



AI Governance in Healthcare: A 5-Step Framework for Leaders



TABLE OF CONTENTS

3	AI Alone Won't Transform Care. Governance Will
4	Radiology: A Leading Example of AI Adoption
5	The Rad AI Governance 5-Step Framework Overview
6	Step 1: Assess Evaluating Needs, Readiness, and Risks
9	Step 2: Define Strategy, Governance Structure and Policies
13	Step 3: Select Choosing AI Use Cases and Solutions
19	Step 4: Execute Implementing and Integrating AI Solutions
22	Step 5: Monitor Ongoing Oversight, Performance Management, and Improvement
25	Making AI Work for Healthcare & Key Takeaways

AI Alone Won't Transform Care. Governance Will.

AI will not transform healthcare on technology alone. The real work is governance. Without clear strategy, accountable roles, and disciplined oversight, even the best models create cost, confusion, and risk. With them, AI can reliably improve care, reduce friction, and earn clinician trust.

This eBook is a practical playbook for executives who need more than pilots. It outlines a 5-step governance lifecycle for AI in healthcare, Assess, Define, Select, Execute, Monitor, so leaders can align vision to outcomes, not headlines. The focus is on people, process, and policy first, then technology.

What you will get:

- Strategic guidance for the healthcare leaders and practical tools for implementation teams
- Checklists, role charters, decision criteria, and vendor due-diligence questions
- Templates for project intake, risk reviews, success metrics, and ongoing monitoring

Radiology is offered as the exemplar because it is one of the more mature clinical adopters of AI and offers repeatable lessons for broader healthcare. The principles and tools apply system-wide, from diagnostics to operations. Follow this framework to turn AI from an experiment into a governed capability that improves patient outcomes, strengthens operations, and meets regulatory expectations.

Radiology: A Leading Example of AI Adoption


Radiology has been at the forefront of healthcare AI, offering a glimpse into the future for other specialties. In fact, more than three-quarters of all FDA-cleared AI medical software is designed for radiology, far more than any other specialty¹. As Dr. Curtis Langlotz of Stanford notes, *“Radiology is leading the way in the development and implementation of AI in clinical practice.”* The field’s embrace of AI is evident: over two-thirds of radiology departments in the U.S. are already using AI, a figure that has roughly doubled since 2019². Radiology’s head start stems from clear use cases (e.g. image analysis, diagnostic support, generative AI) and pressing needs like workforce shortages.

This eBook will highlight radiology examples, such as early adoption of generative AI into their workflow, to illustrate each step. However, the governance principles and tools discussed are applicable across all healthcare sectors (from pathology to hospital operations). Radiology’s experience shows both the potential of AI and the critical importance of governance to ensure these tools are effective, trusted, and aligned with clinical needs.

The Rad AI Governance 5-Step Framework Overview

AI adoption can't be left to chance. To move from pilot experiments to enterprise impact, healthcare organizations need a governance model that balances vision with execution, strategy with operations, and innovation with accountability.

The Rad AI 5-Step Framework provides a practical lifecycle for governing AI in healthcare:



Each step builds on the previous, creating a lifecycle. Governance is a continuous cycle rather than a one-and-done project – after Step 5, organizations loop back to reassess new needs or iterate on existing solutions. In the following sections, we break down each step with practical sub-steps, examples and tools.

STEP 1: ASSESS

Evaluating Needs, Readiness and Risks

Every journey starts with knowing where you stand. In AI, skipping this step is like building a hospital wing without surveying the foundation. Leaders must first assess both the current state of AI in the organization and the environment in which it will operate.

The goal of this phase is simple but critical: **understand where you are, what you have, and what you need before committing to AI projects.** This foundation allows leaders to make informed decisions and avoid costly missteps later in the process. From here, the assessment breaks down into four key activities:



Inventory Current AI Use and Data Assets: The first step is understanding what is already in place. Most large health systems are already using AI or advanced analytics in some form, such as sepsis detection, medical imaging analysis, readmission prediction, or coding. The goal of this review is not to halt existing efforts but to make them visible, identify what is working, uncover redundancies, and ensure compliance with regulations.

