CitiusTech

Shaping Healthcare Possibilities

Designing responsible AI for patient hubs

Faster access without compromising trust

Harisha Cherla
Software Engineer,
CitiusTech

Pooja Nair
Sr Engineering Manager,
CitiusTech

Ravindra Bagli
Engineering Manager,
CitiusTech

Vimal Prasath
Sr Lead Engineer - I,
CitiusTech

Enrolment, OCR, and Predictive Triage

Version 2.1

Contents

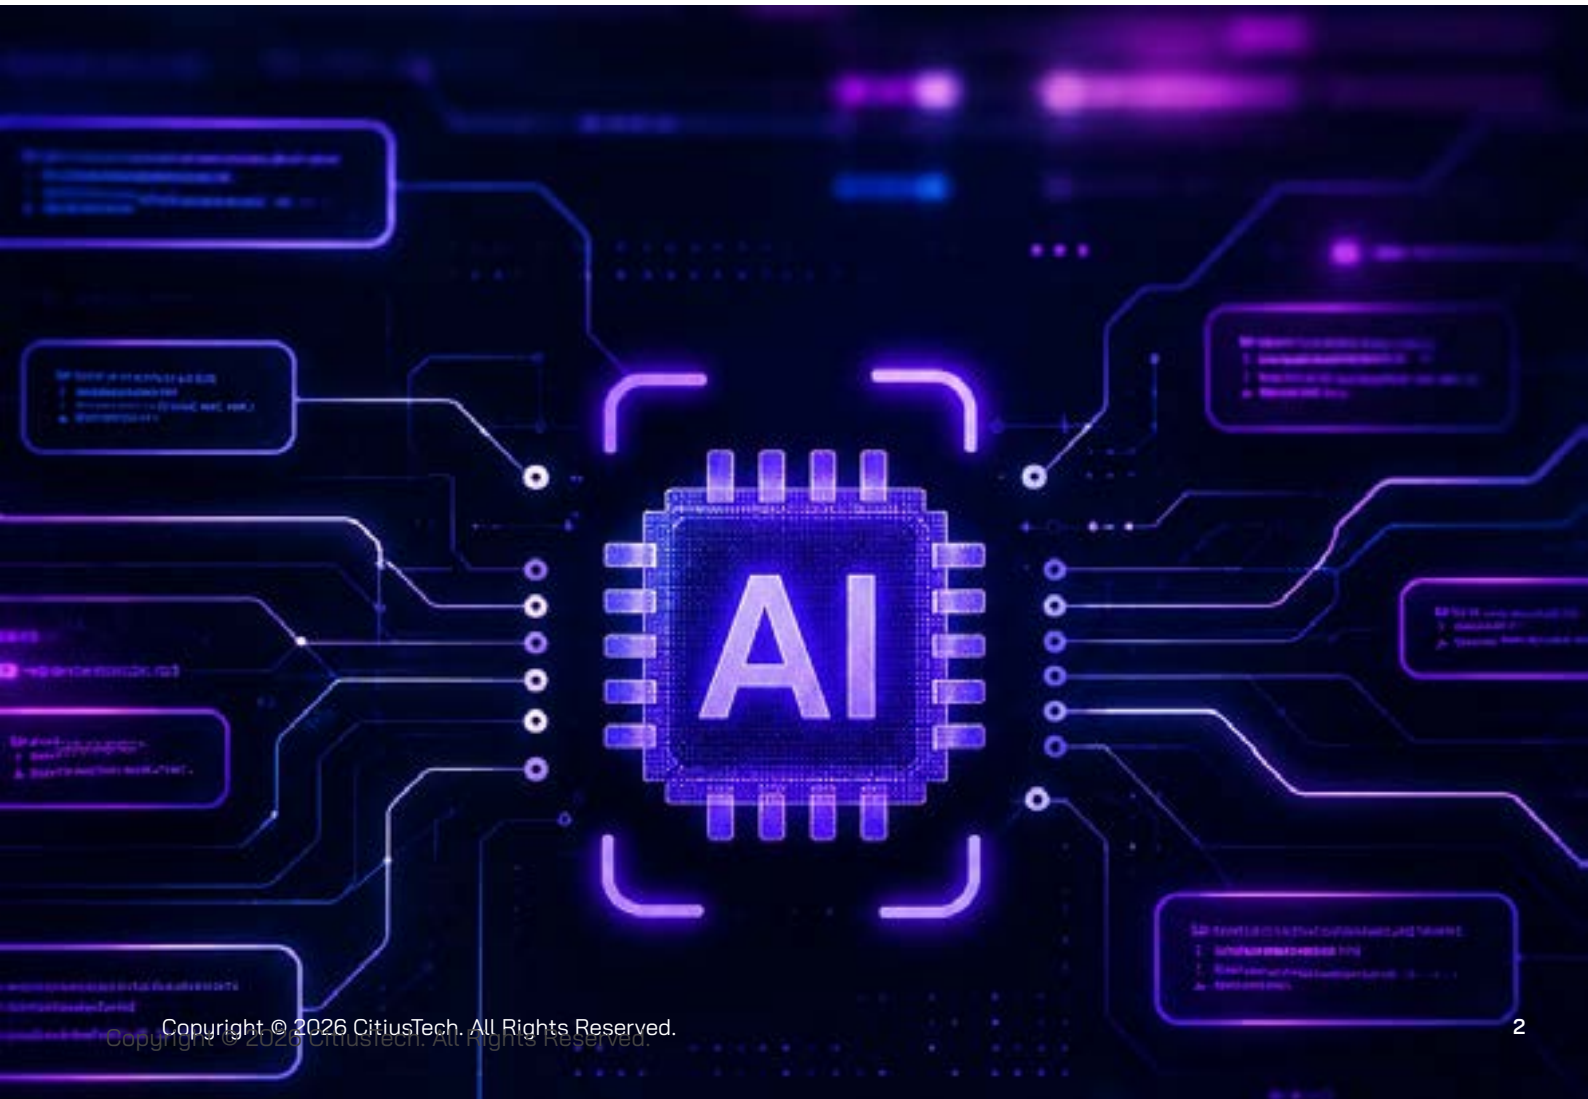

Copyright © 2026 CitiusTech. All Rights Reserved.

Executive Summary

Patient hubs already depend on AI to keep access to therapy moving at pace. The question is no longer whether AI is used, but whether it is governed effectively. Leaders must ensure that enrolment decisions remain unbiased, document capture is accurate, prioritization models are reliable, and patient data is secure. When these controls fail, the consequences such as delays, errors, and erosion of trust, are immediate.

This white paper outlines a practical, end-to-end approach to getting these fundamentals right, enabling hubs to improve speed to therapy without increasing rework or inequity.

Responsible AI is not a theoretical ideal; it is an operational discipline that directly shapes daily hub performance.

Why this matters for hubs

Patient hubs are often the first step in a patient's journey to access therapy, and AI is increasingly embedded across this journey—driving enrolment, capturing documentation, and prioritizing cases. When designed and governed well, these capabilities accelerate access, streamline operations, and reduce manual rework.

When they are not, the impact is immediate and significant. Delays can lead to therapy abandonment, errors can trigger denials and appeals, and inconsistent decisions can erode provider and patient trust. These breakdowns directly affect core hub KPIs—cycle time, rework rates, and equity of access.

Responsible AI allows hubs to move with speed while maintaining fairness, accuracy, and trust—ensuring that operational efficiency does not come at the cost of patient outcomes.

AI in modern patient hubs: Opportunity with obligation

Patient hubs bring together patients, prescribers, payers, and manufacturers, and the decisions made here directly determine whether therapy is approved, how quickly it starts, and who gets access. AI can help hubs move faster and scale consistently, but errors surface immediately at the patient level.

For example:

A misread benefits document or biased enrollment rule can incorrectly flag a patient as ineligible, delaying approval and causing treatment to stall while appeals are filed.

