AI GOVERNANCE FRAMEWORK

For Health Systems

Executive Guidelines · Multi-Service-Line Hospitals

United States Edition | HIPAA · FDA AI/ML · ONC · CMS · NIST AI RMF

Version 1.0 | Starting-from-Scratch Implementation Edition

HOW TO USE THIS DOCUMENT

This framework is ready to adopt as-is for most U.S. hospitals. Fields highlighted in **amber** require your organization to supply specific information — such as your hospital name, dates, designated roles, and contact details. All other content is written for immediate use.

✎ **Amber fields = Organization-specific — must be completed before adoption.**

FIELD	DETAIL
Organization Name	✎ *[Insert Health System Name]*
Framework Title	AI Governance Framework for Health Systems
Version	1.0
Effective Date	✎ *[Insert Effective Date]*
Next Review Date	✎ *[Insert Date + 12 Months]*
Framework Owner	AI Governance Committee / AI Program Office
Approved By	✎ *[CEO / Board Chair Name and Title]*
Regulatory Alignment	HIPAA · FDA AI/ML SaMD · ONC · CMS · NIST AI RMF

SECTION

PREAMBLE

Artificial Intelligence is no longer an emerging technology in healthcare — it is here, it is consequential, and it carries real risk if deployed without structure. This framework provides hospital executives with a comprehensive, implementable AI governance system designed for organizations with no prior formal governance in place.

This document is structured to be adopted incrementally. Health systems do not need to implement everything at once. A phased approach — described in Section 9 — allows governance to grow alongside AI maturity. Fields highlighted in amber require your organization's specific input before this document can be formally adopted.

SECTION 1

SCOPE AND APPLICABILITY

1.1 Who This Governs

This framework applies to all artificial intelligence, machine learning, and algorithmic decision-support tools used within or by the health system, including:

- Clinical Decision Support (CDS) tools
- Ambient documentation and voice AI (e.g., Nuance DAX, Suki, Abridge)
- Diagnostic imaging AI
- Predictive analytics (readmission, deterioration, sepsis)
- Revenue cycle and coding automation
- Scheduling and operational optimization tools
- Patient-facing chatbots and virtual assistants
- Any third-party or vendor-supplied AI embedded in EHR systems (Epic, Oracle Cerner, etc.)

1.2 Who Is Bound by This Framework

All employees, medical staff, vendors, contractors, and affiliated partners who develop, procure, deploy, or use AI systems within the health system are subject to this framework regardless of their service line.

1.3 Service Line Coverage

This framework is explicitly designed for hospitals operating five or more service lines. Each service line is required to maintain a designated AI Steward role (Section 4). The service lines covered by this framework are:

✎ Covered Service Lines

List all service lines in scope (e.g., Acute Care, Emergency Medicine, Surgical Services, Behavioral Health, Oncology, Cardiovascular, Obstetrics, Rehabilitation, Ambulatory).

SECTION 2

GOVERNING PRINCIPLES

All AI deployment decisions must be evaluated against the following seven principles:

1. **Patient Safety Above All**

   No AI system will be permitted to operate if it creates an unacceptable risk of direct patient harm. Safety validation is a non-negotiable prerequisite to deployment.

2. **Human Accountability**

   AI tools augment clinical and operational judgment; they do not replace it. A qualified human must remain accountable for every consequential decision influenced by an AI system.

3. **Transparency**

   Clinicians, patients, and administrators must be able to understand what an AI system does, what data it uses, and what its limitations are. Black-box deployment is not permitted.

4. **Equity**

   AI systems must be evaluated for disparate impact across race, ethnicity, gender, age, disability status, insurance type, and language.

5. **Privacy and Data Stewardship**

   All AI systems accessing PHI must comply fully with HIPAA. PHI may not be used to train third-party AI models without explicit, audited authorization.

6. **Regulatory Compliance**

   The health system will comply with all applicable federal and state regulations including FDA guidance on SaMD, ONC interoperability rules, and CMS conditions of participation.