From there, build a centralized inventory of all AI applications or algorithms across the enterprise, including small pilots that may be operating in silos. This provides a single source of truth and highlights both overlap and opportunities for scale. At the same time, take stock of the data assets that could fuel future AI initiatives, including clinical datasets, imaging archives, and EHR data. Assessing the quality and accessibility of these datasets is essential for determining which AI projects are feasible and where further investment is needed.



Identify Priority Use Cases and Pain Points: Engage stakeholders to find the clinical or operational challenges that AI might help solve. Where are the bottlenecks or gaps in care quality? In radiology, for example, high reading volumes resulting in clogged worklists and delayed interpretations are pain points that AI-assisted detection tools aim to address. List potential AI use cases (e.g. imaging triage, predictive analytics for ICU deterioration, automating administrative tasks) and estimate their value to the organization. Ensure each potential use case aligns with organizational priorities (such as the Quintuple Aim of improving outcomes, satisfaction, equity, and reducing cost). This will guide where to focus first.



Assess Organizational Readiness (People, Process, Technology): AI success depends less on algorithms and more on whether the organization is prepared to adopt them. Without executive sponsorship, clear processes, reliable data, and regulatory alignment, even the best tools will fail. This checklist helps leaders quickly evaluate readiness across four critical areas: **People, Process, Technology and Compliance.**

Assess Organizational Readiness: Checklist



People

- Executive sponsor identified and actively supporting AI initiatives
- Clinical champion(s) engaged to advocate for adoption
- Access to AI or data science expertise (internal team or external partners)
- Stakeholder buy-in confirmed across clinicians, IT, and compliance
- Clinicians' AI literacy assessed and training needs identified
- Cultural readiness evaluated (openness to digital tools, history with prior tech adoption)



Process

- Intake process for new AI/innovation projects defined
- Clear decision-making structure for approving or denying AI projects
- Strategic alignment documented: AI projects linked to goals such as patient care, efficiency, or growth
- Governance committee in place (or plan to establish one in Step 2)
- Documentation standards set for AI projects (charters, success criteria, review cadence)



Technology

- IT infrastructure reviewed for AI readiness (server infrastructure, network capacity)
- Integration points identified (EHR, PACS, departmental and end-user workflow software)
- Data sources inventoried (clinical datasets, imaging archives, EHR data)
- Data quality and accessibility evaluated
- Security and data governance policies reviewed for AI applications




Regulatory & Compliance

- HIPAA and privacy requirements confirmed for AI use
- FDA clearance needs assessed for clinical AI tools
- Emerging regulations monitored (e.g. state AI regulation considerations)
- Risk factors identified (patient safety, bias, cybersecurity) and mitigation plans drafted

This assessment checklist can be used to flag any areas that need attention (for instance, if data is a weakness, you might plan a data quality improvement project before deploying AI).

SWOT Analysis for AI Adoption:

It can be useful to perform a basic Strengths, Weaknesses, Opportunities, Threats analysis regarding AI in your organization. Strengths and weaknesses are internal (e.g. a strength could be a robust IT department or a digitally savvy radiology team; a weakness might be limited budget or lack of data scientists). Opportunities and threats are external (e.g. opportunity to improve patient outcomes or capture market share with AI-driven services; threats could include competitors adopting AI faster, or regulatory penalties for misuse). This strategic view helps executives frame the importance and urgency of AI efforts.



STEP 2: DEFINE

Strategy, Governance Structure and Policies

Once you understand your baseline, the next step is to put direction and discipline in place. **Define** is about turning assessment into action: setting a clear AI vision, aligning it with organizational priorities, and putting the structures and policies in place to guide adoption.

This is where leaders move from ideas to accountability. A strong governance framework makes AI adoption more than a collection of pilots. It makes it a program that is strategic, responsible, and measurable.



Establish the AI Vision and Goals:

Articulate a clear vision for AI in your organization. For example, your vision might be “to *augment clinical decision-making and improve patient outcomes through responsible AI.*” Set specific goals that support this vision, tied to your earlier assessment. Goals could include improving diagnostic accuracy by X%, reducing turnaround time, cutting administrative costs, or expanding access to care via AI tools. Having defined goals helps communicate the “*why*” to stakeholders and will guide decisions (focus on AI projects that drive these goals). Ensure these goals align with overall corporate strategy and patient care objectives.



