MAINSTAY TECHNOLOGIES

The Healthcare AI Readiness Guide

A role-based playbook for turning AI interest into governed, useful action.

- **Use-Case Triage**
- **Vendor Diligence**
- **30-Day Pilot**
- **Role Checklists**
- *and much more!*

For practice leadership, operations, finance, compliance, and internal IT.

Table of Contents

WHY THIS GUIDE

AI does not need to be all-or-nothing

The practical challenge is to separate helpful assistance from decisions that can create clinical, privacy, operational, financial, or reputational harm.

Healthcare organizations are already encountering AI in documentation, scheduling, communications, analytics, coding, support, security, and vendor platforms. The safest response is neither uncontrolled experimentation nor a policy that simply says "no." It is a repeatable process for deciding where AI belongs, what it may touch, who reviews it, and how performance is monitored.

No AI use case moves forward until the organization can name the purpose, the data involved, the accountable owner, the human review step, the evidence it will retain, and the conditions that would pause or stop use.

Use this guide in three ways

1. **LEADERSHIP DISCUSSION:** Align on what responsible adoption means before teams buy tools independently. Agree on risk tolerance, owners, and escalation paths.
2. **USE-CASE REVIEW:** Triage a proposed workflow as low, moderate, or high risk. Identify the approvals and evidence needed before a pilot.
3. **PILOT PLANNING:** Run a 30-day controlled test with baseline metrics and stop conditions. Decide to scale, revise, or retire based on evidence.

RESPONSIBLE ADOPTION

What responsible AI is - and is not

Responsible AI is an operating discipline, not a product claim.

It is

- A business and clinical governance discipline.
- A documented decision process.
- A set of controls around people, data, vendors, and workflows.
- Continuous review as tools and use cases change.

It is not

- A vendor checkbox.
- A promise that outputs are always correct.
- A replacement for legal, clinical, privacy, or security judgment.
- A one-time policy that sits on a shelf.

A practical standard

Responsible AI does not eliminate uncertainty. It makes purpose, ownership, data boundaries, review, evidence, and stop conditions explicit enough to manage that uncertainty.

THE PRACTICAL MODEL

Use the CARE framework

A simple operating model for safer AI decisions in healthcare.

C: CLARIFY THE OUTCOME

- Name the workflow and user.
- Define the decision AI will support - not make.
- Set a measurable baseline.

A: ASSESS DATA AND RISK

- Classify the data.
- Map patient, workforce, and operational impact.
- Identify applicable requirements.

R: REQUIRE CONTROLS AND ACCOUNTABILITY

- Use an approved tool.
- Define human review, access, logging, and fallback.
- Document vendor evidence.

E: EVALUATE IN PRACTICE

- Test representative cases.
- Track quality, risk, adoption, and cost.
- Scale, revise, or stop.

Framework note

CARE is a practical planning model, not a regulatory standard. It is designed to reinforce the same operating rhythm emphasized by the NIST AI Risk Management Framework: govern, map, measure, and manage AI risk.[1][2]

CONTROL BASELINE

Seven non-negotiables for every approved use

These are the minimum conditions for a use case to move from interest to controlled action.

1. NAMED OWNER - One person is accountable for outcome, risk, and review.

2. APPROVED TOOL - The platform and account type are authorized for the intended data.

3. DATA BOUNDARY - Users know what may and may not be entered, retrieved, stored, or shared.

4. HUMAN CHECKPOINT - A qualified person reviews outputs before they affect a patient, employee, payment, record, or decision.

5. TEST PLAN - The pilot uses representative cases and a documented baseline.

6. TRACEABILITY - The organization can reconstruct who used the system, what changed, and what was approved.

7. CHANGE CONTROL - Material vendor, model, integration, or workflow changes trigger re-review.

Good governance test

Could a new team member understand the approved use, the prohibited use, the review step, and the escalation path without asking the person who bought the tool?

USE-CASE TRIAGE

Start with the use case, not the tool

Risk should follow consequence, data sensitivity, and reversibility - not how impressive the product demo looks.

GREEN | ASSISTIVE

TYPICAL CHARACTERISTICS - Low consequence; no sensitive data; reversible; human reviews before use.

HEALTHCARE EXAMPLES - Brainstorming with public information; formatting; drafting internal materials; summarizing non-sensitive meetings.

MINIMUM DECISION GATE - Business owner approval + approved tool + basic data rules.

