



AI Primer for the Board of a Hospital or Health System

Including Owned Physician Practices and Ambulatory Services

Understanding AI Risk, Governance, and Board Oversight in a Rapidly Evolving Clinical Environment





Table of Contents

Why AI Matters to Hospital Boards 3

AI as a Strategic Opportunity for Risk Mitigation 3

The Different Types of AI the Board May Encounter 4

What “AI” Means in the Hospital and Owned Practice Context..... 4

Where AI Is Already Entering the Organization..... 5

Hospital and Practice-Specific AI Risks 5

What “Good” AI Governance Looks Like in a Hospital or Health System 6

Board vs. Management Responsibilities 6

Questions Board Members Should Ask 6

Expected Next Steps 7

APPENDIX

AI FAQ for Hospital and Health System Boards 8

1. Why AI Matters to Hospital Boards

Artificial intelligence (AI) is no longer a technology experiment in healthcare. It is already embedded in clinical decision support, documentation workflows, diagnostic tools, revenue cycle operations, and the software platforms your vendors provide every day. Boards are increasingly expected to understand how AI is used across the enterprise and whether appropriate oversight exists.

For community hospitals, health systems, and their owned practices, AI risk most often presents as:

- **Patient safety and care quality risk**
- **HIPAA and patient privacy exposure**
- **Clinician decision-making and escalation failures**
- **Third-party vendor and EHR platform risk**
- **Regulatory, accreditation, and liability exposure**
- **Reputational and trust impact with patients, medical staff, and the community**

BOARD FRAMING

AI governance is now a patient safety, compliance, and fiduciary issue—not a technology committee agenda item.

2. AI as a Strategic Opportunity for Risk Mitigation

From a board perspective, AI should be viewed not only as a source of new risk, but also as a potential tool for addressing the operational and clinical challenges hospitals already face.

Mitigating Workforce Shortages and Clinician Burnout

- AI can reduce documentation burden, automate prior authorizations, and streamline administrative workflows—allowing clinicians and nurses to spend more time on direct patient care.
- AI-assisted scheduling and staffing tools can improve workforce resilience without replacing clinical judgment or reducing staffing below safe levels.

BOARD ROLE

Ensure AI is used to support clinical capacity and care quality—not as a substitute for adequate staffing or physician judgment.

Reducing Clinical and Operational Error

- AI can assist with standardizing order sets, flagging incomplete or inconsistent documentation, and prompting escalation when clinical patterns suggest deterioration.
- Properly governed AI can help reduce preventable harm events while preserving human oversight for complex, high-acuity clinical decisions.

BOARD ROLE

Ask how AI is being used to improve safety and reliability without automating decisions that require physician or clinical judgment.

Strengthening Financial and Operational Resilience

- AI can improve revenue cycle performance, reduce denials, and enhance coding accuracy—addressing margin pressures that are among the most urgent challenges facing community hospitals.
- AI can support forecasting for volumes, staffing, and supply chain, enabling leadership to anticipate challenges and respond proactively.

BOARD FRAMING

AI should function as a resilience and efficiency tool, paired with governance that prevents over automation or misplaced clinical trust.

3. The Different Types of AI the Board May Encounter

Generative AI (Creates Content)

Examples: Drafting clinical notes or discharge summaries; writing patient communications, policies, or prior authorization letters; ambient documentation tools that transcribe patient encounters.

Primary risks: Clinically inaccurate content; unintentional inclusion of protected health information (PHI); over-reliance on AI-generated summaries in lieu of direct patient assessment.

Predictive / Analytical AI (Forecasts or Scores)

Examples: Sepsis early warning tools; readmission risk scoring; deterioration alerts; financial or operational performance forecasting.

Primary risks: Bias in underlying training data; alert fatigue leading clinicians to ignore valid warnings; false confidence in predictions; poor escalation when results are ambiguous.