Form an AI Governance Committee (or Working Group):

Strong governance requires clear structure and accountability. Most health systems begin with an interdisciplinary **AI Governance Committee** or working group, which may evolve into a permanent body or be integrated into existing IT/innovation committees. The committee’s role is to set policies, prioritize projects, evaluate proposals, and oversee AI performance throughout its lifecycle.

Building Your AI Governance Committee

Key roles and responsibilities:

Executive Sponsor (C-suite leader such as CIO, CMO, or CMIO): Champions the AI program, secures resources, aligns strategy, and communicates progress to leadership.

Clinical Champion (e.g. Radiology or Physician Leader): Provides frontline perspective, validates clinical relevance, advocates for adoption, and guides peer training.

IT/Data Science Lead: Assesses technical feasibility, manages integration with EHR/PACS, ensures infrastructure readiness and validates data quality and model performance.

Compliance & Legal Officer: Safeguards patient safety, privacy, and ethics; ensures regulatory requirements such as HIPAA and FDA clearance are met; develops transparency and liability guidelines.

Operations/Quality Representative: Aligns AI use with operational workflows, change management and performance improvement initiatives.

Project Manager: Coordinates day-to-day execution, vendor relationships, timelines and cross-department communication.



Key question: Determine whether AI governance should operate as a new, dedicated committee or as part of an existing structure. Larger organizations may even appoint a Chief AI Officer (CAIO), while others extend responsibilities to existing leaders such as the CIO or CMIO.



Develop AI Governance Policies and Processes:

With the team in place, formalize the **governance framework** in writing. This may include:

AI Project Intake & Evaluation Process: How will new AI ideas or vendor proposals be submitted and evaluated? Establish criteria (e.g. alignment to strategy, clinical validation required, ROI potential, risk level) and a workflow for review. *Standardizing the intake and evaluation process for AI tools is crucial.* For example, require departments to fill out an AI proposal form, then the AI committee scores it against set criteria.

Risk Assessment & Mitigation Policies: Define a process to assess risks of each AI (bias, errors, cybersecurity, etc.) and ensure mitigation plans. Possibly adopt or reference frameworks (like the NIST AI Risk Management Framework) to systematically address these concerns at each stage.

Data Use and Security Policies: Update data governance policies to cover AI. This includes how patient data is used to train or run AI (ensuring HIPAA compliance, de-identification if needed, and data retention policies). Also address data quality checks, since AI outcomes depend on good data.

Ethics Guidelines: Incorporate guiding principles such as fairness, transparency, accountability, and “**do no harm.**” For example, policies might require bias testing for algorithms on different demographic groups, or mandate that AI decisions affecting care are reviewable by clinicians (no “black box” blind trust).

Compliance Requirements: Outline steps to ensure regulatory compliance, such as verifying if an AI tool is FDA-approved for its intended use or falls under any regulatory exemptions. If operating in multiple regions, consider EU MDR or other local regulations.

Procurement and Vendor Management Procedures: Define how contracts with AI vendors will ensure responsibilities for quality and support (e.g. requiring service level agreements for model performance, vendor’s process for providing updates or addressing issues).

Integration with Existing Governance: If relevant, document how this AI governance intersects with other committees (IT governance, clinical quality, etc.) – for example, an AI project might need sign-off from the Pharmacy & Therapeutics committee if it affects medication decisions, etc.



Success Metrics and KPIs:

As part of strategy, decide how you will **measure the success** of AI implementations. Establish key performance indicators (KPIs) for both the program and individual projects. At a program level, a KPI might be “number of AI solutions successfully deployed and adopted by clinicians” or “percentage of AI projects meeting their ROI targets.” At the project level, metrics could be clinical (e.g. improvement in diagnostic accuracy or patient outcomes due to AI), operational (time saved per task, reduction in backlog), or financial (cost savings, revenue growth from new AI-enabled services). Define baseline measurements now (pre-AI) so that improvement can be quantified later. *For example, ask “What does success look like 5 years from now, and how will AI generate value for the organization?”*



Governance “Definition” Checklist:

Before moving to execution, ensure you have the following in place:

AI Strategy Documented: Vision, priority use cases, and goals are written and approved by leadership.

Governance Committee Formed: Members identified and charter approved (with leadership backing).

Roles Clearly Assigned: Each role (executive sponsor, clinical lead, etc.) knows their responsibilities in the AI initiative.

