FERZ INC.

AI GOVERNANCE EXECUTIVE GUIDE SERIES · VOLUME 2

AI Governance for Healthcare

Clinical AI Safety, FDA Compliance, and Patient Protection

A guide for FDA SaMD alignment, HIPAA and ONC compliance, and enterprise clinical AI risk management.

> *Governance is enforced at the authorization boundary: the point where an output becomes an effect-bearing action, and where evidence, not explanations, must exist.*

PUBLISHER
FERZ, Inc.

VERSION
1.1 · June 2026

FOR

Chief Medical Officers · Chief Information Officers · Heads of Clinical AI · Compliance Officers · Clinical Informaticists · Quality and Patient Safety Officers

FRONT MATTER

License and Terms of Use

Copyright and License

This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). You are free to share (copy and redistribute the material in any medium or format) and adapt (remix, transform, and build upon the material for any purpose, including commercial purposes) under the following terms:

Attribution: You must give appropriate credit, provide a link to the license, and indicate if changes were made. Attribution must not suggest endorsement by the licensor.

The full license text is available at creativecommons.org/licenses/by/4.0.

Attribution Guidance

When citing or reusing this work, attribution must include: FERZ, Inc., “AI Governance Executive Guide: Healthcare,” 2026. Available at ferz.ai. If the material is modified, adapted, or excerpted, the attribution must clearly indicate that changes were made.

Commercial Scope Notice

This publication is provided as a public reference framework for AI governance and risk management. While reuse and adaptation are permitted under the CC BY 4.0 license, commercial services, tooling, certification programs, managed governance offerings, and implementation accelerators derived from this framework are offered separately by FERZ, Inc. or its authorized partners. This notice does not restrict rights granted under the license; it clarifies the separation between open reference material and commercial offerings.

Framework Scope and Citation Currency

This guide is a foundational governance framework, not an exhaustive policy tracker or current-law reference. It cites statutes, regulations, memoranda, standards, and guidance only where they materially affect the general governance model, compliance posture, evidence requirements, oversight structure, or audit-readiness framework. Use-case-specific, agency-specific, jurisdiction-specific, or rapidly changing requirements may apply and should be verified separately before reliance or implementation.

No Warranty or Endorsement

This document is provided “as is,” without warranties of any kind, express or implied. Use of this framework does not constitute regulatory, legal, or compliance advice, nor does it imply endorsement by FERZ, Inc. of any implementation, product, or service that references it.

Limitation of Liability

To the maximum extent permitted by law, FERZ, Inc. is not liable for any direct, indirect, incidental, consequential, special, exemplary, or other damages arising from use of, reliance on, adaptation of, or implementation based on this publication. Users are responsible for obtaining appropriate legal, regulatory, procurement, technical, and compliance review before applying any framework, checklist, clause language, or procedure described in this guide.

Canonical Reference

The authoritative version of this document is maintained by FERZ, Inc. Subsequent adaptations or derivatives should reference the original source to ensure conceptual fidelity and version traceability.

EXECUTIVE BRIEF

Executive At-a-Glance

Audience

Chief Medical Officers, Chief Information Officers, Heads of Clinical AI, Compliance Officers, Clinical Informaticists, and Quality and Patient Safety Officers at health systems, hospitals, medical device manufacturers, and healthcare technology companies.

Core Thesis

> **Governance is enforced at the authorization boundary: the point where an output becomes an effect-bearing action, and where evidence, not explanations, must exist.**

Two Governance Tracks

TRACK	PRIMARY RISK	OWNERSHIP
Decisioning AI	Clinical risk (FDA SaMD, patient safety)	Clinical Governance / Quality
Communications AI	Records and privacy (HIPAA, state laws)	Compliance / Privacy Office

Four-Tier Classification

TIER	RISK LEVEL	GOVERNANCE INTENSITY
Tier 1	Critical	Full validation: clinical testing, FDA pathway, quality committee approval
Tier 2	Elevated	Proportionate: validation, annual review, clinician oversight
Tier 3	Standard	Operational controls, data security, HIPAA compliance
Tier 4	Minimal	Acceptable use policy only

90-Day Implementation

Weeks 1 to 4 (Foundation): Clinical AI inventory sweep, gap assessment, tiering criteria, governance structure.

Weeks 5 to 8 (Critical Systems): Tier 1 focus, FDA and safety gap analysis, bias testing, validation approach.

Weeks 9 to 12 (Operationalization): Policy deployment, training, Examiner Pack compilation, mock survey.

Maturity Endpoint

Level 5: Independently Verifiable. The organization is designed to produce tamper-evident authorization artifacts for consequential clinical AI actions, with evidence sufficient for independent replay.

EXECUTIVE SUMMARY

Why Clinical AI Governance Now

Healthcare organizations face a convergence challenge: AI and machine learning systems now directly influence clinical decisions, yet existing quality and safety frameworks, designed for traditional medical devices and clinical processes, require substantial extension to address the unique characteristics of modern AI systems.

The gap is significant. FDA guidance on AI/ML-based Software as a Medical Device (SaMD) provides a foundation, and the agency continues to mature its approach to continuously learning systems. HIPAA addresses data privacy but not algorithmic decision-making. ONC has begun regulating transparency for predictive algorithms in certified health IT. The Joint Commission evaluates quality systems but has limited AI-specific standards. Organizations must interpret and extend these frameworks themselves.

This creates both risk and opportunity. Organizations that proactively extend their clinical governance frameworks to address AI gain regulatory credibility, reduce patient safety risk, and establish operational advantages. Those that wait face increasing scrutiny, potential adverse events, and competitive disadvantage as AI becomes standard of care.

