated as a medical device under India's Medical Devices Rules, 2017. On 21 October 2025, the Central Drugs Standard Control Organisation (CDSCO) issued a draft Guidance Document on Medical Device Software covering classification, technical documentation, quality management, post-market surveillance and change management for AI/ML-based software. The draft clarifies how the existing Rules apply; it does not create new AI-specific requirements, and it has not been finalised.

3. AI-IRF™ v1.0 Control Mapping

Control definitions are those in Section 7 of ODA3-2026-06-TCR-HAI-001. A mapping shows where a control can contribute evidence; it does not establish compliance with the mapped obligation.

AI-IRF™ v1.0 controls	Related obligations	How the controls contribute
Regulatory operations (REG-1–4)	21 CFR Part 803; HIPAA Breach Notification Rule; EU MDR Art. 87; EU AI Act Art. 73	A jurisdictional decision tree routes each AI incident to the applicable device, breach and AI reporting regimes, with deadline calculation and retained evidence
Detection, triage and severity (MSR-1–5, REL-1–3, SEV-CLS, SEV-RPT)	EU MDR Art. 83 post-market surveillance; FDA post-market obligations	Deterministic, versioned severity and reportability classification supports consistent, timely decisions on whether an event is reportable
Model lifecycle and supply chain integrity (MLC-1–5, BOM-1–5)	FDA PCCP guidance; FD&C Act §524B; EU MDR Annex I (software and IT security)	Signed model identifiers, artefact hashes, runtime manifest validation, dual-approval retraining and tested rollback keep model changes traceable
Hallucination detection (HAL-1–5)	EU AI Act Arts. 14–15; FDA clinical decision support guidance	Confidence thresholds with human review, parallel verification for medical outputs and domain fact-checking for diagnosis support
Content safety and bias (CSS-5–7)	EU AI Act Art. 10; HHS Section 1557 rule (45 CFR 92.210)	Monthly bias analysis, quarterly demographic impact assessment and vulnerable-population assessment for high-risk systems
Agentic permissions and zero trust (APA-1–5, ZTA-1–4)	HIPAA Security Rule §164.312; EU AI Act Art. 14	Workload identity, per-request authorisation, audit trails and human approval for PII access and other high-impact actions
Discovery (SAI-1–5)	HIPAA Security Rule; minimum necessary standard	Discovery of unsanctioned AI tools, with data-loss prevention on AI egress points where PHI could leave the organisation

AI-IRF™ v1.0 controls	Related obligations	How the controls contribute
Physical AI safety (PHY-1–4)	EU MDR Annex I; FDA device safety requirements	Physical-harm playbooks, joint cyber-safety command and fail-safe modes for AI that controls devices such as surgical or rehabilitation robots

4. Sector-Specific Control Guidance

4.1 Diagnostic Imaging AI

- **HAL-2 (confidence and human review):** outputs below the 0.7 confidence threshold for consequential decisions require human review, and uncertainty should be shown to the reporting clinician.
- **HAL-3 (parallel verification):** high-stakes medical outputs are checked against a second model or rule-based verifier, with divergence above 10% flagged for review.
- **CSS-6 (demographic impact assessment):** imaging AI has documented performance differences across patient groups. High-risk systems need a quarterly assessment of demographic parity and equalised odds.
- **MLC-2 and MLC-3 (rollback):** if a model update degrades performance, the previous signed version can be restored within 15 minutes, with validation gates before clinical use resumes.

4.2 Clinical Decision Support AI

- **SEV-CLS and SEV-RPT (classification and reportability):** missed alerts — for example, sepsis or deterioration that the system failed to flag — should be classified and assessed for reportability with a deterministic, versioned method.
- **HAL-4 (error-rate tracking):** track the error rate through clinician correction feedback, alerting above 3% and escalating to corrective action above 5%.
- **MLC-5 (drift-triggered retraining):** drift should trigger a dual-approval retraining workflow. Where a PCCP applies, changes must stay within its authorised scope.

4.3 Health Data Platforms

- **SAI-1 to SAI-5 (discovery):** detect unsanctioned AI tools used with PHI, apply data-loss prevention to AI egress and quarantine rogue instances within 15 minutes of classification.
- **BOM-3 and BOM-4 (supplier assurance):** AI suppliers handling PHI need vendor attestation alongside the business associate agreement, plus incident playbooks with vendor notification within 24 hours.
- **REG-4 (evidence retention):** retain incident evidence with cryptographic integrity. HIPAA requires required documentation to be kept for six years, and longer periods may apply under state or clinical record rules.

4.4 AI-Enabled Medical Devices and Robotics

- **BOM-1 and BOM-5 (signed manifests):** signed model manifests validated on every load support the software bill of materials required for cyber devices under §524B.
- **PHY-1 to PHY-4 (physical safety):** AI that moves or actuates — surgical, rehabilitation or pharmacy robotics — needs physical-harm playbooks, joint cyber-safety command, fail-safe modes and safe-stop on attestation failure.

5. Implementation Pathway

Step	Activity	Healthcare context
1	Scope definition	Identify AI systems that are regulated devices, that process PHI, or that support clinical decisions. Prioritise those with direct patient-safety implications.
2	Reporting map	For each system, map which regimes an incident could trigger: 21 CFR Part 803, HIPAA breach notification, EU MDR Article 87 and, for non-device AI, EU AI Act Article 73. Configure the REG-1 decision tree accordingly.
3	Preparation	Complete PREP-1 to PREP-7, including pre-approved playbooks, designated human approval authority and rollback drills.
4	Clinical integration	Involve clinical informatics, biomedical engineering and the medical device vigilance function alongside security and legal teams.
5	Exercise	Run tabletop exercises for a missed-alert scenario and a PHI exposure through an AI system; retain the records as evidence.
6	Use in existing programmes	Feed AI-IRF™ evidence into ISO 13485 corrective action, post-market surveillance files and HIPAA risk analysis, as each regime permits.

6. Frequently Asked Questions

Question	Answer
Does applying AI-IRF™ provide regulatory clearance?	No. FDA clearance or approval, CE marking under the MDR/IVDR and CDSCO licensing are separate regulatory processes. AI-IRF™ controls can produce incident response evidence that supports post-market obligations; they do not substitute for regulatory submissions.
Do AI-IRF™ controls cover HIPAA compliance?	They contribute to the AI-specific aspects of the Security Rule and breach response — access control, audit trails, supplier assurance and incident handling — but do not replace a HIPAA compliance programme.
Can a hospital apply AI-IRF™ to internally developed clinical AI?	Yes. AI-IRF™ v1.0 applies to organisations that develop, deploy or procure AI systems, including in-house predictive models for sepsis, deterioration or patient flow.
How does AI-IRF™ relate to ISO 13485?	They are complementary. ISO 13485 covers the quality management system for medical device design and manufacture. AI-IRF™ v1.0 covers incident response for AI systems, which can feed the quality system's corrective-action and post-market processes.

7. Notably Absent

- No certification or audit service is operating. Certification is not currently offered. This note describes control implementation only.
- No regulatory, legal or clinical determination under FDA rules, HIPAA, the EU MDR/IVDR, the EU AI Act or CDSCO rules is made or implied.
- No claim that CDSCO has issued final or AI-specific requirements; its software guidance remains in draft.

- No claim that regulators, notified bodies or accreditation bodies such as the Joint Commission or NABH accept AI-IRF™ evidence in any particular form.
- No regulator or standards-body endorsement of AI-IRF™ v1.0 is claimed.

GAISSF Ecosystem developed by ODA3 Institute, licensed under GEL v1.0. | © 2026 ODA3 Pvt Ltd | oda3.org