Policies Approved: AI project intake, risk management, ethics guidelines, and compliance check processes are defined and disseminated.

Resource Plan: Budget and resources (people, time) for initial AI projects are allocated. Responsibilities for working with vendors or developing models are clear.

Communication Plan: A plan to communicate AI governance and strategy to the wider organization (so that everyone understands the approach and how to propose ideas or raise concerns).

STEP 3: SELECT

Choosing AI Use Cases and Solutions

With strategy and governance in place, the next step is deciding where to place your bets. Selection is the point where decisions become real: which use cases to prioritize, which vendors to trust, and which tools to invest in. It is also where risk is highest — a poor choice can drain resources, frustrate clinicians, and damage trust in AI across the organization.

This phase is not about chasing shiny technology. It is about doing the homework to ensure the AI you choose is clinically effective, operationally practical, technically secure, and culturally adoptable.



Prioritize High-Impact Use Cases:

Likely, the assessment in Step 1 generated multiple potential AI opportunities. Now, prioritize them. Consider criteria such as: expected benefit (clinical impact or efficiency gain), feasibility (availability of data and technology, readiness of staff), alignment with strategic goals, and risk level. One approach is using a scoring system or matrix to rank projects. *For example, you might prioritize a radiology AI for flagging urgent findings over a more experimental research AI, because the former addresses an immediate patient safety issue (stroke or hemorrhage detection) and has proven tech.* Ensure diverse input in prioritization – what do frontline clinicians feel would help them most? Sometimes quick wins (projects that are easier to implement) can be prioritized to build momentum, as long as they align with strategy. The output of this step is a shortlist of AI projects/use cases to pursue first (with others in a backlog for later).



Make “Build vs Buy” Decisions:

For each chosen use case, decide whether to build a solution in-house (if you have a strong data science team) or acquire a third-party AI tool (from a vendor or open-source) or perhaps collaborate with a partner (academic or startup) to co-develop. In healthcare, buying or partnering is common due to the specialized nature of algorithms (e.g. for medical imaging AI, many hospitals work with vendors who have FDA-cleared tools). Building in-house might be feasible for non-clinical AI (like hospital operations optimization) if you have the talent. Weigh the cost, time, and expertise required. Many organizations use off-the-shelf AI tools but ensure they can be integrated and validated for their setting.



Evaluate and Select Vendors/Products:

Choosing the right AI partner is one of the most important governance decisions. The wrong choice can waste resources, erode clinician trust, and even harm patients. Use this framework to guide vendor due diligence. It combines core evaluation criteria with critical questions to ask so leaders can make confident, evidence-based decisions.

Vendor Evaluation Framework

Clinical Effectiveness

- Tool validated on patient data similar to ours
- Accuracy, sensitivity, specificity documented and independently verifiable
- Local validation possible before deployment

Ask: *What problem does this AI solve, who are the end-users, and how will we measure success after implementation?*

1

Workflow Integration

- Integrates seamlessly with PACS/EHR or existing systems
- Outputs presented where clinicians already work
- Configurable to reduce false positives and alert fatigue

Ask: *How will this tool fit into daily workflow? Will it save time or add steps?*

2

Technical & Security Requirements

- Infrastructure needs clear (on-prem, cloud, GPUs, storage)
- Compatible with current IT environment
- Cybersecurity and uptime protections in place

Ask: *What infrastructure is required, and how will data be secured and maintained?*

3

Regulatory & Compliance

- FDA clearance/CE mark verified for intended use
- HIPAA/privacy obligations documented with BAA if needed
- Transparent data handling (location, retention, retraining use)

Ask: *What approvals or certifications does this solution carry, and how will it remain compliant as regulations evolve?*

4

Vendor Reputation & Support

- Case studies and references available from similar institutions
- Financial stability and proven track record in healthcare
- Training, adoption support, and 24/7 technical assistance included

Ask: *What ongoing support, training, and monitoring do you provide after go-live?*

5

Cost & ROI

- Total cost of ownership defined (license, hardware, services, maintenance)
- ROI modeled across best- and worst-case scenarios
- Measurable benefits projected (efficiency gains, quality improvements, revenue opportunities)

Ask: *What evidence do you have of ROI from similar health systems?*

6

Risk Management

- False positive/negative rates disclosed and validated
- Liability and error-handling responsibilities contractually clear
- Revalidation and update process defined