Two Governance Tracks

AI governance in healthcare operates across two distinct tracks, each with different regulatory drivers and organizational ownership:

TRACK	PRIMARY RISK	KEY REGULATIONS	TYPICAL OWNERSHIP
Decisioning AI	Clinical / patient safety	FDA SaMD, QMSR (21 CFR 820), state practice acts	Clinical Governance / CMO
Communications AI	Privacy / records	HIPAA, 21 CFR Part 11, state privacy laws	Compliance / Privacy Office

Decisioning AI: Clinical decision support, diagnostic AI, treatment recommendations, risk stratification, resource allocation, triage prioritization. Governed primarily through clinical validation, FDA pathways, and patient safety frameworks.

Communications AI: Patient communication drafts, clinician copilots, discharge summaries, prior authorization assistance, adverse event reporting, clinical documentation. Governed primarily through HIPAA, medical records requirements, and professional standards.

While communications AI may not directly execute clinical decisions, its regulatory risk profile can be equal or higher due to HIPAA, medical records integrity, professional liability, and patient trust exposure. Do not assume that communications use cases are inherently lower risk than decisioning use cases.

This distinction matters because surveyors and regulators ask different questions, and internal ownership often differs. A single AI system may span both tracks. A clinician copilot that drafts clinical notes

(records) based on diagnostic analysis (clinical) requires coordinated governance.

The Core Challenge

Healthcare AI governance has three pillars: Development, Validation, and Governance. Each requires extension for AI:

PILLAR	TRADITIONAL APPROACH	AI / ML EXTENSION REQUIRED
Development	Device design controls, clinical protocols	Training data provenance, algorithm selection, bias assessment, explainability approach
Validation	Clinical trials, bench testing	Demographic subgroup testing, drift monitoring, adversarial robustness, real-world performance
Governance	Quality committees, credentialing	Expanded inventory (embedded AI, vendor), new risk categories, continuous monitoring

What This Guide Provides

Regulatory mapping: How FDA SaMD guidance, the QMSR, HIPAA, ONC transparency requirements, CMS requirements, and the EU AI Act apply to clinical AI systems.

Risk framework: Four-tier classification system with explicit Tier 1 triggers to prevent governance arbitrage.

Governance structure: Clinical governance model extended for AI oversight.

Lifecycle management: Development, validation, deployment, monitoring, and retirement processes.

Survey readiness: Documentation standards, Examiner Pack template, and response playbooks.

GenAI controls: Dedicated control set for generative AI including hallucination, documentation integrity, and disclosure.

Board metrics: KRIs and KPIs for executive clinical AI reporting.

Implementation roadmap: 90-day plan with maturity progression to independently verifiable governance.

PART I

The Regulatory Landscape

CHAPTER 1

Existing Frameworks That Apply Today

Healthcare organizations do not operate in a regulatory vacuum when deploying AI. Multiple existing frameworks apply, even in the absence of comprehensive AI-specific regulation. Understanding these frameworks, and how regulators are extending them to AI, is essential for defensible governance.

FDA AI/ML Guidance: The Foundation

The FDA has issued guidance on AI/ML-based Software as a Medical Device (SaMD), establishing expectations for AI systems that meet the definition of a medical device. The framework continues to mature, and it provides the foundation for clinical AI governance.

Key principle: FDA defines SaMD as software intended to be used for medical purposes without being part of a hardware medical device. AI systems that diagnose, treat, cure, mitigate, or prevent disease may meet this definition regardless of whether they are marketed as decision support.

The FDA's Total Product Lifecycle Approach

FDA promotes a total product lifecycle (TPLC) approach to AI-enabled devices, supported by a maturing set of guidance documents:

1. **Good Machine Learning Practice (GMLP):** The FDA, Health Canada, and the UK MHRA jointly published GMLP guiding principles, adopted in final form by the International Medical Device Regulators Forum (IMDRF) in January 2025. The principles cover multidisciplinary expertise, representative data, independent datasets, reference standards, model design tailored to data, human-AI team performance, clinical study testing, transparency to users, monitoring of deployed models, and cybersecurity safeguards.
2. **Predetermined Change Control Plans (PCCP):** FDA has issued final guidance on the information a manufacturer should include in a Predetermined Change Control Plan (PCCP) for AI-enabled device software functions, allowing manufacturers to pre-specify the types of changes an AI-enabled device may undergo without a new submission, enabling continuous improvement while maintaining safety.
3. **AI-enabled device lifecycle guidance:** FDA issued comprehensive draft guidance in January 2025 on lifecycle management and marketing submissions for AI-enabled device software functions, addressing transparency and bias across the total product lifecycle. This guidance remains draft and is expected to be finalized.

Risk Classification for SaMD

FDA uses a risk-based framework drawn from the IMDRF:

State of the healthcare situation: Critical, Serious, or Non-serious.

Significance of the information: Treat or diagnose, Drive clinical management, or Inform clinical management.

Higher-risk combinations require more rigorous validation and a more demanding regulatory pathway (510(k), De Novo, or PMA).

The boundary between regulated SaMD and exempt clinical decision support (CDS) is nuanced. Under the 21st Century Cures Act, which amended Section 520(o)(1)(E) of the Federal Food, Drug, and Cosmetic Act, certain CDS software functions are excluded from the device definition when they meet four statutory criteria, including that the software enables a health care professional to independently review the basis for the recommendation. Many AI systems do not qualify. When regulatory status is ambiguous, treat the system as Tier 1 until a formal determination is obtained.

QMSR: Quality Management for Devices (formerly 21 CFR Part 820)

FDA's final rule amending the device Quality System Regulation took effect February 2, 2026, replacing most of the prior 21 CFR Part 820 requirements with the Quality Management System Regulation (QMSR), which incorporates ISO 13485 by reference. For AI-enabled devices, design controls, documentation, and quality system expectations now run through the QMSR and ISO 13485 framework. Legacy references to "21 CFR 820" should be read as the QMSR.