YELLOW | CONTROLLED

TYPICAL CHARACTERISTICS - Sensitive data or operational impact; errors could create delay, rework, privacy risk, or inequitable service.

HEALTHCARE EXAMPLES - Patient-message drafts; scheduling or intake support; coding or revenue-cycle assistance; records-request classification; workforce workflows.

MINIMUM DECISION GATE - Workflow owner + IT/security + privacy/compliance review; limited pilot; explicit human checkpoint.

RED | ESCALATE

TYPICAL CHARACTERISTICS - Can influence diagnosis, treatment, triage, patient safety, access, legal rights, or other hard-to-reverse decisions.

HEALTHCARE EXAMPLES - Clinical decision support; autonomous patient advice; diagnostic or treatment recommendations; prioritization that affects care access; regulated device functions.

MINIMUM DECISION GATE - Clinical owner + executive sponsor + privacy/security + legal/regulatory review; formal validation and monitoring plan.

Five questions that usually change the risk tier

Use these questions before assigning a color, approving a pilot, or expanding an existing workflow.

1. What happens if the output is wrong?

Consider patient harm, access delays, revenue impact, privacy exposure, and trust.

2. Can a person meaningfully review it?

A ceremonial "human in the loop" is not a control if the reviewer lacks time, context, or authority.

3. What data enters and leaves the system?

Include prompts, attachments, retrieved data, logs, analytics, and vendor support access.

4. Is the decision reversible?

The harder it is to correct after the fact, the stronger the pre-use controls must be.

5. Will the workflow change over time?

Model updates, integrations, new users, and expanded scope can turn yesterday's low-risk use into today's high-risk use.

Clinical ownership is not optional

Any use that can influence diagnosis, treatment, triage, or patient safety needs an accountable clinical owner - even when the tool is purchased by operations, finance, compliance, or IT. Products that function as medical devices may also require FDA-specific review.[7]

BEFORE THE PILOT

Put the rules around the workflow

A responsible pilot starts with a controlled operating environment – not just a test account.

1. Data boundary

- Classify the information the workflow may use.
- Apply minimum-necessary thinking.
- Prohibit copy/paste into unapproved personal or public accounts.
- Document retention and deletion.

2. Vendor and contract

- Determine whether the vendor is a business associate.
- Execute a BAA when required - and review the underlying service terms.
- Confirm subcontractors, training use, support access, breach notice, and exit rights.

3. Security controls

- Use managed identities, least privilege, MFA, and role-based access.
- Require encryption, audit logs, and administrative visibility.
- Define approved integrations and block unmanaged connectors.

BEFORE THE PILOT

Build human, validation, and monitoring controls

The pilot should be reviewable, reversible, and observable from the first day.

4. Human oversight

- Name the reviewer and required qualifications.
- Specify what must be verified before use.
- Define fallback workflow and escalation authority.
- Do not automate a decision that cannot be meaningfully reviewed.

5. Validation and equity

- Test representative cases, edge cases, and failure modes.
- Check accuracy, consistency, accessibility, and differential impact.
- Document limitations and prohibited uses.

6. Monitoring and change

- Track output quality, complaints, incidents, and near misses.
- Re-review after material vendor, model, data, or workflow changes.
- Maintain stop conditions and an incident response path.

BEFORE THE PILOT

Regulatory anchors to keep in view

These reference points do not replace qualified legal, privacy, security, clinical, or regulatory review.

HIPAA Security Rule

Covered entities and business associates must protect ePHI through appropriate administrative, physical, and technical safeguards. A risk analysis is foundational to that work.[3][4]

Business associates

A written business associate contract is required when a vendor performs covered functions involving PHI on behalf of a covered entity or business associate. A BAA is necessary where applicable, but it is not the full governance program.[5]

Health data outside HIPAA

Some health apps and related entities outside HIPAA may fall under the FTC Health Breach Notification Rule. Do not assume "not HIPAA" means "not regulated."[6]

Certified health IT and medical devices

HTI-1 establishes algorithm-transparency requirements for predictive decision support in certified health IT; AI-enabled device functions may raise FDA considerations.[7][8]

As of July 2026

OCR's proposed HIPAA Security Rule modifications remain a proposal. The current Security Rule remains in effect. Build to current requirements, monitor the proposal, and avoid presenting future requirements as final.[9]

VENDOR DUE DILIGENCE

Ask the questions the demo will not answer

A polished interface does not tell you how data is used, how models change, or what happens when the system fails.