Ask: *How are errors reported and managed, and what happens if performance drops?*

7

Scalability & Roadmap

- Capable of expanding across sites and specialties
- Vendor roadmap shows active investment and regulatory strategy

Ask: *How will this product evolve, and what is the long-term roadmap?*

8

Bias & Equity

- Performance tested across demographics (age, sex, race, socioeconomic groups)
- Results of bias audits shared or local testing supported

Ask: *How does this tool ensure fair performance across diverse patient populations?*

9

Transparency & Explainability

- Documentation available describing model training data, assumptions, and limitations
- Clinician-facing explanations provided where feasible

Ask: *How transparent is the model's decision-making, and can clinicians understand outputs?*

10

Governance In Practice

When evaluating AI partners, executives should look for transparent, accessible information that builds confidence. An example is **Rad AI's Trust Center**, which consolidates the key areas every health system or enterprise leader should expect from their vendor:



Data Privacy & Security

Clear policies on how patient and organizational data are collected, stored, protected, and deleted.



Model Transparency & Risk Management

Documentation of performance, limitations, and safeguards against bias or misuse.



Governance & Oversight

Defined accountability for decisions, with visible policies on ethics and responsible use.



Regulatory & Compliance Alignment

Evidence of adherence to HIPAA, GDPR, and other applicable standards.

Key takeaway: Ask every AI vendor to show where this information lives, how often it's updated, and who in their organization owns it. A dedicated Trust Center is a best-practice benchmark for vendor transparency.

These questions should form the backbone of any RFP or vendor questionnaire. Require vendors to provide detailed, evidence-based answers — for example, performance data specific to your specialty or results from validation studies. Where possible, arrange a limited trial or retrospective test on your own data to verify claims. Use a structured vendor scorecard to rate each criterion side by side, enabling the selection committee to make systematic, comparable decisions rather than relying on anecdotes or marketing materials.



Proof-of-Concept and Pilot Testing:

Pilots and proof-of-concepts can be valuable for testing an AI solution in a controlled way, but they are not without cost. Even when offered for free, pilots require staff time, IT resources, and workflow adjustments that can impact day-to-day operations. They can also create “pilot fatigue” if not well-structured or tied to clear success criteria. Retrospective testing on historical cases is one way to minimize disruption and validate performance before going live. The key is recognizing that pilots are useful tools, but they carry tradeoffs — and should be approached with the same rigor as any other investment.



Decision and Approval:

After evaluation and pilots, the AI governance committee (and possibly senior leadership, depending on investment size) makes the final go/no-go decision. If approved, ensure a formal sign-off on procurement or project kickoff is obtained, and that all stakeholders are informed. If an AI tool did not meet expectations during evaluation, either decide to abandon that solution and maybe revisit the use case later, or select the next best option. The governance process should allow a no-go decision without blame — not every promising AI will pan out, and it's better to stop a questionable project now than to push forward and have it fail in production.

Vendor Evaluation Scorecard

Criteria	Key Questions to Ask	Vendor Response / Evidence	Score (1–5)	Notes
Clinical Effectiveness	What clinical problem does this AI solve? Who are the end-users? How was it validated (population, metrics)? Can we test locally before go-live?			
Workflow Integration	How will the tool fit into existing workflows (PACS/EHR)? Does it reduce or add steps? Can alerts and outputs be customized?			
Technical & Security Requirements	What infrastructure is required (on-prem, cloud, GPUs)? Is it compatible with our IT environment? How is data secured and uptime maintained?			
Regulatory & Compliance	Is the solution FDA-cleared/CE-marked for intended use? How does it meet HIPAA/privacy standards? Where is data stored and how is it used?			
Vendor Reputation & Support	What similar customers are using this tool? Are case studies or references available? What training and 24/7 support is included?			
Cost & ROI	What is the total cost of ownership (licenses, hardware, services)? What ROI evidence exists from similar health systems?			
Risk Management	What are the false positive/negative rates? Who is liable if the AI fails? How are errors handled and revalidated over time?			