7. **Continuous Oversight**

   All active AI systems are subject to ongoing monitoring, drift detection, and periodic revalidation. Deployment is a beginning, not a conclusion.

SECTION 3

GOVERNANCE STRUCTURE

3.1 Executive Sponsor

The Executive Sponsor holds ultimate accountability for the organization's AI governance posture. The sponsor approves the governance framework, chairs or delegates the AI Governance Committee, and ensures adequate resourcing.

✎ Executive Sponsor

Name, Title (e.g., Chief Executive Officer or designate: CMO / COO)

✎ Sponsor Email

email@healthsystem.org

✎ Effective From

MM/DD/YYYY

3.2 AI Governance Committee (AIGC)

The AIGC is the primary decision-making body. It meets at minimum quarterly and has authority to convene emergency sessions as needed.

ROLE	NAME / DESIGNEE	STATUS
Chief Medical Officer (Chair or Co-Chair)	✎ *[Name]*	✎ *[Confirmed / TBD]*
Chief Nursing Officer	✎ *[Name]*	✎ *[Confirmed / TBD]*
Chief Information Officer	✎ *[Name]*	✎ *[Confirmed / TBD]*
Chief Privacy Officer	✎ *[Name]*	✎ *[Confirmed / TBD]*
Chief Compliance Officer	✎ *[Name]*	✎ *[Confirmed / TBD]*
Chief Financial Officer	✎ *[Name]*	✎ *[Confirmed / TBD]*
Legal Counsel	✎ *[Name]*	✎ *[Confirmed / TBD]*
VP/Director of Quality and Patient Safety	✎ *[Name]*	✎ *[Confirmed / TBD]*
VP/Director of Human Resources	✎ *[Name]*	✎ *[Confirmed / TBD]*
Service Line AI Stewards (rotating)	✎ *[Name]*	✎ *[Confirmed / TBD]*
Patient/Community Advocate	✎ *[Name]*	✎ *[Confirmed / TBD]*

3.3 AI Program Office (APO)

The APO is the operational hub for governance execution. It is responsible for running the governance lifecycle, maintaining the AI inventory, coordinating reviews, and supporting service lines.

✎ AI Program Director / APO Lead

Name, Title, Department, Contact Email

✎ APO Email / Reporting Channel	✎ APO Established Date
aigovernance@healthsystem.org	*MM/DD/YYYY*

3.4 Clinical AI Safety Officer (CASO)

Required if the health system operates AI tools that directly influence clinical decisions. Typically a physician or senior clinical leader with informatics training.

✎ Clinical AI Safety Officer
Name, Title, Credentials (e.g., MD, CMIO — or note 'Combined with [Role]' if shared function)

SECTION 4

SERVICE LINE AI STEWARDS

Each service line must designate an AI Steward — a clinical or operational leader who serves as the primary accountability point for AI use within their domain. AI Stewards do not need to be technical experts but must hold a recognized leadership role and complete required AI literacy training.

SERVICE LINE	AI STEWARD NAME	TITLE / ROLE	EMAIL
Acute Medical/Surgical	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Emergency Services	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Surgical Services	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Behavioral Health	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Oncology	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Cardiovascular	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Women's & Obstetrics	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
Ambulatory/Outpatient	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
✎ *[Add Service Line]*	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*
✎ *[Add Service Line]*	✎ *[Name]*	✎ *[Title]*	✎ *[Email]*

SECTION 5

AI SYSTEM LIFECYCLE — FROM REQUEST TO RETIREMENT

5.1 Risk Tiering

All AI systems are classified into one of three risk tiers upon intake registration.

TIER 1 — LOW RISK

Definition: Operational or administrative AI with no direct patient care impact and no PHI access, or PHI limited to de-identified/aggregate data.

Examples: Staff scheduling optimization, supply chain forecasting, facility management.

Review Required: APO administrative review only. No AIGC approval required.

Timeline: 2–4 weeks.