21 CFR Part 11: Electronic Records

21 CFR Part 11 establishes requirements for electronic records and signatures in FDA-regulated contexts. For clinical AI, this means:

Audit trails: Complete, timestamped records of system actions and changes.

Access controls: Role-based permissions with authentication.

Data integrity: Prevention of unauthorized modifications.

Electronic signatures: Legally binding sign-offs for approvals.

ONC HTI-1: Transparency for Predictive Algorithms

ONC's Health Data, Technology, and Interoperability (HTI-1) final rule established the first federal transparency requirements for AI and other predictive algorithms in certified health IT. This is HHS's first attempt to regulate AI tools beyond those FDA regulates as medical devices, and it is current law, not a future development.

Decision Support Interventions (DSI): HTI-1 replaced the prior Clinical Decision Support criterion with a Decision Support Intervention criterion at 45 CFR 170.315(b)(11), adding a predictive DSI category covering AI and machine learning.

Source attributes: Developers that supply predictive DSIs must disclose a defined set of source attributes describing how the intervention was developed, the data used to train it, intended use, validation, and known limitations.

Risk management: Developers must conduct and attest to risk management for predictive DSIs, addressing validity, fairness, robustness, and how data are acquired, managed, and used.

Base EHR definition: DSI became part of the Base EHR definition as of January 1, 2025, so the transparency criterion now reaches the certified systems used across most hospitals and clinicians.

The DSI source attributes and risk-management expectations map directly onto the evidence this guide calls for: provenance of inputs, validation, fairness, and a record of how a predictive output was produced.

HIPAA: Privacy and Security

HIPAA's Privacy and Security Rules apply to AI systems that process protected health information (PHI):

Minimum necessary: AI systems should access only the PHI required for their function.

Business Associate Agreements: Third-party AI vendors must sign BAAs.

Security safeguards: Technical, administrative, and physical protections for PHI in AI systems.

Breach notification: AI-related breaches trigger HIPAA notification requirements.

CMS Conditions of Participation

Hospitals participating in Medicare and Medicaid must meet CMS Conditions of Participation, which increasingly implicate AI:

Quality Assessment and Performance Improvement (QAPI): AI systems affecting patient outcomes fall under QAPI oversight.

Medical staff credentialing: Clinicians using AI may need competency documentation.

Patient rights: Patients may have rights to information about AI use in their care.

The Joint Commission

Joint Commission standards for accredited organizations increasingly apply to clinical AI:

Leadership standards: Oversight of technology affecting patient care.

Performance improvement: Monitoring and improving AI system performance.

Information management: Integrity of clinical information including AI outputs.

State Medical Practice Acts

State medical boards regulate the practice of medicine, which may include AI-assisted clinical decisions:

- AI may not independently practice medicine without licensed clinician oversight.
- Clinicians remain responsible for AI-informed decisions.
- Some states are developing specific AI guidance for practitioners.

The standard of care remains the clinician's responsibility. Use of AI does not shift the duty of care from the clinician to the technology or vendor. Clinicians must exercise independent professional judgment, and AI outputs should inform, not replace, clinical decision-making.

CHAPTER 2

What's Coming

The regulatory landscape for AI in healthcare is evolving rapidly. Organizations should anticipate and prepare for emerging requirements rather than waiting for final rules.

EU AI Act: Extraterritorial Impact

The EU AI Act entered into force in 2024 and applies in phases rather than on a single activation date. It has extraterritorial reach affecting U.S. healthcare organizations that serve EU patients or deploy AI systems whose outputs are used in the EU.

High-risk classification: Medical devices and in vitro diagnostic devices that incorporate AI are classified as high-risk, requiring conformity assessments, risk management systems, data governance, technical documentation, human oversight, accuracy and robustness measures, and registration. For AI embedded in products already regulated under the Medical Device Regulation and In Vitro Diagnostic Regulation, the high-risk obligations under Article 6(1) and Annex I are tied to an extended medical-device transition, currently August 2, 2027, and are layered on top of MDR and IVDR rather than replacing them. A Digital Omnibus proposal announced in November 2025 may adjust these dates. Plan against the phased timeline, not a single enforcement date.

Anticipated U.S. Regulatory Evolution

While the specific regulatory trajectory depends on administration priorities, several developments are likely:

FDA AI lifecycle guidance: Finalization of the January 2025 draft guidance on lifecycle management and marketing submissions for AI-enabled device software functions, and continued development of the TPLC and PCCP framework.

HIPAA Security Rule modernization: OCR published a Notice of Proposed Rulemaking in January 2025 that would strengthen the Security Rule for the first time since 2013, including making implementation specifications mandatory and adding specific cybersecurity controls. As a proposed rule, it is not a current obligation; the existing Security Rule remains in effect until any final rule is issued.

CMS payment and cybersecurity policy: Reimbursement considerations for AI-assisted care, and anticipated cybersecurity conditions of participation for Medicare and Medicaid providers.

State legislation: Growing state-level AI transparency and accountability requirements.

Importantly, supervisory expectations often evolve through survey practice and enforcement actions before formal rulemaking occurs. Organizations should not wait for explicit AI regulation to implement governance frameworks.

PART II

Risk Framework for Healthcare AI

CHAPTER 3

AI System Classification

Expanding the Clinical AI Inventory

The first challenge in AI governance is defining what counts as a clinical AI system subject to governance. Traditional inventories focused on formal medical devices. AI governance requires expanding scope significantly.

For governance purposes, any system that materially influences clinical decisions, patient care, or regulated communications is treated as subject to governance, regardless of technical implementation. This avoids definitional debates that delay governance while risk accumulates.