A faulty prioritization model can push urgent cases to the back of the queue, increasing abandonment risk.

These are not abstract failures—they translate into delays, rework, and loss of confidence from both patients and providers.

That's why responsible design in hub AI isn't optional; it is essential to protecting patients while keeping access to therapy moving.

Core AI-Enabled hub workflows

Enrollment AI

What the AI does

Enrollment AI determines patient eligibility, routes cases for insurance verification, and prioritizes which cases should be handled first. These systems typically rely on historical operational and insurance data, which may reflect access limitations rather than true medical need.

Risk example

The system uses prior treatment adherence to estimate the likelihood of insurance approval. Patients from underserved areas often have fewer pharmacy options, which affects their treatment history. As a result, these patients are deprioritized despite being equally medically eligible.

Primary risks

- Hidden bias driven by location, insurance type, or historical utilization
- Limited explainability of eligibility and prioritization decisions
- Excessive automation without sufficient human review or escalation

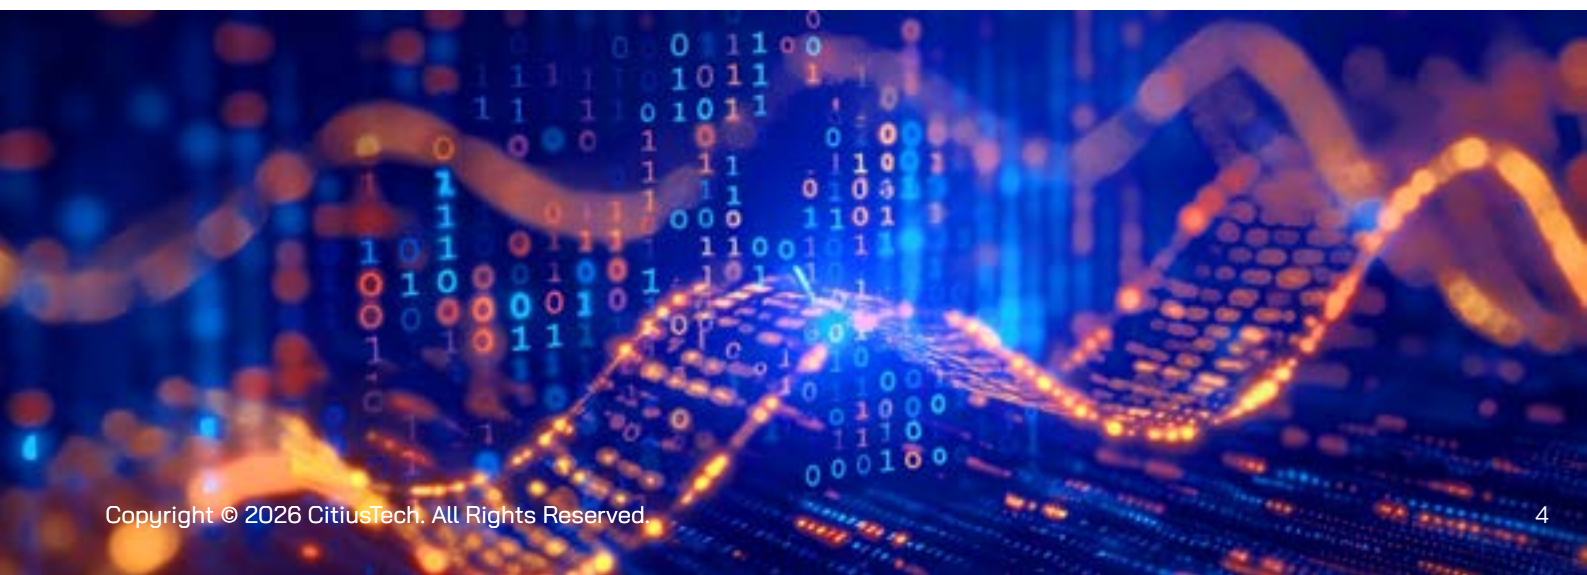

OCR-Driven Document Capture

What the AI does

OCR systems convert faxed, scanned, and handwritten documents into structured digital data. This accelerates processing but allows small recognition errors to propagate silently through downstream workflows