TIER 2 — MODERATE RISK

Definition: Systems that influence care delivery decisions, access PHI, generate clinical documentation, or produce output reviewed by clinicians prior to action.

Examples: Ambient documentation AI, patient deterioration alerts, readmission prediction, prior authorization automation.

Review Required: APO review + Clinical AI Safety Officer sign-off + AIGC notification.

Timeline: 4–8 weeks.

TIER 3 — HIGH RISK

Definition: Systems that autonomously generate clinical outputs, directly influence treatment decisions without mandatory human review, or are classified as FDA SaMD Class II or III.

Examples: Diagnostic imaging AI (FDA-cleared), AI-driven treatment recommendation engines, autonomous triage systems.

Review Required: Full AIGC approval, CASO sign-off, Legal review, Privacy review, equity analysis.

Timeline: 8–16 weeks. May require FDA compliance documentation review.

5.2 Lifecycle Stages Overview

1 · Intake & Registration	Submit to APO before any pilot, procurement, or deployment. BAA confirmation required for all PHI-accessing systems.
2 · Risk Tiering	APO assigns Tier 1/2/3 within 10 business days of complete submission.
3 · Due Diligence	Tier 2 & 3: clinical validation, equity analysis, privacy/security review, vendor accountability review. (See Section 6 — Checklist.)
4 · Pilot	Tier 2: min 60 days. Tier 3: min 90 days. Performance metrics set before pilot begins.
5 · Deployment	AIGC approval required. Training delivered before go-live. Monitoring activated.

6 · Ongoing Monitoring	Annual review (all). Quarterly AIGC reporting + 6-month drift assessment (Tier 2 & 3).
7 · Decommissioning	Formal request to APO. Data disposition plan. BAA termination.

SECTION 6

VENDOR & AI SYSTEM DUE DILIGENCE CHECKLIST

Complete this checklist for every Tier 2 and Tier 3 AI system before AIGC approval. One checklist per system. Retain completed checklists in the AI system registry. Fields marked ✎ in amber must be completed by your organization.

✎ Health System Name
[Insert Health System Name]

AI System Name	✎ *[Enter system name]*	**Vendor Name**	✎ *[Enter vendor]*
Version / Release	✎ *[Version number]*	**Review Date**	✎ *[MM/DD/YYYY]*
Risk Tier	✎ *[Tier 2 / Tier 3]*	**Service Line**	✎ *[Service line name]*
Intended Use Case	✎ *[Describe the intended clinical or operational use case]*	**AI Steward**	✎ *[Name, Title]*
Requestor	✎ *[Name, Title, Department]*	**APO Reviewer**	✎ *[Name]*

PART A — REGULATORY & FDA STATUS

■ Has the FDA regulatory status of this system been determined?
Classification: ■ Not a Medical Device ■ SaMD Class I ■ SaMD Class II ■ SaMD Class III
■ If SaMD Class II or III: Has FDA 510(k) clearance, De Novo authorization, or PMA approval been confirmed?
■ Is the FDA clearance letter or authorization documentation on file?
■ Does the system fall under the FDA's Clinical Decision Support software exclusion criteria (21st Century Cures Act)?
■ If subject to FDA oversight: Has the vendor's Predetermined Change Control Plan (PCCP) been reviewed?
■ Are there any outstanding FDA safety communications, recalls, or adverse event reports for this system?
■ Does the system comply with ONC Health IT certification requirements if it interfaces with certified EHR technology?

Notes / FDA Documentation Reference:

✎ FDA Status Notes
Record FDA classification decision rationale, clearance numbers, or document references here.