What counts as a clinical AI system:

- **Diagnostic AI:** Radiology interpretation, pathology analysis, symptom checkers, risk prediction.
- **Treatment AI:** Clinical decision support, drug interaction checking, dosing recommendations, care pathway optimization.
- **Operational AI:** Triage prioritization, resource allocation, scheduling optimization, readmission prediction.
- **Generative AI:** Clinical documentation, patient communication, prior authorization, discharge instructions.
- **Embedded AI:** AI features within EHRs, medical devices, monitoring systems, and vendor platforms.
- **Third-party AI services:** APIs, cloud services, remote interpretation, outsourced analytics.

The Four-Tier Framework

Healthcare AI governance requires a four-tier system that calibrates oversight intensity to patient safety risk. This framework aligns with FDA's risk-based approach while extending to non-device AI.

TIER	CRITERIA	EXAMPLES	REGULATORY TOUCHPOINTS
Tier 1 (Critical)	Direct patient impact, diagnostic or treatment decisions, safety-critical	Diagnostic imaging AI, sepsis prediction, treatment recommendations, surgical robotics	FDA SaMD, QMSR, clinical validation, EU AI Act high-risk
Tier 2 (Elevated)	Clinical workflow impact, patient-facing with clinician review	Triage prioritization, readmission risk, clinical documentation assist	HIPAA, CMS quality reporting, ONC DSI transparency, medical records integrity

TIER	CRITERIA	EXAMPLES	REGULATORY TOUCHPOINTS
Tier 3 (Standard)	Operational efficiency, no direct patient decisions	Scheduling optimization, supply forecasting, staff allocation	Operational controls, data security
Tier 4 (Minimal)	Personal productivity, no PHI, no clinical context	General research, administrative drafting, education	Acceptable use policy only

AUTOMATIC TIER 1 TRIGGERS

Tiering frameworks fail when departments tier down to avoid validation overhead. To prevent governance arbitrage, the following conditions automatically classify a system as Tier 1, regardless of other assessments:

1. Impacts clinical decisions, diagnosis, treatment recommendations, or patient safety.
2. Drives or materially influences quality reporting (CMS measures, Joint Commission, Leapfrog).
3. Produces outputs that become part of the medical record, even if clinician reviewed.
4. Involves protected health information in clinical decisioning logic.
5. Feeds into patient safety or clinical quality calculations.

Third-Party Clinical AI Risk

Third-party clinical AI presents unique governance challenges. Organizations remain responsible for the safety and effectiveness of AI they deploy, regardless of vendor claims.

Vendor Minimum Evidence Package

To make third-party oversight actionable, require the following evidence from vendors, scaled by tier:

EVIDENCE REQUIRED	TIER 1	TIER 2	TIER 3
Regulatory status / clearance documentation	Required	Required	Requested
Intended use statement and contraindications	Required	Required	Requested
Training data demographics and representativeness	Required	Required	N/A
Clinical validation study results	Required	Required	N/A
Bias testing across demographic subgroups	Required	Risk-based	N/A
Predictive DSI source attributes (where applicable)	Required	Required	Requested
Post-market surveillance data / incident history	Required	Required	Requested
BAA and audit cooperation language	Required	Required	Required

This is the difference between “we have a vendor risk policy” and “we can survive a survey.”

CHAPTER 4

AI-Specific Risk Categories

Clinical AI systems present risk categories that extend beyond traditional medical device risk. Effective governance requires understanding and addressing each category.

Clinical Risk (Extended)

Training Data Risk: AI model quality depends fundamentally on training data. Risks include non-representative patient populations leading to biased performance, historical data encoding past disparities, data from different clinical settings not generalizing, and incomplete documentation of data provenance.

Diagnostic and Prognostic Risk: AI may produce false positives leading to unnecessary interventions or false negatives leading to missed diagnoses. Performance may vary across patient subgroups.

Drift Risk: Clinical AI can degrade over time as patient populations change, practice patterns evolve, or EHR documentation practices shift. Both data drift and concept drift require detection and remediation.

Explainability Risk: Clinicians may be unable to assess AI recommendations without understanding the reasoning. Black-box AI may undermine clinical judgment and informed consent.

Operational Risk

Workflow integration failures: AI systems failing to integrate properly with clinical workflows.

Alert fatigue: Excessive AI alerts leading to clinician desensitization.

Latency and availability: AI not meeting performance requirements for time-sensitive clinical decisions.

Rollback capabilities: Inability to quickly revert to non-AI workflows when issues arise.

Compliance Risk

HIPAA violations: Unauthorized access to or disclosure of PHI through AI systems.

Medical records integrity: AI-generated content compromising record accuracy or completeness.

Professional liability: Clinician responsibility for AI-informed decisions.

Informed consent: Patient rights to know about AI involvement in their care.

Cybersecurity and Data Risk

Training data poisoning: Adversaries corrupting training data to influence clinical AI behavior.

Adversarial attacks: Inputs crafted to cause AI misclassification, such as adversarial images.

PHI leakage: Sensitive information exposed through AI outputs or model inversion.

Prompt injection: Manipulating generative AI through crafted inputs to bypass safety rails.

CHAPTER 5

Generative AI Controls

Generative AI presents unique control requirements beyond traditional clinical AI governance. Because GenAI can produce novel outputs, such as clinical notes, patient communications, and discharge instructions, that may enter the medical record or influence care, a dedicated control set is required.

Data Security Controls

PHI input filtering: Prevent unnecessary PHI from entering GenAI prompts; apply minimum necessary principles.

Output monitoring: Detect and prevent PHI leakage or hallucinated patient information in outputs.

Prompt injection defense: Technical controls against manipulation attempts that could generate harmful clinical content.

Session isolation: Ensure user sessions do not leak patient information across conversations.

Clinical Documentation Controls

For AI-generated clinical documentation, medical records integrity requirements apply:

Draft retention: Capture AI-generated drafts as well as final documentation.

Clinician attestation: Require explicit clinician review and sign-off before documentation enters the record.

Authorship indication: Clear indication when AI assisted in generating documentation.