Risk example

A handwritten dosage on a prior authorization form is misread by the OCR engine. The incorrect value is carried forward into insurance verification, resulting in a denial and delaying patient treatment

Primary risks

- Inconsistent accuracy across fields and document types
- No systematic flagging of low-confidence extractions for review
- Increased exposure of sensitive patient data across systems

OCR-Driven Document Capture

What the AI does

Predictive triage tools help patient hubs determine which cases require immediate attention, based on urgency, clinical risk, or expected benefit from intervention.

Risk example

The model places too much weight on prior healthcare utilization. Newly diagnosed patients with high medical risk but limited history are ranked lower, delaying timely intervention.

Primary risks

- Model performance degradation as data patterns change
- Reinforcement or amplification of existing inequities
- Limited transparency for frontline teams into prioritization logic

Responsible AI risk categories

Across hub workflows, Responsible AI risks typically fall into four distinct categories. While they often interact, separating them helps teams identify root causes and apply the right controls.

1. Bias and fairness

What can go wrong

- Unequal outcomes driven by proxy features such as geography, payer type, or historical utilization
- Reinforcement of historical access gaps through skewed training data

Why it matters

Bias and fairness risks most often originate in the data used to make enrollment, eligibility, or prioritization decisions. Models may rely on seemingly neutral signals—like past healthcare usage or insurance history—that correlate with access rather than medical need. Over time, this can lead to patients with similar clinical profiles experiencing different outcomes, simply because the system learned from historically uneven patterns.

2. Data accuracy and integrity

What can go wrong

- OCR misreads, transformation errors, or inconsistent field-level accuracy
- Error propagation across downstream workflows

Why it matters

Accuracy and integrity risks are most visible in OCR-driven processes. While digitization accelerates intake, even small extraction errors—such as an incorrect dosage or date—can silently flow into benefits verification, enrollment, or triage. Without confidence scoring or validation checks, these issues often surface only after a denial, rework cycle, or treatment delay has already affected the patient.

3. Explainability and transparency

What can go wrong

- AI-driven decisions that cannot be clearly interpreted or justified
- Limited visibility for operations, compliance, or appeals teams

Why it matters

When AI systems cannot explain why a case was deprioritized or denied, frontline and oversight teams struggle to respond to patient questions, resolve appeals, or support audits. Lack of transparency erodes trust and increases operational friction, even when the underlying decision may be technically correct.

4. Security and privacy

What can go wrong

- Concentration of sensitive PHI across AI pipelines
- Increased exposure to cyber and privacy risks

Why it matters

Hub workflows often require AI systems to access and exchange large volumes of sensitive patient data. Without strong access controls, monitoring, and governance, this concentration of PHI increases the impact of security failures and privacy breaches, creating both regulatory and reputational risk.

These risks must be addressed **systematically, not reactively,** through governance, monitoring, and human-in-the-loop controls that are built into hub workflows from the start—so speed to therapy improves without compromising fairness, accuracy, or patient trust.

Governance framework for responsible AI

Primary objective

A well-designed Responsible AI governance framework in healthcare ensures that AI solutions improve patient outcomes without compromising clinical safety, ethics, privacy, or regulatory compliance. At its core, governance provides clear accountability for how AI is selected, deployed, and monitored in real-world clinical settings.

Healthcare organizations typically achieve this through formal oversight structures that review AI use cases for clinical appropriateness, data protection, and fairness. For example, CommonSpirit Health's Ethics Data Advisory Group (EDAG) is often cited as an illustrative model of how multidisciplinary review bodies can evaluate AI solutions before and after deployment. Such examples are not prescriptive, but they demonstrate how governance can be operationalized in practice.