PART B — HIPAA & PRIVACY COMPLIANCE

■ Does this system access, process, store, or transmit Protected Health Information (PHI)?
■ Is a signed Business Associate Agreement (BAA) in place with the vendor?
■ Has the BAA been reviewed by Legal Counsel within the past 12 months?
■ Does the system comply with the HIPAA Minimum Necessary standard for PHI access?
■ Is PHI transmitted outside the organization's network or to third parties?
If yes, document the legal basis: ■ BAA ■ Patient Authorization ■ De-identification ■ Other
■ Has a HIPAA Security Risk Analysis been performed covering this system?
■ Are appropriate access controls, audit logging, and encryption in place for PHI handled by this system?
■ Does the vendor's data retention and deletion policy comply with organizational and HIPAA requirements?
■ Is PHI used to train or re-train the AI model? If yes, has explicit authorization been obtained?
■ Has a Privacy Impact Assessment (PIA) been conducted for this system?

■ Is the data flow diagram for this system documented and current?

BAA Reference & Privacy Notes:

✎ BAA Execution Date & Reference Number
BAA Execution Date: [MM/DD/YYYY] | BAA Reference / Contract #: [Insert] | Reviewed by: [Name]

PART C — CLINICAL VALIDATION

- ■ Has the clinical evidence base for this system been reviewed? (peer-reviewed publications, FDA clearance studies, white papers)
- ■ Was the system validated on a patient population similar to ours in demographics, payer mix, and acuity?
- ■ Are sensitivity and specificity rates documented and clinically acceptable for the intended use?
- ■ Is the false positive rate acceptable given the clinical workflow and downstream burden on staff?
- ■ Is the false negative rate acceptable given the potential for missed diagnoses or delayed interventions?
- ■ Has the Clinical AI Safety Officer reviewed and approved the clinical validation evidence?
- ■ Is there a documented clinical champion or physician sponsor for this system within the relevant service line?
- ■ Has the system been evaluated in a comparable health system (reference site visits or case studies available)?
- ■ Are there published peer-reviewed studies supporting this system's performance? (attach references)
- ■ Is the system's intended use case within the scope of its validation evidence — i.e., no off-label clinical use?
- ■ Does the system provide explainable outputs that clinicians can interrogate and understand?
- ■ Is there a defined clinical override protocol — a clear process for clinicians to override or reject AI recommendations?

Key Performance Metrics (enter values from vendor validation documentation):

METRIC	VENDOR-REPORTED VALUE	ACCEPTABLE THRESHOLD	PASS / FAIL
Sensitivity (Recall)	✎	✎ *[e.g., ≥85%]*	✎
Specificity	✎	✎ *[e.g., ≥80%]*	✎
AUC / AUROC	✎	✎ *[e.g., ≥0.80]*	✎
Positive Predictive Value	✎	✎	✎
Negative Predictive Value	✎	✎	✎
False Positive Rate	✎	✎ *[e.g., ≤15%]*	✎
Validation Dataset Size	✎ *[n=]*	—	—
Validation Population	✎ *[Describe]*	—	—

PART D — EQUITY & BIAS ANALYSIS

- ■ Does the vendor provide disaggregated performance data by race, ethnicity, sex, and age?
- ■ Does the vendor provide disaggregated performance data by insurance status / payer type?
- ■ Does the vendor provide disaggregated performance data by preferred language or LEP status?
- ■ Has the system been specifically evaluated for algorithmic bias or disparate impact?
- ■ If performance disparities exist across subgroups: is a documented mitigation plan in place?
- ■ Was the training dataset representative of diverse patient populations?
- ■ Has an internal equity review been conducted using data from our patient population?
- ■ Does the system comply with Section 1557 of the ACA (non-discrimination in health programs)?
- ■ Is there a process to monitor for emerging equity concerns post-deployment?
- ■ Has the equity analysis been reviewed and signed off by the Clinical AI Safety Officer?

Equity Analysis Summary:

✎ Identified Equity Concerns and Mitigations

Describe any identified performance disparities by subgroup and the planned or confirmed mitigation strategy. Note: 'No disparities identified' is an acceptable entry if vendor data supports this conclusion.