Criteria	Key Questions to Ask	Vendor Response / Evidence	Score (1–5)	Notes
Scalability & Roadmap	Can the solution scale across departments/sites? What is the vendor's roadmap for updates and future features?			
Bias & Equity	Has the solution been tested across diverse patient populations (age, sex, race, socioeconomic)? Will the vendor share bias audit results?			
Transparency & Explainability	Are model cards or documentation available (data sources, limitations, failure modes)? Can clinicians understand or interrogate outputs?			
Change Management Support	What adoption tools, training, and workflow support are provided? How will the vendor partner on culture and process change?			
Monitoring & Partnership	Does the vendor commit to ongoing performance monitoring and benchmarking (e.g., through registries)? What reporting is included?			
Contractual Safeguards	Are there performance guarantees in the contract? What provisions exist for updates, retraining, or regulatory changes?			

How to use

- Scoring: Rate each criterion 1–5 (1 = poor/does not meet needs, 5 = fully meets/exceeds expectations).
- Weighting (optional): Assign higher weight to priority areas (e.g., clinical effectiveness, workflow integration, compliance).
- Comparison: Use the table across multiple vendors to create a side-by-side comparison that drives objective decision-making.




STEP 4: EXECUTE

Implementing and Integrating AI Solutions

Execution is the stage where strategy becomes reality. It is not just about installing software. It is about embedding AI into the daily rhythm of care in a way that works for both people and technology.


The core principle: **AI must adapt to existing workflows, not force clinicians and staff to adapt to it.** Adoption sticks when frontline users feel the tool was designed with them, not imposed on them.

Execution succeeds when two dimensions move in lockstep:



Co-Design Around Workflow

Start execution by understanding the workflow you are entering. Shadow clinicians and staff to see how they work today and identify real bottlenecks. Map processes step by step and bring those users into the design.



If AI adds clicks, slows decisions, or forces workarounds, adoption will stall. The goal is for AI to be invisible in the workflow there when needed, unobtrusive when not.

Build a Phased Implementation Plan

Deploying AI requires discipline. A phased plan keeps projects on track while giving teams time to adapt.



Governance oversight should continue throughout, with checkpoints to review progress, risks, and user adoption.

Prepare and Engage Users

No AI succeeds without trust. Training and engagement are essential to making staff comfortable and confident.

- **Train the Trainers** — Identify super-users who receive deeper training and support peers.
- **Hands-On Training** — Use real cases to show how the AI helps, what it does not do, and how to interpret outputs.
- **Educational Materials** — Provide quick guides, FAQs, and scenarios for reference.
- **Set Expectations** — Reinforce that AI augments clinical expertise; it does not replace it.
- **Feedback Loops** — Give users easy ways to report issues or suggest improvements.

Communicate and Celebrate

Keep communication open across teams.

- Share regular updates during deployment (“Validation is complete,” “Pilot goes live next week”).
- Celebrate quick wins — highlight stories where AI caught a critical finding or reduced workload.
- Share performance metrics early to reinforce that the tool is delivering on its promise.

Confirm Readiness Before Full Go-Live

Before expanding AI system-wide, confirm that the essentials are in place:

Validation complete in our environment.

Workflow integration tested end-to-end.

Staff trained and equipped with guides.

IT and vendor support on standby for launch.

Monitoring and error-reporting active from day one.

Contingency plan ready for downtime or rollback.

If systems will be touching the clinical workflow check pre-defined care pathways and ensure any necessary approvals for information or referral access have been established.

Will AI results be retained and if so in what format.

Evaluate the regulatory framework not just for AI but for workflows related to Stark law and anti-kickback statutes.

Governance Oversight

The Governance Committee should stay involved during execution:

- Review adoption and performance data from pilots.
- Intervene if safety or ethical issues arise.
- Approve workflow adjustments or configuration changes.
- Document lessons learned to strengthen future deployments.

STEP 5: MONITOR

Ongoing Oversight, Performance Management, and Improvement

AI is not “set and forget.” Once live, it enters the most important phase: proving that it consistently delivers value, remains safe, and earns the trust of the people who rely on it. Monitoring is not about policing technology. It is about creating the feedback loop that keeps AI effective and clinicians confident.

Done right, monitoring turns AI from a pilot into a sustainable capability. Done poorly, tools become shelfware or, worse, undermine patient care.

What Monitoring Should Achieve



Performance Assurance — The AI continues to meet accuracy, sensitivity, and speed benchmarks in real-world conditions.