Hallucination detection: Processes to identify fabricated clinical details such as medications, allergies, or history.

Use Case Governance

CATEGORY	EXAMPLES	REQUIRED CONTROLS
Approved	Literature summarization, education, research drafts	Acceptable use acknowledgment, no PHI
Restricted	Clinical note drafts, patient message drafts, discharge instructions	Clinician review, documentation, approval workflow
Prohibited	Autonomous diagnosis, unreviewed patient communications, treatment decisions	Technical enforcement, exception process only

Hallucination and Accuracy Controls

Generative AI may produce plausible but incorrect clinical content. Specific controls include:

- Explicit warnings that AI-generated content must be verified.

- Structured output validation against EHR data where possible.
- Sampling and audit processes for AI-generated documentation.
- Incident reporting for identified hallucinations with clinical impact.

Human Approval Requirements

Define approval gates for clinical GenAI outputs:

- All clinical documentation requires clinician review and attestation before entry to the medical record.
- Patient-facing communications require clinician approval before sending.
- Prior authorization content requires appropriate clinical sign-off.
- Adverse event reports require qualified reviewer approval.

Human review must be meaningful, not rubber-stamp. Reviewers must have the authority, training, and time to override or block AI outputs. Checkbox acknowledgments without substantive review do not satisfy clinical oversight requirements.

> **CLINICIAN ATTESTATION AND THE AUTHORIZATION BOUNDARY**
>
> The runtime authorization boundary governs whether an effect-bearing clinical action may release. Clinician attestation remains mandatory wherever law, clinical policy, medical-record rules, or patient-safety governance requires it. In those cases, clinician sign-off is not a substitute for the authorization boundary; it is a required input to the authorization verdict and must be captured in the authorization artifact.
>
> In practice, the attestation becomes part of the evidence rather than a step beside it. The authorization artifact records who attested, to what, and under which clinical policy, so a later reviewer can confirm not only that a clinician signed off, but that the sign-off was a required and satisfied condition of release.

PART III

Governance Structure

CHAPTER 6

Organizational Model

Clinical AI Governance Structure

Healthcare organizations operate under quality and safety governance structures. AI governance must integrate into these structures rather than creating parallel governance.

FUNCTION	TRADITIONAL ROLE	AI GOVERNANCE EXTENSION
Clinical Departments (First Line)	Own clinical quality and safety in daily operations	AI system ownership, use case documentation, performance monitoring, incident reporting, clinician training
Quality and Compliance (Second Line)	Set policy, monitor quality, independent review	AI policy and standards, independent validation coordination, inventory management, regulatory liaison
Internal Audit (Third Line)	Independent assurance of framework effectiveness	AI governance framework audits, control testing, compliance verification, AI-specific audit techniques

Key Roles

Chief Medical Officer / Chief Medical Informatics Officer: Executive accountability for clinical AI safety and effectiveness.

Clinical AI Governance Lead: Dedicated focus on AI-specific clinical oversight, bias monitoring, and validation coordination.

Clinical AI Validation Team: Specialized validation capability with skills in clinical testing, demographic subgroup analysis, and performance monitoring.

Department AI Champions: Clinician leads who understand both clinical context and AI capabilities in their specialty.

Governance Bodies

Medical Executive Committee: Ultimate clinical oversight of Tier 1 AI deployments; charter should explicitly include AI.

Clinical AI Subcommittee: Working body that reviews AI use cases, validation results, and monitoring data; reports to the MEC.

Pharmacy and Therapeutics (for drug-related AI): Oversight of AI affecting medication decisions.

Health IT Steering Committee: Technical and operational oversight of AI implementations.

Governance Authority and Capability Requirements

Clinical AI governance frameworks fail in two predictable ways: the governance function lacks authority to enforce standards, or it lacks competence to make credible decisions. Both failure modes must be addressed by design.

The Authority Problem

Governance without authority produces policy documents that IT ignores and clinical departments route around. Effective clinical AI governance requires:

Pipeline integration: Governance must have hooks into actual deployment processes: intake forms that trigger before vendor contracts, validation gates that block EHR integration, and monitoring access with real clinical performance data. Policy without pipeline integration is theater.

Escalation path with teeth: When governance raises patient safety concerns, there must be an escalation path to someone who can override IT delivery pressure and departmental adoption pressure. This typically means reporting to the CMO, a Board Quality Committee, or having direct access to the CEO on material patient safety issues.

Information rights: Even where governance cannot block deployments, it must have visibility into all clinical AI systems. Shadow AI in clinical workflows is a governance design failure, not just a clinician behavior problem. Information rights include access to Health IT inventories, vendor contract notifications, EHR integration logs, and clinical decision support configurations.

Budget independence: Governance cannot depend on the goodwill of functions it oversees. Validation resources, monitoring tools, and external clinical expertise must be funded independently of IT or departmental budgets.

The Competence Problem

Authority without competence produces either rubber-stamp approvals or reflexive blocking. Neither serves patients. Effective clinical AI governance requires dual competency:

Clinical and regulatory expertise: Deep understanding of FDA pathways, HIPAA, ONC transparency requirements, clinical workflow integration, and patient safety frameworks. Ability to translate regulatory requirements into operational controls that clinicians will actually follow.

Technical credibility: Sufficient understanding of AI and ML to evaluate vendor validation claims, assess demographic bias testing, and recognize when technical teams are deflecting or oversimplifying. This does not require the governance lead to be a data scientist, but they must be able to hold their own in technical discussions with vendors and informatics teams.

In practice, this dual competency usually requires a team rather than a single individual, often pairing clinical informaticists with quality and regulatory specialists under unified governance leadership. The CMIO role, where it exists, sometimes bridges these domains.