PART E — CYBERSECURITY & TECHNICAL INTEGRATION

- ■ Has the Information Security team reviewed and approved this system?
- ■ Has a formal security risk assessment (e.g., HIPAA Security Rule analysis) been completed for this system?
- ■ Is all data in transit encrypted using TLS 1.2 or higher?
- ■ Is all data at rest encrypted to NIST-recommended standards?
- ■ Are role-based access controls implemented and documented?
- ■ Is audit logging enabled for all PHI access by the system?
- ■ Does the vendor have SOC 2 Type II certification or equivalent?
- ■ Has the vendor's penetration testing / vulnerability assessment been reviewed?
- ■ Is the system integrated with or accessing the EHR? If yes: have EHR access controls been scoped and approved?
- ■ Is there a defined incident response procedure for a security breach involving this system?
- ■ Is the vendor's SLA for uptime, patch management, and vulnerability response acceptable?
- ■ Is there a documented data backup and recovery plan?

✎ Security Review Approver & Date

Approver Name: [Name, Title] | Date: [MM/DD/YYYY] | Assessment Reference #: [Insert]

PART F — VENDOR ACCOUNTABILITY & CONTRACT TERMS

- ■ Does the contract include explicit performance SLAs (accuracy, uptime, response time)?
- ■ Does the contract require the vendor to notify the organization of any model updates, retraining, or version changes?
- ■ Does the contract include provisions for model drift monitoring and disclosure?
- ■ Does the contract specify who is responsible for post-market surveillance and ongoing safety monitoring?
- ■ Does the contract include audit rights — allowing the organization to audit vendor AI practices?
- ■ Does the contract specify data ownership — confirming organizational ownership of patient data?
- ■ Does the contract prohibit the vendor from using organizational PHI to train models for other clients?
- ■ Are indemnification and liability provisions adequate for AI-related patient harm scenarios?
- ■ Does the contract include a termination for cause clause tied to safety or compliance failures?
- ■ Is there a documented vendor escalation pathway for clinical safety concerns?
- ■ Has the vendor provided customer references from comparable health systems?
- ■ Is the vendor financially stable and likely to maintain support for the contract term?

Contract Reference Information:

✎ Contract / MSA Reference

Contract Name: [MSA / Agreement Title] | Contract #: [Insert] | Execution Date: [MM/DD/YYYY] | Term End: [MM/DD/YYYY]

✎ Contract Reviewed By (Legal Counsel)

Name: [Attorney Name] | Date: [MM/DD/YYYY] | Firm / In-House: [Insert]

PART G — PATIENT RIGHTS & DISCLOSURE

- ■ Has the Notice of Privacy Practices been updated to reference this system where applicable?
- ■ Does the informed consent process disclose AI involvement where this system influences diagnosis or treatment planning?
- ■ Is there a mechanism for patients to request human review of AI-generated recommendations?
- ■ If this is a patient-facing AI system: does it clearly identify itself as AI to the patient?
- ■ Does the patient-facing system provide a clear escalation pathway to human staff?
- ■ Has compliance with Section 1557 (non-discrimination) been confirmed for patient-facing functions?

PART H — POST-DEPLOYMENT MONITORING PLAN

Complete before deployment approval. This becomes the active monitoring commitment.

MONITORING REQUIREMENT	FREQUENCY	RESPONSIBLE PARTY	REPORTING TO
Performance vs. baseline metrics	✎ *[e.g., Quarterly]*	✎ *[Name / Role]*	✎ *[AIGC / APO]*
Model drift assessment	✎ *[e.g., Every 6 months]*	✎ *[Name / Role]*	✎ *[AIGC / APO]*
Equity re-analysis	✎ *[e.g., Annually]*	✎ *[Name / Role]*	✎ *[AIGC / APO]*
Incident/concern review	✎ *[e.g., Monthly]*	✎ *[Name / Role]*	✎ *[AIGC / APO]*
Vendor update review	✎ *[e.g., Per release]*	✎ *[Name / Role]*	✎ *[AIGC / APO]*
Full governance re-review	✎ *[e.g., Annual]*	✎ *[Name / Role]*	✎ *[AIGC / APO]*

CHECKLIST SIGN-OFF

All four signatures below are required before this checklist is submitted to the AIGC for approval.