Workflow Confidence — Clinicians see that AI enhances their work rather than slowing them down.



Compliance & Safety — Tools remain aligned with privacy, security, and regulatory standards.



Continuous Improvement — Lessons from real-world use drive updates, refinements, or even retirement of underperforming tools.



How to Monitor in Practice



Track Real-World Performance

- Benchmark outcomes against the KPIs defined in Step 2 (e.g., diagnostic accuracy, turnaround times, efficiency gains, ROI).
- Audit random samples of AI-supported cases to check accuracy and relevance.
- Watch for model drift — performance degradation as patient populations or workflows evolve.
- Where available, participate in external benchmarking registries such as the **American College of Radiology's Assess-AI**. This registry collects real-world performance data across sites, enabling hospitals to compare their AI results against peers and national standards.



Listen to End-Users

- Gather regular feedback from clinicians and staff (surveys, focus groups, quick feedback channels).
- Track adoption rates. If clinicians are ignoring AI outputs, it's a warning sign.
- Use frontline input to recalibrate workflows, alert thresholds, or training needs.



Audit Risks and Compliance

- Ensure the AI is only being used for its intended and approved purposes.
- Stay ahead of evolving regulations (e.g., FDA guidance, EU AI Act).
- Run bias and equity checks to confirm consistent performance across different patient groups.
- Maintain cybersecurity audits and safeguards for data handling.



Plan for Maintenance and Updates

- Revalidate after vendor updates or new model releases.
- Keep clear documentation of versions, validation results, and governance decisions.
- Decide proactively when to scale, retrain, or sunset a tool that no longer delivers value.

Deployment Sustainment Checklist

KPIs tracked and reported regularly

User feedback loop in place

Participation in external benchmarking (e.g., ACR Assess-AI) considered

Compliance and bias audits scheduled

Vendor update/revalidation process documented

Annual portfolio review to expand, refine, or retire tools

The Governance Role

- Monitoring is not just an IT function. It is a governance responsibility. The AI Governance Committee should:
- Review quarterly performance and adoption reports
- Compare results against internal goals and external benchmarks (such as Assess-AI)
- Intervene if accuracy, safety, or ethics are in question
- Approve major updates or expansions
- Ensure successes and lessons are communicated across the organization

Making AI Work for Healthcare

AI will not transform healthcare on its own. Technology is the easy part. The real challenge and the real opportunity lies in governance.

When healthcare leaders apply discipline to how AI is assessed, defined, selected, executed, and monitored, they move beyond hype and create lasting impact. Governance bridges the C-suite vision with frontline practice. It turns scattered pilots into sustainable capabilities. And it ensures that AI adoption is not just fast, but safe, trusted, and valuable.

This is not a one-time project. Each cycle of governance makes the next smarter. A successful radiology deployment can open the door to pathology, cardiology, or operational AI. With every loop, the organization learns, adapts, and raises its standards. Over time, AI governance becomes as ingrained as quality or safety oversight, part of the fabric of how care is delivered.

The message is simple: **AI without governance is a gamble. AI with governance is a growth engine.**

The frameworks, questions, and checklists in this eBook are not meant to be followed blindly. They are starting points to adapt, expand, and make your own. If you use them well, you will not only maximize the return on your AI investments, but you will also do what matters most: deliver better care, build trust with clinicians, and protect patients.

Key Takeaways



- **Governance is the differentiator.** AI success depends less on the algorithm and more on disciplined oversight across people, process, and technology.
- **Adoption must fit workflows.** Tools succeed when they are designed with clinicians and staff, not imposed on them.
- **Monitoring builds trust.** Real-world performance, user feedback, and benchmarking (e.g., ACR Assess-AI) ensure AI stays accurate, safe, and equitable.
- **Governance is a cycle, not a project.** Each round of Assess, Define, Select, Execute, and Monitor makes the next smarter, turning AI from pilots into a lasting organizational capability.

Citations

1. Hill, D. L. G. (2024). AI in imaging: the regulatory landscape. *British Journal of Radiology*, 97(1155), 483–491. <https://doi.org/10.1093/bjr/tqae002>
2. *AI hasn't killed radiology, but it is changing it.* (2025, April). Washington Post. <https://www.washingtonpost.com/health/2025/04/05/ai-machine-learning-radiology-software/>