Data and privacy

1. Exactly what data is collected, generated, logged, retained, or shared?
2. Are prompts, files, and outputs used to train or improve any model?
3. Can training use be disabled contractually and technically?
4. What are the retention, deletion, export, and legal-hold options?

Security and access

1. Does the service support SSO, MFA, role-based access, and managed accounts?
2. What administrative logs are available and how long are they retained?
3. How are data encrypted in transit and at rest?
4. Who - including subcontractors and support staff - can access customer data?

Evaluate performance, workflow, and exit

The operating model matters as much as the feature set.

Model and performance

1. What is the intended use, and what uses are explicitly unsupported?
2. How was performance evaluated for this context and population?
3. What limitations, known failure modes, and bias risks are documented?
4. How are model or version changes communicated, tested, and rolled back?

Workflow and operations

1. How does the tool integrate with the EHR, portal, identity, records, or other systems?
2. What happens when the service is unavailable or produces uncertain output?
3. Can staff review, correct, and override outputs before downstream action?
4. What support, incident notification, and escalation commitments are documented?

Commercial and exit

1. What is the fully loaded cost: licenses, implementation, integration, security, training, and support?
2. What usage limits or price escalators can change the economics?
3. Who owns prompts, configurations, workflows, outputs, and derived data?
4. Can the organization exit cleanly and retrieve or delete its data?

Evidence to request

Do not rely on verbal assurances when the decision depends on data handling, performance, or control.

- Security documentation.
- Data-flow diagram.
- BAA and service terms.
- Subprocessor list.
- Retention and deletion settings.
- Model or system card.
- Validation summary.
- Change-notice process.
- Audit-log sample.
- Incident and business-continuity procedures.
- Accessibility information.
- Customer references for comparable use.

Review principle

The evidence should be specific enough to support the intended use, not merely demonstrate that the vendor has a general security or AI program.

THE ROLE MAP

One guide, five decisions

Responsible AI becomes practical when each leader owns a distinct decision.

Practice administrator

- Likely concern: staff frustration, communication delays, daily disruption.
- Outcome: a smoother workday.
- Decision to own: where AI helps staff without adding rework.
- First metric: rework and escalations.

Operations leader

- Likely concern: workflows, access, phones, portals, records.
- Outcome: reliable care-supporting operations.
- Decision to own: how AI changes the process and fallback.
- First metric: cycle time and exceptions.

Finance leader

- Likely concern: cost, internal expertise, tradeoffs.
- Outcome: clear short- and long-term economics.
- Decision to own: whether value exceeds total cost and risk.
- First metric: fully loaded cost and verified benefit.

Compliance contact

- Likely concern: PHI, HIPAA, evidence, breach readiness.
- Outcome: confidence requirements were considered.
- Decision to own: what data, evidence, approval, and notice are required.
- First metric: evidence completeness and open findings.

Internal IT leader

- Likely concern: capacity, security, applications, planning.
- Outcome: specialist leverage without losing control.
- Decision to own: architecture, access, integration, logging, and support boundary.
- First metric: MTTR, change failure, and access exceptions.

Shared governance and a useful meeting agenda

Responsible AI becomes practical when each leader owns a distinct decision and the group follows a consistent review sequence.

Shared governance

A business or clinical owner approves the workflow. IT controls the platform and integration. Compliance/privacy sets data and evidence requirements. Finance validates the economics. Practice and operations leaders confirm the workflow actually works. Add a clinical owner whenever patient-care judgment or safety could be affected.

A useful meeting agenda

Use this sequence to turn a broad AI conversation into a documented decision.

1. PURPOSE: What workflow are we improving and for whom?

2. RISK: What could go wrong, who could be affected, and how reversible is it?

3. DATA: What information enters, leaves, and remains in the system?

4. CONTROLS: Who reviews, what is logged, and what is the fallback?

5. EVIDENCE: What will we retain to show the decision and performance?

6. DECISION: Pilot, revise, defer, or decline - with owner and next review date.

ROLE-BASED ACTION PLAN

Practice Administrator

Primary outcome: A smoother workday for clinical and administrative teams. "Your staff should be able to get help without chasing updates or translating IT issues on their own."

What this role is trying to protect

Staff frustration, unresolved issues, patient communication delays, and day-to-day disruption.

What responsible AI means here

AI should remove friction from a clearly defined task without creating extra review, tool-switching, or uncertainty for already-busy staff.