Reporting Line Tradeoffs

Where the clinical AI governance function reports determines its capture risk:

REPORTS TO	ADVANTAGE	CAPTURE RISK
CMO / Chief Medical Officer	Clinical authority, patient safety mandate	May lack technical access; dependent on IT cooperation
CIO / CMIO	Technical access, EHR integration easier	Captured by implementation pressure; asked to approve what is already contracted
CNO / Clinical Departments	Clinical workflow context, adoption focus	Captured by departmental pressure; governance becomes checkbox
Compliance / Legal	Regulatory credibility, independence	May lack clinical and technical depth; legalistic rather than operational
CEO / Board Quality Committee	Maximum independence and authority	May lack day-to-day clinical connection; reserved for largest systems

The most effective structures typically have clinical AI governance reporting to the CMO with a strong dotted line to the CIO or CMIO for technical liaison, or embedded within an empowered Quality function with direct Board Quality Committee access for Tier 1 patient safety escalations.

The Effective Challenge Test

Effective governance requires critical analysis by objective, informed parties. For clinical AI governance, apply this test to your organizational design:

- Can your governance function say no to a Tier 1 clinical AI deployment that a department chair wants to launch?
- Can your governance function obtain validation data that a vendor does not want to provide?
- Does your governance function learn about clinical AI initiatives before the vendor contract is signed, or after go-live?
- Can your governance function compel remediation of patient safety findings, or only recommend?
- Does your governance function have the clinical and technical vocabulary to challenge vendor performance claims?

If the answer to any of these is no, your governance structure has a design flaw that will manifest during a survey, an FDA inspection, or a patient safety event.

CHAPTER 7

Policy Framework

Effective AI governance requires a comprehensive policy framework that provides clear guidance while remaining flexible enough to accommodate evolving AI capabilities.

Core Policies Required

1. **Clinical AI Acceptable Use Policy.** Permitted uses by role and clinical context; prohibited uses (autonomous diagnosis, unreviewed patient communications); shadow AI prohibition, with no unapproved AI in clinical workflows; and PHI handling requirements in AI systems.

2. **Clinical AI Development and Acquisition Policy.** Development lifecycle requirements by tier; documentation standards (training data, validation approach, intended use); testing and validation gates before clinical deployment; and vendor due diligence requirements.
3. **Clinical AI Validation Policy.** Validation scope and frequency by tier; validation techniques (clinical accuracy, demographic subgroup testing, real-world performance); independence requirements for validation; and findings classification and remediation timelines.
4. **Clinical AI Monitoring Policy.** Continuous monitoring requirements (performance metrics, drift detection); periodic review cadence by tier; escalation triggers and clinical incident reporting; and post-market surveillance procedures.
5. **Third-Party Clinical AI Policy.** Due diligence requirements before contracting; minimum evidence package by tier; ongoing performance monitoring; and incident notification and cooperation requirements.

PART IV

Lifecycle Management

CHAPTER 8

Development and Validation

Pre-Development Requirements

Before clinical AI development or acquisition begins, several assessments should be completed:

Clinical use case justification: Clear articulation of the clinical problem and why AI is appropriate.

Intended use statement: Specific clinical context, user, and decision support function.

Patient population definition: Who the AI will be used for, and whether training data is representative.

Regulatory pathway assessment: Determination of FDA classification and requirements.

Ethical review: For Tier 1 and Tier 2 systems, assessment of potential bias and equity concerns.

Development Documentation Standards

Clinical AI development requires thorough documentation aligned with FDA expectations:

Training data specification: Source, patient demographics, selection criteria, preprocessing, known limitations.

Algorithm documentation: Architecture selection rationale, alternatives considered, hyperparameter choices.

Clinical endpoint definition: What the AI is predicting or classifying and how ground truth was established.

Performance characterization: Sensitivity, specificity, PPV, and NPV across relevant subgroups.

Intended use and indications: Specific clinical context, user qualifications, contraindications.

Validation Framework

Validation requirements vary by tier:

VALIDATION ACTIVITY	TIER 1	TIER 2	TIER 3	TIER 4
Clinical accuracy testing	Required	Required	N/A	N/A
Demographic subgroup analysis	Required	Risk-based	N/A	N/A
Local validation on site data	Required	Required	Recommended	N/A

VALIDATION ACTIVITY	TIER 1	TIER 2	TIER 3	TIER 4
Workflow usability testing	Required	Required	Basic	N/A
Independent clinical review	Required	Required	Self-assess	N/A

CHAPTER 9

Deployment and Monitoring

Deployment Controls

Staged rollout: Phased clinical deployment with monitoring gates before full implementation.

Rollback procedures: Documented ability to revert to non-AI clinical workflows quickly.

Performance baseline: Establish baseline metrics before deployment for comparison.

Clinician training: Ensure users understand AI capabilities, limitations, and appropriate use.

Patient notification: Determine requirements for informing patients about AI use in their care.

Patient notification requirements are risk- and jurisdiction-dependent. For Tier 1 decisioning AI (diagnostic, treatment, prognostic), involve legal counsel and ethics or consent governance to determine appropriate disclosure. Some states and payers are developing specific AI disclosure requirements.

Ongoing Monitoring

REVIEW TYPE	TIER 1	TIER 2	TIER 3	TIER 4
Performance monitoring	Continuous	Monthly	Quarterly	N/A
Demographic subgroup review	Quarterly	Annual	N/A	N/A
Full revalidation	Annual	18 months	As needed	N/A
Clinical incident review	Within 24h	Within 72h	As needed	N/A

Audit Trails: Logs versus Evidence

A critical distinction exists between operational logging and authorization artifacts:

CONCEPT	DEFINITION	REGULATORY VALUE
Logs	Operational trace of system activity: timestamps, user actions, errors, performance metrics	Necessary but not sufficient; can show what happened but not whether it was clinically permitted
Authorization Artifacts	Replayable, version-bound, policy-bound authorization records with tamper-evident integrity	Independently verifiable; can answer “how do you know?” for surveyors, FDA, and malpractice defense

For Tier 1 clinical systems, the goal is to move beyond “audit logging enabled” toward tamper-evident authorization artifacts that show whether an effect-bearing clinical action was permitted before release under the applicable clinical policy, authority chain, provenance condition, and clinician-attestation condition where required. This distinction becomes critical at maturity Level 5.