Effective governance frameworks share a few common characteristics:

- Clear executive ownership and decision rights for AI initiatives
- Involvement of diverse stakeholders, including clinicians, compliance teams, and patient representatives
- Ongoing monitoring, audits, and escalation paths for identified risks

Together, these elements help organizations reduce bias, increase transparency, and build sustained trust in AI-enabled healthcare systems.

Healthcare-specific risk considerations

- Patient harm caused by inaccurate, incomplete, or inappropriate AI recommendations
- Algorithmic bias affecting clinical decisions across race, gender, age, or socioeconomic groups
- Exposure or misuse of protected health information (HIPAA/GDPR risks)
- Limited explainability, reducing clinician confidence and adoption
- Non-adherence to healthcare regulations and standards (e.g., FDA guidance, Clinical Decision Support rules)

Responsible AI principles

Domain	Reinterpreted Guideline
Clinical Harm Prevention	AI should support (not override) medical decision-making by qualified clinicians
Medical Soundness	Models must be grounded in validated clinical evidence and accepted medical knowledge
Decision Transparency	Care providers must be able to interpret and justify AI-driven insights
Data Stewardship	Patient data must be safeguarded with explicit consent visibility and lifecycle controls
Equity by Design	AI outcomes should be consistently fair across diverse patient demographics
Human Accountability	Clear clinician ownership is required for all AI-influenced care decisions
Legal & Policy Alignment	AI use must conform to applicable healthcare and AI regulations globally
Decision Traceability	All AI-assisted clinical actions must be reviewable and auditable end-to-end

Healthcare AI Governance Structure

Clinical AI Steering Committee	Clinical AI Ethics & Risk Board	AI Product & Model Owners
CMIO/CMO	Clinicians	Clinical Accuracy
CIO/CTO	Biostats	Model Behavior
Safety Officer	PV	Post-Market Monitoring
Compliance and Legal Team	Patient Safety	
Data Protection Officer		
QA / Validation Head		

Reference Architecture with Embedded Controls

The Reference Architecture with Embedded Controls illustrates a layered approach to deploying AI in healthcare with safety, compliance, and clinical accountability built in by design. At the foundation, the AI Governance Layer establishes regulatory and ethical oversight by aligning AI usage with FDA guidance, Clinical Decision Support (CDS) policies, HIPAA requirements, and enterprise ethics standards. Centralized governance, supported by audit logs and policy enforcement, ensures transparency and traceability across all AI initiatives.

The Clinical AI Control & Assurance Layer translates governance into operational safeguards for patient-impacting use cases. This layer manages clinical use case registration and risk tiering, enforces PHI governance and consent management, and requires robust model documentation supported by clinical evidence. Validation, QA, and GxP-aligned controls are applied alongside human-in-the-loop mechanisms, ensuring clinicians retain authority to review and challenge AI outputs. Continuous post-market surveillance monitors real-world performance, bias, and safety signals to maintain clinical reliability over time.

Supporting these controls, the Health Data & AI Platform Layer integrates EHR/EMR systems, claims data, real-world evidence, ML models, LLMs, feature stores, and MLOps pipelines to enable scalable AI development within controlled environments. At the point of care, clinicians, operations teams, and patients interact with AI through decision-support experiences that emphasize review and override rather than automation. Together, the architecture ensures AI augments clinical decision-making while protecting patient safety, maintaining regulatory compliance, and building trust across the healthcare ecosystem.

Reference Architecture with Embedded Controls

Human oversight and operational integration

Human oversight is how hubs keep speed-to-therapy high without letting AI create silent errors. AI should assist decisions, not replace accountability. In patient hubs, every AI output must be reviewable, explainable at a basic level (what/why/confidence), and easy to stop or override when something looks off.

There are different ways this works:

Human-in-the-loop: A person reviews or approves AI results before action is taken. This is used for high-risk decisions.

Human-on-the-loop: AI works on its own, but a person watches it and can step in if needed.

Human-out-of-the-loop: AI runs without regular human involvement. This is not suitable for important healthcare decisions.