Start with these questions

1. Which repetitive task creates the most avoidable staff frustration today?

2. Can the workflow be improved without exposing PHI or moving staff into unmanaged tools?

3. Who reviews the output before it reaches a patient, chart, payer, or outside party?

4. What would make staff stop using the tool and return to the existing workflow?

Practice Administrator: Pilot guidance

Pilot candidates, watch-outs, measures, and the first 30-day action.

Good pilot candidates

- Draft internal SOPs from approved source material.
- Summarize non-sensitive meetings and action items.
- Create staff FAQs from approved content.
- Draft patient communications inside an approved environment with human review.

Watch-outs

- Public or personal AI accounts.
- Copying patient identifiers into unapproved tools.
- Outputs that take longer to fix than to create.
- Different departments inventing different rules.

How to measure progress

- Minutes saved per completed task - verified, not estimated.
- Output rework or correction rate.
- Staff escalations and support requests.
- Staff confidence and adoption after training.

Your first 30-day action

Choose one high-frequency administrative workflow. Map the current steps, identify where time is lost, remove PHI unless it is essential, and define the exact human review before testing a tool.

Operations Leader

Primary outcome: Reliable operations that support care delivery and reduce bottlenecks. "When healthcare workflows depend on technology, IT reliability becomes an operational priority."

What this role is trying to protect

Systems that slow patient access, phones, portals, records requests, scheduling, or other core workflows.

What responsible AI means here

AI is an operational change. It must fit the real process, handle exceptions, degrade safely, and leave staff with a clear fallback when the system is unavailable or uncertain.

Start with these questions

1. Where does work queue up, bounce between teams, or require repeated re-entry?

2. Which exceptions are common enough that a pilot must include them?

3. Could the AI output delay, misroute, or disadvantage a patient?

4. What manual fallback keeps the workflow moving during downtime?

Operations Leader: Pilot guidance

Pilot candidates, watch-outs, measures, and the first 30-day action.

Good pilot candidates

- Classify and route incoming requests.
- Summarize call or intake notes for staff review.
- Support scheduling or capacity analysis.
- Identify bottlenecks in records, referral, or authorization workflows.

Watch-outs

- Testing only the happy path.
- Hidden dependencies on EHR, portal, telephony, or identity systems.
- No downtime or rollback plan.
- Optimizing speed while increasing exceptions or patient confusion.

How to measure progress

- End-to-end cycle time and queue age.
- Exception, reroute, and manual override rate.
- Patient-access or service-level impact.
- Downtime recovery and fallback performance.

Your first 30-day action

Select a workflow with measurable delays. Build a simple process map that includes handoffs, exception paths, systems touched, data used, and the fallback. Use that map as the pilot scope.

Finance Leader

Primary outcome: Better visibility into cost impact, timing, and long-term implications. "Good IT guidance should make the financial tradeoffs clearer before you commit."

What this role is trying to protect

The cost of internal expertise, unpredictable consulting, and unclear short- versus long-term tradeoffs.

What responsible AI means here

AI should be funded as a business change, not just a license. The business case must include implementation, integration, security, governance, training, support, and exit costs - plus the cost of errors and rework.

Start with these questions

1. What baseline cost, delay, or capacity constraint are we actually trying to change?

2. What costs sit outside the vendor quote?

3. How will we verify benefits instead of relying on projected hours saved?

4. What is the downside if adoption is low, accuracy is weak, or the vendor changes pricing?

Finance Leader: Pilot guidance

Pilot candidates, watch-outs, measures, and the first 30-day action.

Good pilot candidates

- Summarize contracts or invoices for qualified review.
- Support budget scenarios and variance explanations.
- Compare procurement options using approved inputs.
- Assist revenue-cycle analysis without making autonomous coverage or payment decisions.

Watch-outs

- License-first buying with no workflow owner.
- Savings based only on self-reported time.
- Duplicate tools and unmanaged departmental subscriptions.
- Ignoring security, compliance, integration, training, and exit costs.

How to measure progress

- Fully loaded monthly and one-time cost.
- Verified capacity released or cycle-time improvement.
- Cost per completed transaction or task.
- Error, rework, and support cost.
- Payback range with upside, base, and downside cases.

Your first 30-day action

Require a one-page business case before purchase: current baseline, target outcome, fully loaded cost, risk assumptions, success measures, decision date, and exit cost.

Compliance / Privacy Contact

Primary outcome: Confidence that requirements were considered before decisions were made. "Before choosing a tool or process, clarify the security, HIPAA, and evidence questions it needs to satisfy."