Assistive / Decision-Support AI (Supports Clinical Judgment)

Examples: Diagnostic imaging AI for radiology or pathology; medication interaction alerts; care gap identification in owned practices.

Primary risks: Clinicians treating AI recommendations as definitive; missed edge cases involving high-acuity patients; liability questions when AI-assisted decisions lead to adverse outcomes.

Automation / Agentic AI (Acts Without Human Initiation)

AI that executes tasks autonomously, not just recommends.

Examples: Automated prior authorization submission; autonomous scheduling or referral routing; AI-initiated patient outreach.

Primary risks: Reduced human oversight in clinical pathways; difficulty explaining or reversing automated actions; elevated liability when automation intersects with clinical decisions.

BOARD TAKEAWAY

As AI moves from assistive to automated, the organization's oversight obligations—and its liability exposure—increase substantially.

4. What “AI” Means in the Hospital and Owned Practice Context

AI in community hospitals and owned practices is typically embedded in existing platforms, often without explicit labeling as ‘AI.’ The board should assume AI is already present in many of the tools staff and clinicians use daily.

Common examples include:

- Ambient documentation and clinical note generation (e.g., Nuance DAX, Abridge, Suki)
- EHR-embedded clinical decision support and early warning alerts (e.g., Epic, Oracle Health)
- Diagnostic imaging AI for radiology, pathology, and cardiology
- Revenue cycle tools: coding assistance, denial prediction, prior authorization automation
- Patient engagement and outreach automation
- Workforce scheduling and staffing optimization
- Productivity tools such as Microsoft Copilot used by administrative and clinical staff

Critical Distinction for the Board

- Some AI advises
- Some AI prioritizes or flags
- Some AI automates

Risk escalates quickly as AI moves closer to direct patient care or autonomous clinical action.

5. Where AI Is Already Entering the Organization

Even when not formally reviewed or approved, AI commonly enters hospitals and owned practices through:

- Embedded AI features in EHR and clinical platforms (often enabled by default in vendor updates)
- Revenue cycle and billing software with AI-assisted coding or denial management
- Diagnostic imaging and specialty AI tools procured at the department level
- Workforce productivity software adopted by administrative teams
- Owned practice staff using consumer AI tools for documentation or communications
- Individual clinician or staff experimentation (“shadow AI”—unofficial AI use not formally approved by the organization)

KEY INSIGHT

Most hospitals cannot accurately answer “Where is AI used today—including in our owned practices?” without a structured inventory. The answer is almost certainly: more places than leadership realizes.

6. Hospital and Practice-Specific AI Risks

Patient Safety and Clinical Quality

- Missed or delayed escalation of patient deterioration when AI alert systems malfunction or are ignored
- Over-reliance on AI-generated clinical summaries or diagnostic suggestions
- Algorithmic bias producing inequitable recommendations across patient populations
- AI-assisted documentation errors that propagate through the clinical record

Privacy, HIPAA, and Data Governance Risk

- Patient PHI entered into AI tools without proper data agreements
- Unclear vendor practices for AI model training using patient data
- De-identification failures in data shared with AI vendors or research partners

Regulatory, Accreditation, and Liability Exposure

- FDA oversight requirements for AI-enabled clinical decision support classified as medical devices
- CMS and payer scrutiny of AI-assisted coding and billing practices
- Medical staff credentialing and peer review implications when AI influences clinical decisions
- Increased plaintiffs’ attorney focus on AI-influenced adverse outcomes

Workforce and Operational Risk

- Clinician distrust or over-trust of AI tools, both of which introduce safety risk
- Inconsistent AI use across departments, campuses, or owned practices
- Automation displacing clinical judgment in time-sensitive or high-acuity scenarios

Reputational and Community Trust Risk

- Patient and community loss of trust following an AI-related care incident
- Inability to explain AI-influenced clinical decisions to patients, families, or medical staff
- Heightened scrutiny during accreditation surveys, payer audits, or transactions

BOARD TAKEAWAY

These risks are not theoretical—they are already occurring in health systems nationally. The board’s role is to ensure leadership has a governance plan before an incident, not after.