ROLE	NAME & SIGNATURE	DATE	DECISION
Service Line AI Steward	✎	✎	■ Approve ■ Conditional ■ Reject
AI Program Office Reviewer	✎	✎	■ Approve ■ Conditional ■ Reject
Clinical AI Safety Officer	✎	✎	■ Approve ■ Conditional ■ Reject
AIGC Chair / Co-Chair	✎	✎	■ Approve ■ Conditional ■ Reject

✎ **Conditions / Notes (if conditional approval)**

Describe any conditions that must be resolved before full deployment approval is granted.

SECTION 7

WORKFORCE TRAINING & AI LITERACY

TIER	AUDIENCE	CONTENT	FORMAT	FREQUENCY
A	All Staff	What AI is, how it is used, how to report concerns, patient rights	30-min eLearning	Annual
B	Clinical Staff & Service Line Leaders	Tier A + critical evaluation of AI output, override protocols, service-line-specific tools	60-min hybrid module	Annual
C	AI Stewards & AIGC Members	Tiers A & B + governance framework, risk tiering, vendor evaluation, regulatory landscape	4-hr onboarding + quarterly briefings	Annual + Ongoing
D	AI Program Office Staff	Full technical and regulatory competency: FDA PCCP, ONC, NIST AI RMF, PIA methodology	Individual development plan	Annual review

✎ **Training Platform / LMS**

Name of LMS or training platform where modules will be housed (e.g., HealthStream, Workday Learning, custom LMS)

SECTION 8

INCIDENT REPORTING & RESPONSE

8.1 What Must Be Reported

- An AI system output that contributes to a patient safety event or near-miss
- Unexpected, erroneous, or biased AI behavior inconsistent with documented purpose
- Unauthorized access to or disclosure of PHI by or through an AI system
- A vendor changes an AI model without disclosure or prior approval
- Staff feel pressured to follow an AI recommendation they believe is clinically inappropriate
- Any AI-related HIPAA breach or suspected breach

✎ AI Incident Reporting Channel

Reporting email / portal URL / safety system module name and URL (e.g., aigovernance@healthsystem.org or [SafetyReportingSystem] module link)

8.2 Response Timelines

TRIGGER	RESPONSE REQUIREMENT	TIMELINE
Any incident report received	APO acknowledges receipt	2 business days
Any incident report received	Initial triage completed	5 business days
Tier 3 system or patient harm event	AIGC and CASO notified immediately; emergency convening if warranted	Immediate
Any event contributing to patient harm	Formal root cause analysis (RCA)	Per RCA policy
Confirmed HIPAA breach	HIPAA Breach Notification Rule process activates	Per HIPAA (60 days)

The health system strictly prohibits retaliation against any staff member who in good faith reports a concern about an AI system.

SECTION 9

IMPLEMENTATION ROADMAP (PHASED)

Phase 1 — Foundation

Months 1–3

- ■ Executive Sponsor designated; framework adopted at Board or C-Suite level
- ■ AIGC constituted with initial membership
- ■ APO function established (fractional acceptable in Year 1)
- ■ Comprehensive inventory of all AI tools currently in use — this is often a surprising discovery exercise
- ■ Interim risk tier assigned to all existing tools
- ■ AI incident reporting channel activated (email address is sufficient initially)
- ■ BAA status confirmed for all vendors with PHI access
- ■ **MILESTONE: Completed AI system inventory submitted to AIGC**

Phase 2 — Structure

Months 4–6

- ■ Service Line AI Stewards designated and onboarded
- ■ Intake and risk tiering process formally launched for new AI requests
- ■ Tier A staff training module developed and deployed
- ■ Vendor review templates and due diligence checklist activated
- ■ Governance policies reviewed by Legal and Compliance
- ■ Board briefing on AI governance posture
- ■ **MILESTONE: First AIGC quarterly meeting under full framework**