What this role is trying to protect

Sensitive patient information, HIPAA expectations, funder or contract requirements, breach readiness, and evidence.

What responsible AI means here

Responsible AI means being able to show what was approved, what data is allowed, which vendor terms apply, who reviews outputs, how incidents are handled, and when the use case was last reassessed.

Start with these questions

1. Does the workflow involve PHI, other sensitive data, or a business-associate relationship?

2. What evidence proves access, training, validation, and review are working?

3. Does the vendor use customer data to train models or permit support access?

4. What event triggers privacy, security, legal, clinical, or regulatory escalation?

Compliance / Privacy Contact: Pilot guidance

Pilot candidates, watch-outs, measures, and the first 30-day action.

Good pilot candidates

- Index and summarize compliance evidence for human review.
- Compare policies, contracts, or questionnaires against approved requirements.
- Draft risk-register updates from verified source material.
- Prepare audit support materials without replacing legal or compliance judgment.

Watch-outs

- Treating a BAA as the entire approval.
- No record of data flows, approvals, or model changes.
- Policy text generated and adopted without qualified review.
- No incident-notice, evidence-retention, or reassessment process.

How to measure progress

- AI use register completeness.
- Open findings and average closure time.
- Training completion and policy acknowledgement.
- Privacy/security incidents and near misses.
- Vendor evidence freshness.

Your first 30-day action

Create an AI use register. For every proposed use, capture owner, tool, data, risk tier, vendor/BAA status, human review, metrics, approval, and next review date.

Internal IT Leader

Primary outcome: More capacity and specialist access without losing internal knowledge or control. "Your internal team should not have to be the expert in every security, application, infrastructure, and planning question."

What this role is trying to protect

A capable but stretched team balancing tickets, security, applications, infrastructure, vendor change, and planning.

What responsible AI means here

Responsible AI needs an approved technical pattern: managed identities, access control, logging, integration standards, test and rollback, support ownership, and a clean boundary between platform support and workflow execution.

Start with these questions

1. What is the approved account, tenant, identity, and access pattern?

2. Which integrations, connectors, plugins, or retrieval sources are allowed?

3. Can we observe use, investigate an incident, and roll back a change?

4. Who supports the platform, the workflow, the data, and the underlying business process?

Internal IT Leader: Pilot guidance

Pilot candidates, watch-outs, measures, and the first 30-day action.

Good pilot candidates

- Summarize tickets and draft knowledge articles.
- Assist scripts or configuration review under change control.
- Triage alerts or logs for qualified analyst review.
- Create test cases, documentation, and runbook drafts.

Watch-outs

- Personal accounts and unmanaged API keys.
- Auto-executing code or changes without review.
- Opaque connectors with broad permissions.
- Silent model changes, no test environment, or no rollback.
- Undefined support boundaries after a pilot.

How to measure progress

- Mean time to resolve and time to triage.
- Change failure and rollback rate.
- False positives, false negatives, and manual overrides.
- Access exceptions and unmanaged-tool discoveries.
- Hours shifted from reactive work to higher-value priorities.

Your first 30-day action

Publish one approved AI architecture pattern covering identity, data classes, logging, integrations, test/production separation, change control, and support boundaries. Make exceptions explicit.

CONTROLLED ADOPTION

A 30-day pilot that creates evidence

Keep the pilot small, reversible, observable, and tied to a real workflow.

DAYS 1-5 | DEFINE

Name the workflow, owner, users, baseline, data boundary, risk tier, human review, and stop conditions.

Evidence at exit: pilot charter + current-state baseline.

DAYS 6-10 | APPROVE

Complete vendor diligence, contract or BAA review, security configuration, access, logging, and integration approval.

Evidence at exit: approval record + vendor evidence file.

DAYS 11-15 | PREPARE

Build representative test cases, edge cases, failure scenarios, user instructions, and fallback workflow.

Evidence at exit: test set + acceptance criteria + training.

DAYS 16-23 | RUN

Pilot with a limited cohort. Review outputs, capture corrections, monitor incidents, and hold short feedback huddles.

Evidence at exit: usage log + quality, risk, and adoption data.

DAYS 24-30 | DECIDE

Compare results with baseline. Review costs, complaints, near misses, and control performance. Scale, revise, or stop.

Evidence at exit: decision memo + next review date.

PILOT CHARTER

Complete before access is granted

Capture the decision in plain language before the pilot begins.