Most hospitals cannot accurately answer “Where is AI used today
—including in our owned practices?”

7. What “Good” AI Governance Looks Like in a Hospital or Health System

Boards are not expected to evaluate or approve individual clinical AI tools. They are expected to ensure that a defensible system of oversight and control exists across the enterprise, including owned practices.

A mature AI governance program in a community hospital or health system typically includes:

- Clear executive and clinical leadership ownership (e.g., CMO, CNO, CMIO, CIO accountable by domain)
- A documented inventory of AI tools in use across the hospital and owned practices
- A risk-based intake, clinical review, and approval process for new AI tools—including EHR vendor updates
- Policies defining acceptable and prohibited AI use for clinical and administrative staff
- Human-in-the-loop expectations (a clinician must review and approve before AI recommendations affect patient care) for all patient-impacting decisions
- Medical staff engagement and education on AI tools in clinical use
- Periodic governance committee and board-level reporting on AI use, incidents, and maturity

The goal is **safe clinical enablement**—not restricting beneficial technology.

8. Board vs. Management Responsibilities

Board Responsibilities

- Set expectations for AI risk tolerance, particularly for patient-facing and clinical applications
- Ensure governance, accountability, and clinical oversight exist across the enterprise
- Require visibility into AI use, incidents, and governance maturity
- Receive periodic AI risk and program updates, including owned practice scope

Management and Clinical Leadership Responsibilities

- Implement governance processes with defined clinical and operational accountability
- Approve and monitor AI use cases, including those embedded in vendor platforms
- Train clinical and administrative staff; address unsafe or unauthorized AI use
- Maintain audit-defensible documentation of AI review and approval decisions

Clear separation protects patients, preserves medical staff trust, and fulfills the board's fiduciary obligations.

9. Questions Board Members Should Ask

- Where is AI currently used across our hospital, departments, and owned practices—including tools embedded in vendor platforms?
- Who owns decisions when AI influences a clinical recommendation, a billing submission, or a care pathway?
- How do we prevent unsafe, unauthorized, or clinically inappropriate AI use?
- How is patient privacy protected when AI tools interact with PHI?
- How would we explain our AI controls to CMS, accreditors, plaintiffs' counsel, or a potential acquirer?
- Do our vendor contracts address how AI is used with patient data—including whether vendors may use that data to train their models?
- If an AI-related adverse event or privacy breach occurred today, do we have a defined response process and clear accountability?
- Are high-acuity clinical settings and owned practices held to appropriately differentiated AI governance standards?

The goal is **safe clinical enablement**—not restricting beneficial technology.


10. Expected Next Steps

Common board-appropriate next steps include:

- Requesting an AI readiness or risk snapshot from management, including owned practice scope
- Directing leadership to formalize AI governance with defined clinical and operational accountability
- Establishing expectations for periodic AI oversight reporting to the quality, compliance, or full board
- Directing legal and compliance counsel to review vendor contract language for AI data use, training rights, and liability provisions
- Engaging the CMO, CMIO, and CNO to ensure clinical AI governance is integrated with existing patient safety and quality structures

AI governance is not a one-time project. As EHR vendors and clinical technology platforms add new AI capabilities—often through routine updates—the board’s role is to ensure that oversight keeps pace. That means setting clear expectations now, receiving periodic updates from management, and revisiting governance standards as the organization’s AI footprint grows across both the hospital and its owned practices. The goal is not to restrict beneficial innovation—it is to ensure that innovation is clinically sound, equitable, and accountable.

AI governance is not a one-time project. As EHR vendors and clinical technology platforms add new AI capabilities—often through routine updates—the board’s role is to ensure that oversight keeps pace.



AI FAQ for Hospital and Health System Boards

Clinical Settings and Risk Differentiation

1. Is AI risk the same across the inpatient hospital and owned practices?

No. The risk profile differs materially by setting.