Observability Is Not Authorization

Logging, monitoring, and post-hoc review can show that an event occurred. They do not establish that an effect-bearing clinical action was permitted before it released. Logs show what happened; the authorization artifact shows whether the effect-bearing clinical action was permitted before release, under the applicable clinical policy, authority chain, provenance requirements, and clinician-attestation condition. If the authorization artifact cannot be produced, the action fails closed.

Combining clinical monitoring, model-performance tracking, telemetry, and post-market or post-deployment surveillance may strengthen oversight, but it does not by itself produce a determination, made before release, of whether an effect-bearing clinical action was permitted. For effect-bearing clinical AI actions, require the authorization artifact, not only monitoring data or logs.

> **HEALTHCARE DOCTRINE ANCHOR**
>
> In clinical AI, the authorization boundary does not replace clinician judgment. It governs whether an AI-mediated clinical action may release, and it records the required clinician attestation, provenance, policy state, authority chain, and verdict in a tamper-evident authorization artifact.

PART V

Survey and Examination Readiness

CHAPTER 10

Documentation Standards

Clinical AI Documentation Package (Tier 1)

Each Tier 1 clinical AI system should have a complete documentation package containing:

- System identification (unique ID, clinical owner, vendor or developer, tier rationale).
- Clinical context (use case, workflow integration, patient population, clinical decisions affected).
- Technical documentation (algorithm, training data, performance characteristics).
- Validation documentation (scope, clinical testing results, subgroup analysis, limitations).
- Monitoring documentation (metrics, thresholds, escalation procedures, incident history).

Examiner Pack: Standard Binder Structure

For each Tier 1 clinical AI, maintain a ready-reference binder to minimize scrambling during surveys or inspections:

TAB	CONTENTS
1	Inventory entry, tier classification rationale, regulatory status
2	Intended use statement, patient population, clinical workflow
3	Validation summary with clinical testing results and sign-offs
4	Performance monitoring dashboard, thresholds, last three months of data
5	Change log, version history, algorithm updates
6	Clinical incidents, patient safety events, corrective actions
7	Vendor artifacts (if applicable): FDA clearance, validation studies, BAA

CHAPTER 11

Survey and Inspection Preparation

Common Survey and Inspection Questions

Prepare responses to common surveyor and inspector inquiries:

Inventory and Governance

- Show me your complete clinical AI inventory.
- How do you identify shadow AI in clinical workflows?
- Walk me through your clinical risk tiering methodology.
- Who has authority to approve Tier 1 clinical AI?

Validation and Safety

- Show me validation documentation for a specific AI system.
- How do you test for demographic bias in diagnostic AI?
- What is your approach to ensuring AI works for your patient population?
- How do you validate vendor AI claims before deployment?

Monitoring and Incidents

- Show me evidence of ongoing performance monitoring.
- What triggers a clinical AI safety review?
- Walk me through a recent AI-related patient safety event.
- How do you detect when AI performance degrades?

Survey Response Playbook

Documentation retrieval: Know where all clinical AI documentation lives and who can access it.

Subject matter experts: Pre-identify clinical and technical SMEs for each Tier 1 system.

Response coordination: Designate a survey coordinator to manage requests.

Issue escalation: Define an escalation path for unexpected findings.

CHAPTER 12

Board and Executive Reporting

Clinical AI Governance KRIs and KPIs

Executives adopt what they can track. The following metrics support board-level clinical AI governance reporting:

METRIC	DEFINITION	TARGET
Inventory coverage	Percent of known clinical AI use cases tiered and documented	Over 95%
Tier 1 validation currency	Percent of Tier 1 AI with current validation (within 12 months)	100%
Performance alerts	Count of performance degradation alerts by severity	Trending down
Demographic disparity findings	Count of significant subgroup performance gaps	Zero unresolved

METRIC	DEFINITION	TARGET
Vendor concentration	Percent of Tier 1 AI from a single vendor	Under 40%
Shadow AI detections	Count of unapproved clinical AI discovered	Trending down
Patient safety events	AI-related patient safety events by severity	Zero serious
Time-to-remediate	Average days from finding to resolution by severity	Under 7 days for critical

Report these metrics quarterly to the Clinical AI Subcommittee and annually, with trends, to the Medical Executive Committee and Board Quality Committee.

PART VI

Implementation Roadmap

CHAPTER 13

90-Day Implementation Plan

Phase 1: Foundation (Weeks 1 to 4)

Weeks 1 to 2: Inventory and Gap Assessment

- Conduct a clinical AI system inventory sweep across all departments.
- Map existing quality and safety frameworks to clinical AI requirements.
- Identify policy and governance gaps.
- Assess organizational readiness (clinical informatics, vendor management).

Weeks 3 to 4: Governance Structure

- Establish a Clinical AI Subcommittee or expand an existing committee charter.
- Define roles and responsibilities.
- Establish tiering criteria including automatic Tier 1 triggers.
- Draft core policies.

Phase 2: Critical Systems (Weeks 5 to 8)

Weeks 5 to 6: Tier 1 System Focus

- Identify and prioritize Tier 1 clinical AI systems.
- Conduct gap analysis against FDA and clinical validation requirements.
- Develop remediation plans.
- Begin enhanced documentation.

Weeks 7 to 8: Validation Enhancement

- Implement demographic subgroup testing for diagnostic and treatment AI.
- Establish clinical performance monitoring.
- Enhance drift detection capabilities.
- Compile vendor minimum evidence packages.