1. Use case and intended outcome

What workflow is being tested, for whom, and what result should improve?

2. Accountable owner / clinical owner if applicable

Name the person who owns the outcome, risk, and review.

3. Approved users, tool, and data classification

Specify who may use what account, with which data.

4. Required human review and fallback

State what must be checked and how work continues if the tool is unavailable or uncertain.

5. Success measures and baseline

Record the starting point, target, data source, and review frequency.

6. Stop conditions and decision date

Define what pauses the pilot and when leaders will decide to scale, revise, or stop.

Pilot discipline

Do not expand the user group, data class, integration, or decision authority during the pilot without written re-approval. Expansion is a scope change, not a casual next step.

MINIMUM VIABLE GOVERNANCE

Six artifacts that make AI governable

You do not need a large committee. You do need a durable record of decisions, ownership, and evidence.

1. AI USE REGISTER: A single inventory of proposed, piloted, approved, paused, and retired uses.

2. ACCEPTABLE-USE GUIDE: Plain-language rules for approved tools, data, verification, prohibited uses, and escalation.

3. PILOT CHARTER: Purpose, owner, users, data, controls, metrics, timeline, and stop conditions.

4. VENDOR EVIDENCE FILE: Contracts, BAA, sub-processors, data flow, security, validation, changes, incidents, and exit.

5. DECISION LOG: Who approved what, on which evidence, with which limitations and next review date.

6. INCIDENT PATH: Who to contact, what to preserve, how to contain, and who decides whether use may resume.

Governance principle

The record should be simple enough to maintain and specific enough to explain what was allowed, why it was allowed, and what must happen when conditions change.

Starter AI use register and ownership model

Use text fields rather than a grid when building the first version.

Starter AI use register

USE CASE _____

ACCOUNTABLE OWNER _____

TOOL AND ACCOUNT TYPE _____

DATA CLASSIFICATION _____

RISK TIER _____

REQUIRED HUMAN REVIEW _____

CURRENT STATUS _____

NEXT REVIEW DATE _____

Ownership model

EXECUTIVE SPONSOR: Sets risk tolerance, resolves tradeoffs, and funds the work.

WORKFLOW OWNER: Owns outcome, process, users, fallback, and adoption.

CLINICAL OWNER: Required when clinical judgment, patient safety, diagnosis, treatment, or triage may be affected.

IT / SECURITY: Owns platform pattern, access, integration, logging, monitoring, and technical incident response.

PRIVACY / COMPLIANCE / LEGAL: Interprets data, contract, evidence, notification, and regulatory requirements.

FINANCE / PROCUREMENT: Validates total cost, terms, concentration risk, and exit.

MEASUREMENT

Measure the pilot like an operational change

Hours saved are useful, but they are not enough. A responsible decision balances outcome, quality, risk, adoption, and cost.

Outcome

- Did the workflow improve the intended result?
- Example measures: cycle time, throughput, backlog, service level.
- Evidence: baseline and pilot comparison.

Quality

- Are outputs accurate, consistent, usable, and reviewable?
- Example measures: correction rate, false positive/negative, reviewer agreement.
- Evidence: sample review and error log.

Risk

- Did the pilot create privacy, security, clinical, equity, or operational concerns?
- Example measures: incidents, near misses, complaints, access exceptions.
- Evidence: risk and incident log.

Adoption

- Can people use the process correctly and consistently?
- Example measures: active use, completion, training, support requests.
- Evidence: usage and training records.

Cost

- Does value exceed the fully loaded cost?
- Example measures: license + implementation + support, verified capacity, rework.
- Evidence: TCO and benefit range.

Decision rules and predefined stop conditions

A pilot is successful when it produces a clear decision - not simply when it reaches day 30.

Scale

- Outcome improved against baseline.
- Quality and control thresholds were met.
- Users followed the approved workflow.
- Costs and support model are understood.

Revise

- Value is promising but errors, adoption, workflow fit, or controls need adjustment.
- The next test is bounded and documented.

Stop

- Patient, privacy, security, equity, or operational risk exceeds tolerance.
- Outputs cannot be validated or reviewed.
- Vendor terms, logging, or support are inadequate.
- The workflow creates more rework than value.

Predefined stop conditions

- PHI or other restricted data is entered into an unapproved account, tool, connector, or workflow.
- Required human review is bypassed, rushed, or cannot be performed meaningfully.
- A patient-facing, financial, workforce, or clinical output causes material error, delay, unfair impact, or complaint.
- The vendor materially changes data use, model, subprocessors, integration, retention, or pricing without review.
- Audit logs, access controls, rollback, or required evidence are unavailable.
- The pilot exceeds its approved users, data class, integration, or decision authority.