Inpatient / High-Acuity Settings

- Patients are often unable to self-advocate during acute illness
- AI-influenced decisions can affect diagnosis, treatment, and escalation with life-safety implications
- Errors have higher potential for immediate patient harm
- Regulators, accreditors, and medical staff expect conservative, well-documented safeguards

Owned Practices / Ambulatory Settings

- Patients have greater agency and typically more time for informed decision-making
- AI is used more for documentation, coding, scheduling, and care gap identification
- Risk is generally lower acuity but still present, particularly for diagnostic AI and prior authorization tools

BOARD TAKEAWAY

AI governance must be risk-based and setting-specific, not a single policy applied uniformly across the enterprise.

2. Should inpatient clinical AI have stricter governance than ambulatory AI?

Yes—intentionally, and as a matter of patient safety. Best practice applies graduated controls:

- Stricter clinical review and approval for AI affecting inpatient or high-acuity populations
- Clear human-in-the-loop requirements for all patient-impacting clinical decisions
- More conservative thresholds for AI tools used in the ICU, ED, or surgical settings

This enables innovation in lower-risk settings while protecting vulnerable patient populations.

3. Can AI replace physician or clinical judgment?

No. Governance should explicitly prohibit this in all clinical contexts.

AI may support awareness, documentation, and pattern recognition, but physician and nursing judgment, assessment, and escalation authority must remain primary. This is not only a governance principle—it reflects current FDA guidance on clinical decision support and establishes a defensible liability posture.

4. What AI use cases are most sensitive in the inpatient and ED settings?

- Sepsis, deterioration, and early warning systems
- Diagnostic imaging AI (radiology, pathology, cardiology)
- Medication dosing support or interaction alerts
- Clinical documentation tools that summarize or auto-populate patient records

These require clear clinical ownership, defined escalation protocols, medical staff education, and documented governance approval.

AI may support awareness, documentation, and pattern recognition, but physician and nursing judgment, assessment, and escalation authority must remain primary.

5. Are productivity tools lower risk in owned practices?

Lower risk does not mean no risk. Common concerns in the ambulatory and practice setting include:

- Practice staff entering patient information into AI tools without data agreements
- Clinicians over-trusting AI-generated prior authorization language or clinical summaries
- Inconsistent understanding of acceptable AI use across the practice network

Governance should address acceptable use, data handling, and practice-level training.

6. Do patients need to be informed about AI use in their care?

Increasingly, yes—particularly for AI that influences clinical decisions, diagnostic interpretation, or care planning.

Transparency with patients and families preserves trust and is increasingly expected by regulators, accreditors, and the public. Many health systems are proactively disclosing AI use in clinical settings as a best practice.

7. Does HIPAA risk differ by setting?

The law is the same; the exposure differs.

- **Inpatient settings** present higher sensitivity due to the volume and acuity of PHI involved, the complexity of data-sharing with AI vendors, and documentation practices.
- **Owned practice** settings risk more commonly stems from productivity tools, prior authorization workflows, and patient communication platforms where PHI boundaries are less clearly enforced.

8. What will regulators, accreditors, or attorneys ask after an AI-related incident?

Typical questions include:

- Did the organization know where AI was in use, including in vendor platforms?
- Were clinical safeguards risk-based, documented, and approved?
- Were physicians, nurses, and staff trained on AI limitations and escalation expectations?
- Who owned the decision to deploy AI in that clinical context?
- Was there a human-in-the-loop requirement, and was it followed?

9. What does a reasonable board posture look like?

A defensible posture includes:

- Awareness that AI use is already occurring across the hospital and owned practices
- Risk-differentiated governance with stronger safeguards for high-acuity clinical settings
- Clear delegation to clinical and operational leadership, with defined accountability
- Periodic oversight reporting to the appropriate board committee

Boards do not evaluate individual clinical tools—but they must ensure that governance exists, is documented, and is functioning across the full enterprise.