Phase 3: Operationalization (Weeks 9 to 12)

Weeks 9 to 10: Policy Deployment

- Finalize and approve policies through appropriate committees.
- Conduct training for clinical staff and informatics teams.

- Implement acceptable use controls.
- Establish shadow AI detection.

Weeks 11 to 12: Survey Readiness

- Compile Examiner Pack binders for Tier 1 systems.
- Conduct a mock survey or inspection.
- Remediate gaps.
- Establish ongoing monitoring cadence and board reporting.

CHAPTER 14

Maturity Progression

Five Levels of Clinical AI Governance Maturity

LEVEL	NAME	WHAT IT LOOKS LIKE	TIMELINE
1	Foundational	Basic inventory exists, policies drafted but inconsistently applied	Current state
2	Structured	Complete inventory with tiering, policies approved, validation for Tier 1	6 to 12 months
3	Managed	Automated inventory, consistent validation, continuous monitoring, survey-ready	12 to 24 months
4	Optimized	Real-time performance visibility, automated compliance checking, predictive safety	24+ months
5	Independently Verifiable	The organization is designed to produce tamper-evident authorization artifacts for consequential clinical AI actions, with evidence sufficient for independent replay	Infrastructure investment

Level 5: Independently Verifiable

At Level 5, the organization is designed to produce tamper-evident authorization artifacts for consequential clinical AI actions, with evidence sufficient for independent replay. The maturity endpoint is not better monitoring. It is an independently reconstructable verdict: whether the action was authorized before it executed.

This is the logical endpoint of the authorization-boundary principle: governance is enforced at the point where an output becomes an effect-bearing clinical action, and evidence, not explanations, must exist. Organizations pursuing Level 5 maturity should evaluate infrastructure capable of producing tamper-evident authorization artifacts, fail-closed release behavior, and replayable audit trails that support FDA, survey, and malpractice-defense evidence needs.

TECHNICAL FOUNDATION

The Five Tests Standard

OPEN SPECIFICATION

The evidence standard in this guide aligns with the Five Tests Standard (5TS), an open specification for verifiable AI governance. 5TS defines five requirements: Stop, Ownership, Replay, Escalation, and Provenance. Provenance means origin, not truth. It is a test, not a verdict.

The verdict space is ALLOW, DENY, and ABSTAIN. An ABSTAIN blocks execution pending authorized human override; escalation is the consequence of ABSTAIN, not a fourth verdict.

Five Tests Standard (5TS) · Open standard DOI: 10.5281/zenodo.21040295 · Explanatory foundation DOI: 10.5281/zenodo.21048661

APPENDICES

Reference Material

APPENDIX A

Clinical AI System Intake Form

Every clinical AI use case must submit this form before deployment.

FIELD	DESCRIPTION
Use Case Name	Free text, descriptive name
Clinical Owner	Name, title, department; single accountable clinician
AI System / Tool	Vendor or product name, or internally developed description
Governance Track	Decisioning AI / Communications AI / Both
Clinical Purpose	What clinical problem does this solve? What decisions or outputs?
Patient Population	Who will this AI be used for?
Output Type	Diagnostic / Treatment / Prognostic / Documentation / Communication
Data Classes	PHI / De-identified / Imaging / Genomic / Device data
Regulatory Status	FDA cleared / FDA exempt / Under review / Unknown
Automatic Tier 1 Trigger?	Yes (specify which) / No
Proposed Tier	Tier 1 / Tier 2 / Tier 3 / Tier 4
Tier Justification	Brief explanation

APPENDIX B

Tier 1 Clinical AI Control Checklist

All controls required before Tier 1 clinical AI go-live:

TIER 1 CRITICAL: ALL CONTROLS REQUIRED BEFORE CLINICAL GO-LIVE

- Designated clinical owner with documented accountability.
- Intended use statement and patient population defined.
- Automatic Tier 1 trigger assessment documented.
- FDA regulatory status determined and documented.
- Training data demographics and representativeness documented.
- Clinical validation completed with performance metrics.
- Demographic subgroup analysis completed.
- Pre-execution authorization boundary configured for effect-bearing clinical actions.
- Authorization artifact evidence configured for effect-bearing clinical actions, not only logging.
- Clinician attestation captured where law or clinical policy requires it.
- Clinical incident reporting path defined.
- Vendor minimum evidence package collected (if third-party).
- HIPAA and privacy review completed.
- Clinical AI Subcommittee or MEC approval documented.
- Examiner Pack binder compiled.
- Performance monitoring and review schedule established.

APPENDIX C

Regulatory Reference Summary

REGULATION / GUIDANCE	SCOPE	KEY CLINICAL AI REQUIREMENTS
FDA SaMD guidance	Medical device software	Risk classification, GMLP, clinical validation, PCCP, post-market surveillance
QMSR (21 CFR 820, effective Feb 2, 2026)	Medical devices	Quality management system incorporating ISO 13485; design controls; documentation
21 CFR Part 11	Electronic records	Audit trails, access controls, electronic signatures, data integrity
ONC HTI-1 (DSI, 45 CFR 170.315(b)(11))	Certified health IT	Predictive DSI transparency: source attributes, risk management, data governance; Base EHR criterion since Jan 1, 2025
HIPAA	Covered entities	Privacy, security, minimum necessary, BAAs, breach notification
CMS CoP	Medicare providers	QAPI, medical staff, patient rights, information management
Joint Commission	Accredited orgs	Leadership, performance improvement, information management

REGULATION / GUIDANCE	SCOPE	KEY CLINICAL AI REQUIREMENTS
EU AI Act	EU market access / EU use cases	High-risk classification for AI in medical devices and IVDs; conformity assessment; phased application with an extended medical-device transition (currently Aug 2, 2027), layered on MDR and IVDR

For questions on implementation or to discuss evidence-based governance infrastructure for clinical AI: ferz.ai