WHEN SOMETHING GOES WRONG

Respond to AI issues like operational incidents

Contain first. Preserve evidence. Classify the impact. Do not silently correct the output and move on.

1. CONTAIN: Pause the use case or affected workflow; revoke access or disable integration if needed.

2. PRESERVE: Retain prompts, inputs, outputs, logs, timestamps, model or version, users, and vendor notices.

3. CLASSIFY: Determine privacy, security, clinical, operational, financial, legal, and patient-communication impact.

4. NOTIFY: Engage the workflow owner, IT/security, privacy/compliance, clinical leadership, legal, insurer, or vendor as required.

5. REMEDIATE: Correct downstream records or communications, reduce access, fix configuration, retrain users, or change the workflow.

6. DECIDE: Resume with changes, narrow the use, revalidate, or retire. Document the decision and next review.

Decision memo

At day 30, write one page: what changed, what did not, what went wrong, what evidence exists, what controls remain, and why the organization chose to scale, revise, or stop.

Scenario prompts

Use these examples to identify the immediate action and the primary escalation path.

PHI entered into an unapproved tool

- **IMMEDIATE ACTION:** Stop use, preserve evidence, assess disclosure and vendor access.
- **PRIMARY ESCALATION:** Privacy/compliance + security.

Incorrect patient communication

- **IMMEDIATE ACTION:** Stop automated or drafted send, correct communication, assess affected patients.
- **PRIMARY ESCALATION:** Operations + clinical + privacy.

Biased or inconsistent output

- **IMMEDIATE ACTION:** Pause affected decision path, review population and test cases.
- **PRIMARY ESCALATION:** Clinical/operations + compliance.

Vendor terms or model change

- **IMMEDIATE ACTION:** Freeze expansion, review data use, validation, and contract impact.
- **PRIMARY ESCALATION:** IT + compliance + procurement.

Workflow or integration outage

- **IMMEDIATE ACTION:** Use fallback process; isolate failed connector; protect data integrity.
- **PRIMARY ESCALATION:** IT + operations.

Important

Whether an event is a reportable breach, safety event, contractual incident, or other notifiable event depends on the facts and applicable requirements. Escalate to qualified privacy, security, clinical, and legal resources rather than making that determination informally.

READINESS CHECK

Twelve questions before you scale

Score 0 = no, 1 = in progress, 2 = yes. This is a planning tool, not a compliance assessment.

1. We have an executive sponsor and a named owner for each AI use case.

[ ] 0 [ ] 1 [ ] 2

2. We maintain an inventory of approved, piloted, paused, and prohibited uses.

[ ] 0 [ ] 1 [ ] 2

3. Users know which AI tools and account types are approved.

[ ] 0 [ ] 1 [ ] 2

4. We have clear rules for PHI, sensitive data, prompts, attachments, outputs, and retention.

[ ] 0 [ ] 1 [ ] 2

5. Vendor review covers contracts or BAA, data use, subprocessors, security, logging, model change, and exit.

[ ] 0 [ ] 1 [ ] 2

6. Each use case has an explicit human review step and fallback workflow.

[ ] 0 [ ] 1 [ ] 2

Complete the assessment and interpret the result

Use the score to choose the next operating step.

7. Clinical leadership owns any use that can influence clinical judgment or patient safety.

[ ] 0 [ ] 1 [ ] 2

8. Pilots use representative test cases, edge cases, and documented acceptance criteria.

[ ] 0 [ ] 1 [ ] 2

9. We can investigate use through identities, logs, versions, and decision records.

[ ] 0 [ ] 1 [ ] 2

10. Staff receive role-specific training and know how to report an AI concern or near miss.

[ ] 0 [ ] 1 [ ] 2

11. We measure outcome, quality, risk, adoption, and total cost against a baseline.

[ ] 0 [ ] 1 [ ] 2

12. Material changes to tool, model, data, integration, or workflow trigger re-review.

[ ] 0 [ ] 1 [ ] 2

Interpret the result

0-8 | BUILD THE FOUNDATION: Do not scale. Establish approved tools, data rules, ownership, vendor diligence, and incident reporting. Choose a low-risk workflow for a controlled start.