In healthcare, AI systems must be designed so users clearly understand what the AI can and cannot do. Staff should be able to spot unusual results, avoid blindly trusting AI output, and stop or override the system when needed. This requires clear screens, explanations, alerts, and simple controls such as a "stop" option.

Governance and roles

Make ownership explicit:

- assign a single accountable owner for each AI workflow **(Enrollment, OCR, Triage)** responsible for performance, fairness checks, incident response, and approving model/threshold changes.

Role-based oversight (what gets reviewed, what escalates, who can override):

Decisions that must have human review

- **Enrollment:** exceptions, missing/conflicting info, and any AI-driven deprioritization that could delay access steps.
- **OCR:** critical fields when confidence is low (patient identifiers, payer, drug/dose, dates, signatures).
- **Triage:** "defer" recommendations and any change to urgency/SLA class.

Escalation triggers (auto-route to a human queue)

- Confidence below threshold for a critical field, or mismatch vs the hub record.
- Sudden increase in overrides/rework, performance drop, or fairness variance beyond tolerance.
- Any privacy/security anomaly involving PHI or audit trail gaps.

Override ownership (no ad hoc overrides)

- **Hub Ops Lead / Supervisor:** case-level overrides and queue decisions; ensures escalations meet SLA.
- **Clinical Safety / Medical Oversight:** overrides tied to patient-risk or urgency logic.
- **Compliance & Privacy Lead:** policy exceptions, audit readiness, and PHI handling controls.
- **AI Model Owner:** approves threshold/model changes after review; owns monitoring and drift response.

Hub examples

- **OCR:** if critical-field confidence is <90%, route to "Verify" queue; if <80%, block downstream submission until corrected.
- **Triage:** if the model suggests "standard/defer" but urgency markers exist (e.g., repeated denials or time-sensitive therapy window), auto-escalate to the clinical/case manager queue.

Implementation checklist

- **Governance setup:** Set up an AI governance group with clear ownership.
- **Policy & SOPs:** Define policies and procedures for AI use.
- **Human oversight design:** Decide how humans will review or override AI outputs.
- **Escalation pathways:** Establish clear incident reporting and response steps.
- **Training & competency:** Train staff and confirm competency.
  Adoption notes: Responsible AI only works if frontline teams are trained to challenge outputs, use override paths, and report anomalies without fear—treat this as change management, not just system deployment**.**
- **Audit trails & logging:** Log all AI inputs, outputs, and decisions.
- **Validation & testing:** Test and validate AI before use.
- **Change Management:** Control and document all changes.
- **Monitoring & maintenance:** Monitor performance and risks regularly
- **Privacy & security:** Always ensure data privacy and security.

Metrics, KPIs, and evidence of responsible AI

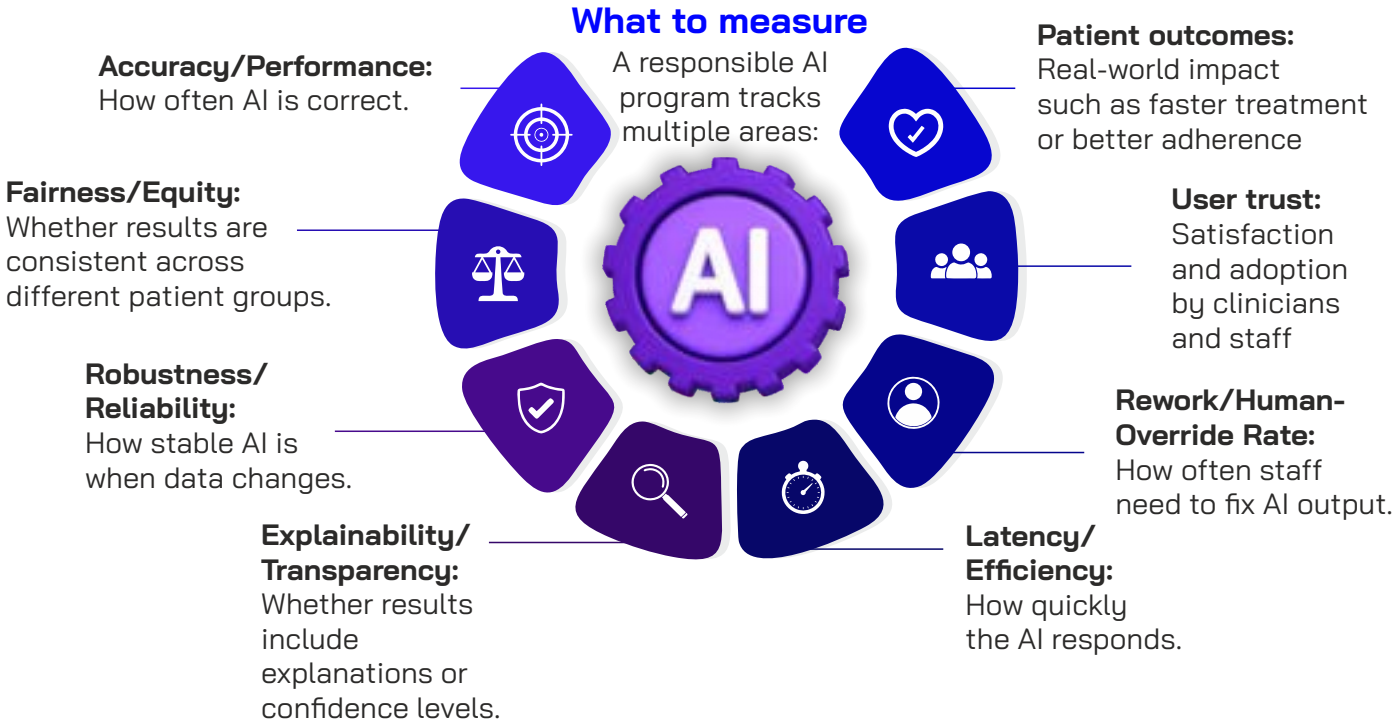

Key metrics for patient hubs

Leader focus: time to therapy (median), manual rework/override rate, outcome equity ratio, and document processing accuracy. These four signals usually have a surface access impact, safety risk, and operational drag early—before you scale.

Common metrics may include:

KPI	What it measures	How it is calculated	Primary data source	Target/ Benchmark	Review frequency
Document processing accuracy (%)	How accurately documents are read and processed by the system	Correct documents ÷ total documents	QA audits, system logs	⩾95% accuracy	Daily Weekly
Patient enrollment success rate (%)	Percentage of referred patients successfully enrolled	Enrolled patients ÷ total referrals	Hub enrollment system	Program-specific; benchmark vs manual process	Weekly/ Monthly

KPI	What it measures	How it is calculated	Primary data source	Target/ Benchmark	Review frequency
Outcome equity ratio	Consistency of outcomes across patient groups	Lowest group rate ÷ highest group rate	Enrollment data with demographics	Close to 1.0 (minimal variation)	Monthly/ Quarterly
Time to therapy (days)	Time from referral or enrollment to therapy start	Therapy start date – referral date (median)	Scheduling systems, patient records	Defined per therapy/program	Monthly/ Quarterly
System response time (seconds)	Speed of system processing per request	Average processing time per transaction	Application and system logs	Near real-time for operational workflows	Daily
System availability (%)	Percentage of time the system is operational	Uptime ÷ scheduled time	IT monitoring tools	⩾99.9% availability	Monthly
Manual rework rate (%)	How often staff must correct AI output	Reworked cases ÷ total cases	QA logs, error tracking	⩽5%	Weekly/ Monthly
Clinician satisfaction score	Clinician trust and usability perception	Average survey score	User satisfaction surveys	⩾80% satisfaction	Quarterly

Targets should be realistic and reviewed regularly.

Monitoring and reporting

Data should be collected automatically from system logs and user feedback. Reviews should be held at various levels:

- Daily system checks
- Weekly team reviews
- Monthly performance and fairness reviews
- Quarterly governance reviews

Alerts should trigger when performance drops or risks increase.

Evidence and compliance

Keep clear records to show responsible AI use, including:

- Logs of AI decisions
- Test and validation results
- Training records
- Incident reports
- Safety reviews and approvals

Consistent documentation, regular audits, and transparent reporting help demonstrate that AI systems are safe, controlled, and ready for review.

Cybersecurity and data protection

AI-enabled hubs expand the attack surface by concentrating sensitive PHI across OCR pipelines, model APIs, and feedback loops. That increases both the likelihood and blast radius of a breach, so security controls must be embedded by design, not bolted on later.

Minimum controls that make a measurable difference in hubs:

Access control (prevent):

enforce least-privilege role-based access, MFA for privileged users, and strong segregation between environments (dev/test/prod) for any PHI-handling components.

Monitoring (detect):

continuous logging and alerting for PHI access, unusual download patterns, API abuse, and model endpoint anomalies wired into SIEM with audit-ready retention.

Containment (limit impact):

data minimization (only required fields), encryption in transit/at rest, and “break-glass” procedures (kill switch / quarantine queue) to stop automation when suspicious activity or data leakage is detected.

Practically, Responsible AI strengthens security by improving traceability (what data was used, when, and why), constraining automation so issues don't propagate silently, and speeding investigations with complete logs of AI inputs, outputs, and human actions.

The outcome is simple: faster patient access and stronger patient trust because the hub can prove decisions were controlled, data was protected, and operations remain audit-ready under real-world threats.

Implementation roadmap

AI-enabled hubs expand the attack surface by concentrating sensitive PHI across OCR pipelines, model APIs, and feedback loops. That increases both the likelihood and blast radius of a breach, so security controls must be embedded by design, not bolted on later.

Minimum controls that make a measurable difference in hubs:

Phase 1: Foundation	Phase 2: Controlled deployment	Phase 3: Scale and optimize
This focuses on building a sturdy foundation. This includes setting clear rules for how AI will be governed, identifying where risks may exist, and establishing baseline metrics to understand current performance. These basics ensure AI is introduced in a controlled, measurable, and trustworthy way from the start.	This focuses on rolling out AI in a safe and controlled way. Confidence thresholds ensure the system only acts when it is sufficiently sure, while humans step in for uncertain or high-impact cases. Fairness monitoring helps confirm that decisions remain consistent and do not disadvantage certain patient groups as usage scales.	This focuses on running AI safely and continuously improving it. The system is regularly monitored to ensure it performs as expected, and any model updates are made only after proper governance review. Executive dashboards provide leaders with clear visibility into performance, risks, and outcomes over time.
Governance setup	**Confidence thresholds**	**Continuous monitoring**
Risk assessment	**Human-in-the-loop workflows**	**Model updates under governance approval**
Baseline metrics	**Fairness monitoring**	**Executive dashboards**

Conclusion

Responsible AI doesn't slow innovation; it makes safe growth possible. By building fairness, clarity, and security into hub AI workflows, organizations can move patients to therapy faster while keeping trust intact. The most dangerous hub AI is not the one that fails loudly, but the one that quietly scales without accountability.

CitiusTech

Shaping Healthcare Possibilities

About CitiusTech

CitiusTech is a global technology services, consulting, and business solutions enterprise 100% focused on the healthcare and life sciences industry. We enable 140+ enterprises to build a human-first ecosystem that is efficient, effective, and equitable. Leveraging deep domain expertise and next-generation technologies including AI, Cloud, Data, and Intelligent Automation, we assist our clients in realizing their vision, accelerate transformation, and achieve business outcomes. With 7,700+ healthcare technology professionals worldwide, CitiusTech powers digital innovation, business transformation, and industry-wide convergence through next-generation technologies, solutions, and products.Follow CitiusTech on Twitter or LinkedIn.

www.citiustech.com