Phase 3 — Operations

Months 7–12

- ■ All existing Tier 3 systems formally reviewed and approved or suspended
- ■ Tier B and Tier C training deployed
- ■ Ongoing monitoring cadence established for all active systems
- ■ Patient disclosure language updated in Notice of Privacy Practices
- ■ Equity analysis initiated for highest-volume clinical AI tools
- ■ AI governance scorecard developed and reported to Board annually
- ■ **MILESTONE: First annual AI risk report to Board completed**

Phase 4 — Maturity

Year 2+

- Proactive engagement with emerging regulatory developments
- Participation in national AI governance collaborative or benchmarking network
- Internal AI development standards if organization builds proprietary tools
- Advanced equity analytics capability
- Formal NIST AI Risk Management Framework alignment assessment

SECTION 10

REGULATORY REFERENCE SUMMARY

HIPAA Privacy Rule (45 CFR Part 164)	Governs use and disclosure of PHI. All AI systems accessing PHI must operate under Business Associate Agreements. Minimum necessary standard applies to AI data pipelines.
HIPAA Security Rule	Requires administrative, physical, and technical safeguards for ePHI. AI systems accessing ePHI must be covered by the organization's security risk analysis.
FDA AI/ML-Based SaMD Guidance	AI systems intended to diagnose, treat, or prevent disease are subject to FDA oversight. Confirm FDA clearance for all clinical AI tools. Reference: 2021 AI/ML Action Plan and Predetermined Change Control Plan guidance.
ONC 21st Century Cures Act / Information Blocking Rule	AI systems involved in clinical documentation, data access, or interoperability must comply with ONC certification requirements and must not constitute information blocking.
Section 1557 of the ACA	Prohibits discrimination in federally funded health programs. AI producing disparate outcomes for protected populations may constitute discriminatory practice.
CMS Conditions of Participation	AI tools affecting quality of care, patient rights, or medical record integrity may be subject to CMS CoP requirements.
NIST AI Risk Management Framework (AI RMF 1.0)	Voluntary framework aligned with this document's four functions: Govern, Map, Measure, Manage.
FTC Act Section 5	FTC has signaled enforcement interest in deceptive AI practices. Patient-facing AI claims must be accurate and substantiated.

SECTION 11

KEY DEFINITIONS

TERM	DEFINITION
Algorithmic Bias	Systematic error in AI output creating unfair outcomes for identifiable groups, often resulting from biased training data or flawed model design.
Ambient Documentation AI	AI tools that listen to patient-clinician conversations and generate clinical notes automatically (e.g., Nuance DAX, Suki, Abridge).
Business Associate Agreement (BAA)	A HIPAA-required contract between a covered entity and a vendor that creates, receives, maintains, or transmits PHI on its behalf.
Clinical Decision Support (CDS)	Software designed to assist clinicians in patient care decisions, ranging from simple alerts to complex predictive models.
Model Drift	Degradation of an AI model's performance over time as real-world data diverges from its training data.
Protected Health Information (PHI)	Individually identifiable health information governed by HIPAA.
Software as a Medical Device (SaMD)	Software intended for medical purposes without being part of a hardware medical device, as defined by the FDA and IMDRF.
Predetermined Change Control Plan (PCCP)	An FDA mechanism allowing AI/ML developers to pre-specify planned modifications and the controls used to manage those changes.

SECTION 12

FRAMEWORK MAINTENANCE

This framework is a living document. It must be reviewed and updated annually at minimum, within 90 days of any significant regulatory change, following any AI-related patient safety event, and whenever the AI portfolio materially expands or changes.

✎ **Next Scheduled Review Date**

MM/DD/YYYY — Review owner: [Name, Title]

✎ **Version History**

v1.0 — [Effective Date] — Initial adoption v1.x — [Date] — [Brief description of change]

This framework is intended to be adapted to the specific operational, legal, and clinical context of the adopting health system. Legal counsel and compliance review is recommended prior to formal adoption.