9-16 | PILOT CAREFULLY: The basics are forming. Close the highest-risk gaps, limit scope, and use a documented 30-day pilot.

17-24 | READY FOR CONTROLLED ADOPTION: Governance is strong enough to support measured expansion. Continue monitoring and re-review after material change.

NEXT STEP

Make the next AI decision clearer

AI readiness is a workflow, data, governance, and adoption decision before it becomes a software decision.

A practical starting point: Mainstay's Embark Project

Designed for organizations that need platform readiness, practical guardrails, and initial training before they build custom AI workflows.

- Select and configure one approved AI platform.
- Create a practical "How We Use AI Here" acceptable-use guide.
- Train an initial cohort with role-relevant examples.
- Prioritize two or three next opportunities and route them to the right delivery path.

Best fit

- No approved AI platform, or inconsistent use across individual accounts.
- Need practical guardrails and training before wider adoption.
- AI interest exists, but the first use case is unclear or risk-sensitive.
- Leadership wants a contained starting point before custom builds.

Explore AI & Automation at: mstech.com/ai

Choose the right delivery path

Use the readiness discussion to determine whether the next move is platform readiness, a controlled build, ongoing execution, or another path.

Use another path when

- You already have an approved platform and a clearly defined workflow to build.
- You need ongoing build, tuning, maintenance, or workflow support.
- You need formal policy authorship, legal advice, or compliance documentation beyond practical acceptable-use guidance.
- You need dedicated long-term AI strategy and roadmap ownership after results exist.

Book a 20-minute Healthcare AI Readiness Fit Call

Bring one workflow, one concern, and one target outcome.

We will help clarify whether the right next step is readiness, a controlled build, ongoing execution, or a different path.

New client inquiries

(844) 458-4566 • info@mstech.com • mstech.com/contact

SOURCES AND NOTES

Reference points for responsible adoption

Sources reviewed for this guide (July 2026).

[1] NIST, AI Risk Management Framework 1.0 - https://www.nist.gov/itl/ai-risk-management-framework

[2] NIST, AI RMF: Generative Artificial Intelligence Profile - https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

[3] HHS OCR, The HIPAA Security Rule - https://www.hhs.gov/hipaa/for-professionals/security/index.html

[4] HHS OCR, Guidance on Risk Analysis - https://www.hhs.gov/hipaa/for-professionals/security/guidance/guidance-risk-analysis/index.html

[5] HHS OCR, Business Associate Contracts - https://www.hhs.gov/hipaa/for-professionals/covered-entities/sample-business-associate-agreement-provisions/index.html

[6] Federal Trade Commission, Health Breach Notification Rule - https://www.ftc.gov/legal-library/browse/rules/health-breach-notification-rule

[7] U.S. Food and Drug Administration, Artificial Intelligence-Enabled Medical Devices - https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices

[8] ASTP/ONC, HTI-1 Final Rule and Decision Support Interventions resources - https://www.healthit.gov/regulations/hti-rules/hti-1-final-rule/

[9] HHS OCR, HIPAA Security Rule Notice of Proposed Rulemaking - https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/index.html

Important disclaimer

Use this guide as a planning aid and engage qualified professionals for decisions that affect regulated or high-consequence work.

This guide provides general educational and planning information. It is not legal, regulatory, medical, clinical, privacy, security, accounting, or procurement advice; it does not determine whether any organization, product, or use case is compliant, safe, effective, or appropriate.

Requirements depend on the organization, data, workflow, jurisdiction, contracts, and product function. Engage qualified clinical, legal, privacy, security, compliance, and regulatory professionals for decisions that affect patient care, rights, safety, protected information, or regulated products.

Responsible AI is not the total absence of risk. It is the discipline to make the intended use, data, accountability, controls, evidence, and stop conditions clear - and to keep them clear as the tool and workflow change.

Why Mainstay?

Built for the Long View
- Mainstay has always led ahead of the curve—building a full information security program long before it became standard. We bring that same foresight to AI & automation.

Governed From Day One
- Clear ownership, guardrails, and accountability turn AI from experimentation into trusted, repeatable work.

Real Outcomes, Not Hype
- We help teams move from tools and pilots to workflows that support daily operations and strategy.

Expertise Without Overload
- Fractional leadership, hands-on guidance, and automation support—without adding strain to internal teams.

Designed to Scale
- A clear path from first AI wins to AI as a strategic engine, at a pace organizations can sustain.

Visit us at mstech.com/ai or scan the QR code: