Healthcare Leadership Whitepaper

AI Fluency Framework

March 9, 2026 | Jeff Draper, CHTS, Chief Human Capital & Digital Strategy Officer (CHCDSO)

Executive Summary

Healthcare leaders face a fast-moving AI landscape shaped by clinical risk, reimbursement pressures, workforce shortages, and changing regulations. Yet many health systems still treat AI as a string of pilots rather than a strategic capability. **AI fluency** is the missing infrastructure: the shared leadership and operator competence to govern, deploy, and scale AI safely, measurably, and at pace.

This executive whitepaper defines AI fluency for the C-suite and operational leaders, synthesizes leading frameworks, and presents a practical **AI Fluency Framework for Healthcare** spanning five domains—**Governance & Responsible AI; Operator & Workflow Literacy; Technical & Data Foundations; Change Readiness & Org Design; Cross-Functional Collaboration**—with a maturity model, role architecture, KPIs, and an implementation roadmap.

Regulatory momentum underscores the need. The **ONC HTI-1** final rule establishes first-of-its-kind **algorithm transparency** requirements for decision support in certified health IT, raising the bar for oversight and documentation. The **HHS OCR Section 1557** nondiscrimination rule now puts legal responsibility on covered entities to **identify and mitigate** discriminatory impacts when using patient care decision-support tools, including AI. The **FDA** advances a **total product lifecycle** approach for AI-enabled devices and codifies **Good Machine Learning Practice (GMLP)** expectations, which aligns AI governance with post-market monitoring and change-control. The **NIST AI Risk Management Framework (AI RMF)** provides a recognized structure to govern AI risk across the lifecycle (Govern, Map, Measure, Manage). [healthit.gov], [natlawreview.com] [nycdentalsociety.org], [jamanetwork.com] [fda.gov], [fda.gov] [nist.gov]

Bottom line for executives: Treat AI fluency as strategic infrastructure—on par with EHRs and cybersecurity—to de-risk adoption, accelerate value capture, and maintain public trust.

Table of Contents

The AI Fluency Gap in Healthcare

Common failure modes persist across health systems:

- **Pilot proliferation without workflow ownership.** Tools deploy but accountability for outcomes is diffuse; clinicians become skeptical when error handling and escalation paths are unclear. [accu-health.com]
- **Underpowered governance.** Few systems maintain an end-to-end view of algorithm risk, change control, cybersecurity, and human-in-the-loop expectations across service lines. [nist.gov], [fda.gov]
- **Insufficient transparency and bias controls.** HTI-1 and Section 1557 now require visibility and mitigation for algorithmic risks—expect audits of documentation, training data, and impacts on protected classes. [healthit.gov], [nycdentalsociety.org]
- **Misaligned incentives and measurement.** "Go-live" is counted as success rather than **clinically meaningful outcomes, safety, equity, and adoption durability**. [accu-health.com]
- **Fragmented data and technical foundations.** Lifecycle governance (model updates, drift, provenance) is often absent outside regulated SaMD, but regulators increasingly expect it. [fda.gov], [fda.gov]

Why this matters now

- **Regulatory velocity:** Transparency (ONC HTI-1), nondiscrimination (HHS OCR Section 1557), and lifecycle oversight (FDA) are converging. Executives must align operating models accordingly. [healthit.gov], [nycdentalsociety.org], [fda.gov]
- **Trust and safety:** Without explicit human oversight and documented controls, AI erodes clinician and patient trust. Leading labs and industry commitments emphasize red-teaming, evaluations, and safeguards. [openai.com]
- **Competitive advantage:** Systems that embed fluency can scale safely and translate innovation into throughput, quality, and workforce relief. The inverse is capital loss and reputational risk. [nist.gov]

Defining AI Fluency for Healthcare Leadership

Working definition: *AI fluency is the executive and operator capability to plan, govern, and integrate AI into clinical and business workflows so that outcomes are safer, more equitable, and measurably better—under transparent, lifecycle-aware controls.* This blends strategic governance (boards, committees, risk appetites) with **operator-level competencies** (task delegation to AI, prompt quality, critical evaluation, diligence in responsible use). [anthropic.com], [anthropic....illjar.com]

Executive fluency includes: setting governance structures, defining accountability and escalation, aligning to NIST AI RMF functions (Govern/Map/Measure/Manage), and resourcing lifecycle monitoring. [nist.gov]

Operational fluency includes: knowing **what to delegate to AI**, how to **describe** tasks effectively, how to **discern** and verify outputs, and how to apply **diligence** (privacy, safety, equity, documentation). These "**4Ds**" from the Anthropic-supported AI Fluency work are transferable across models and tasks. [anthropic.com], [anthropic....illjar.com]

Synthesis of Existing Frameworks (High-Level Comparison)

The table below distills concepts from authoritative sources to inform a healthcare-specific approach. *(Open Sans font)*

Source / Perspective	Core Idea	Relevance to Healthcare Leaders
NIST AI RMF (1.0 & GenAI Profile)	Governance and lifecycle risk management via **Govern/Map/Measure/Manage**; trustworthiness characteristics (safety, explainability, privacy, fairness). [nist.gov]	Provides the enterprise scaffolding to map risks and mitigations across service lines and vendors; aligns with board-level oversight.
ONC HTI-1 Final Rule	Algorithm transparency for **Decision Support Interventions** in certified health IT; shifts CDS to DSI with baseline information requirements. [healthit.gov], [natlawreview.com]	Drives procurement and integration requirements; necessitates inventory, disclosures, and clinician-facing transparency content.
HHS OCR Section 1557	Requires **identification and mitigation** of discriminatory impacts when using patient care	Makes AI fairness and impact analysis a

	decision-support tools. [nycdentalsociety.org], [jamanetwork.com]	compliance obligation; requires cross-functional monitoring and escalation.
FDA (SaMD, GMLP)	Lifecycle mindset; **predetermined change control plans**, post-market surveillance, bias and data quality expectations. [fda.gov]	Even outside SaMD, these principles guide validation, versioning, and monitoring of AI in clinical contexts.
Anthropic AI Fluency (4Ds)	Operator competencies: **Delegation, Description, Discernment, Diligence** across automation/augmentation/agency. [anthropic.com], [anthropic....illjar.com]	Provides practical behaviors for clinicians and staff to use AI safely and effectively.
Healthcare operator lens (Accu-Health)	AI adoption succeeds when **workflow ownership, operator roles, and measurement** are designed as infrastructure. [accu-health.com]	Anchors "people + process" as the adoption engine, not just tool rollout.

The Proposed AI Fluency Framework for Healthcare

Goal: A leadership and operator blueprint that enables **responsible, scalable adoption** with measurable outcomes.

1) Governance & Responsible AI

- **Board-aligned AI risk appetite** with a systemwide charter that maps to NIST AI RMF and references clinical safety, equity, privacy, cybersecurity, and transparency. [nist.gov]
- **Algorithm Governance Committee** including CMIO/CNIO, CISO, Compliance/Privacy, Quality & Safety, Health Equity, Legal/Regulatory, and frontline representatives; responsibilities: inventory, approvals, post-market monitoring, incident response, and public transparency artifacts per HTI-1 where applicable. [healthit.gov]
- **Responsible AI policy** defining human-in-the-loop, explainability thresholds, red-teaming/evaluations, and vendor obligations aligned to leading industry commitments. [openai.com]
- **Section 1557 controls** for bias identification and mitigation across patient-facing and back-office tools (triage, CDI, collections), with documentation of reasonable efforts and corrective actions. [nycdentalsociety.org]

2) Operator & Workflow-Level Literacy

- **Role-specific pathways** to build the "4Ds" competencies (Delegation, Description, Discernment, Diligence) for clinicians, revenue cycle, contact center, supply chain, and analytics teams. [anthropic.com]
- **Standard work** for verification and escalation (e.g., when LLM-generated note drafts or patient messages must be edited/approved by licensed professionals), in line with leading lab usage policies. [bakerdonelson.com]
- **Prompting playbooks** embedded in workflows (EHR inbox, care coordination, prior auth) with guardrails for PHI handling and patient communications consistent with HIPAA guidance. [hhs.gov]

3) Technical & Data Foundations

- **AI system inventory** tied to data lineage: training/evaluation datasets, provenance, and representativeness; model versioning; deployment environment; monitoring metrics; and PCCP-style update plans for high-impact tools. [fda.gov]
- **Risk-based architecture:** input/output guardrails, retrieval with provenance, drift detection, adversarial testing, and integrated cybersecurity controls. [aws.amazon.com]
- **Privacy & PHI controls:** strict vendor due diligence, BAA rigor, and continuous compliance, recognizing AI-specific HIPAA pitfalls (e.g., de-identification, training data reuse, tracking tech). [hhs.gov], [hipaajournal.com]

4) Change Readiness & Organizational Design

- **Named workflow owners** for every deployment; change budgets cover training-to-competency and process re-design—not one-time training. [accu-health.com]
- **Clinical communication plan** to actively manage trust (benefits, limitations, escalation channels), aligned with HTI-1 transparency content. [healthit.gov]
- **Talent model** for "AI-fluent operators": hybrid roles at the intersection of informatics, operations, and change management who translate strategy into sustained workflows. [accu-health.com]

5) Cross-Functional Collaboration

- **Working groups** connecting Quality & Safety, Health Equity, Data/IT, Clinical Operations, and Finance to align on KPIs and benefits realization.
- **AIM (Allocate-Iterate-Mediate) collaboration lens** to clarify work division, iterative improvement loops, and human mediation/oversight. [aifluency.org]

Governance & Safety Infrastructure

Decision-making and risk management

1. **Tiered risk classification** (clinical impact, patient contact, automation level, data sensitivity). Tie review rigor to tier. Align to NIST AI RMF "Map/Measure" and FDA lifecycle where relevant. [nist.gov], [fda.gov]
2. **Approval gates**: business case → data readiness → bias & safety assessment → usability & workflow validation → pilot with pre-defined success/fail criteria → scale. (HTI-1 artifacts embedded when tool is part of certified health IT.) [healthit.gov]
3. **Post-market monitoring**: real-world performance, drift, user feedback, safety events, and equity impacts; triggers for rollback or model updates; public-facing transparency summaries where appropriate. [fda.gov]

Oversight bodies

- **AI Executive Council (quarterly):** sets strategy and risk appetite, reviews portfolio value and risk.
- **Algorithm Governance Committee (monthly):** approves deployments, reviews incidents and fairness metrics, ensures Section 1557 compliance. [nycdentalsociety.org]
- **Clinical & Operator Communities of Practice (biweekly):** share playbooks, escalate pain points, iterate prompts and standard work.

Risk controls to institutionalize

- **Red-teaming & evaluations** (internal + external SMEs) for safety, misuse, and societal risks consistent with leading lab commitments. [openai.com]
- **Transparency packets** for clinicians and, where appropriate, patients: intended use, limitations, data sources, performance bounds, update cadence—aligned with HTI-1 DSI. [healthit.gov]
- **Equity assessments** and mitigation plans required for deployments implicating Section 1557, with documented "reasonable efforts." [nycdentalsociety.org]

Operator Roles & Change Infrastructure

Role design

- **AI-Fluent Operator (clinical or operational):** Owns workflow integration, standard work, verification steps, and feedback loops; acts as a translator between frontline teams and data/IT. [accu-health.com]
- **Clinical Prompt Steward:** Maintains task libraries and safe prompting patterns for specific specialties; coordinates with privacy and quality. [anthropic.com]
- **Model Owner (IT/Data):** Maintains versioning, telemetry, drift and bias monitoring; coordinates PCCP-style updates and change notices. [fda.gov]

Training pathways

- **Foundations for all leaders and staff:** AI basics, risk, verification, and escalation aligned to the **4Ds**. [anthropic.com]
- **Role-specific enablement:** e.g., ambulatory clinicians (documentation drafting + safety checks), inpatient nurses (handoff summaries + exception handling), schedulers (triage prompts + privacy), revenue cycle (coding assistance + audit).
- **Competency-based validation:** observed skills in workflow, not quiz completion—mirroring workforce-first adoption principles. [accu-health.com]

Workflow integration

- **Design for handoffs and exception paths**; automation bias countermeasures: require second checks for high-impact tasks, and measure override appropriateness. [hai.stanford.edu]
- **Labeling and disclosure** when AI assists in patient communications as required by law and platform policies; ensure licensed review where needed. [bakerdonelson.com]

Measurement & Maturity Model

What to measure: Safety, quality, equity, experience, and value—not just usage. [accu-health.com]

Sample KPI Set

- **Safety & quality:** error rates detected in verification; time-to-escalation; revert/rollback frequency; post-deployment adverse event correlation. [fda.gov]
- **Equity:** disparity deltas in model performance and outcomes (e.g., sensitivity/specificity by demographic group) with documented mitigations per Section 1557. [nycdentalsociety.org]

- **Experience:** clinician time saved (documentation, inbox); Net Trust Score for AI tools; patient comprehension of AI-assisted communications. [healthit.gov]
- **Value:** throughput, length-of-stay improvements, denial reductions, cost-to-serve changes.

AI Fluency Maturity Model (simple)

(Open Sans font)

Level	Characteristics	Executive Actions
Level 1 — Ad hoc	Pilots scattered; no inventory; limited training; no formal bias/equity review.	Establish AI Executive Council; create inventory; freeze high-risk pilots pending governance. [nist.gov]
Level 2 — Emerging	Basic committee; some training; initial transparency docs; reactive incident handling.	Adopt NIST AI RMF mapping; publish DSI transparency packets; stand up Section 1557 mitigation processes. [healthit.gov], [nycdentalsociety.org]
Level 3 — Managed	Tiered risk reviews; operator roles defined; post-market monitoring; equity KPIs.	Implement lifecycle monitoring and PCCP-style updates; expand Communities of Practice. [fda.gov]
Level 4 — Integrated	Portfolio-level value tracking; trust communications; cross-functional cadence; external red-teaming.	Tie funding to KPI attainment; publish annual transparency report. [openai.com]
Level 5 — Scaled	AI fluency embedded in job architecture; rapid, safe scale across sites; continuous improvement loops.	Extend to partner network; codify vendor requirements; benchmark externally. [nist.gov]

Capability Matrix (Illustrative)

(Open Sans font)

Domain	Starter (Quarter 1)	Build (Quarters 2–3)	Scale (Quarters 4–6)
Governance	Inventory; charter; risk tiers	DSI transparency packets; Section 1557 processes	Annual transparency report; external red-team
Operators	Foundations + 4Ds; role mapping	Competency checks in workflow; prompt libraries	Recertification; cross-site playbooks
Tech/Data	Basic telemetry; access controls; PHI guardrails	Drift & bias dashboards; PCCP-style plan	Automated alerts; vendor SLAs for updates

Change	Named workflow owners; change budget	Clinical comms + patient disclosures	Embedded improvement loops and KPIs
Collaboration	Equity + Safety + IT triage huddles	Communities of Practice active	External benchmarking and learning network

[anthropic.com], [nist.gov], [healthit.gov], [nycdentalsociety.org], [fda.gov]

Implementation Roadmap for Executives (Phased Approach)

Phase 0 (Weeks 0–2): Set the intent

- Appoint an **executive sponsor** (COO/CMIO), establish the **AI Executive Council**, and confirm risk appetite. [nist.gov]
- Issue an **adoption pause** for high-risk pilots lacking governance; require rapid inventory submission.

Phase 1 (Weeks 2–8): Baseline & guardrails

- **Inventory all AI systems** (internal and vendor): use, data sources, PHI, clinical impact, automation level; flag DSI within certified health IT for HTI-1 transparency work. [healthit.gov]
- Stand up **Algorithm Governance Committee**, adopt NIST AI RMF mapping, and publish **Responsible AI Policy** including human-in-the-loop requirements and incident handling. [nist.gov]
- Launch **Section 1557 risk process**: identify tools with potential disparate impact; define mitigation steps and documentation. [nycdentalsociety.org]
- Initiate **operator foundations training** (4Ds) for priority workflows. [anthropic.com]

Phase 2 (Months 3–6): Pilot with proof

- Select **3–5 high-value workflows** (e.g., documentation support, care navigation messaging, prior auth summarization).
- For each: define **success metrics** (safety/equity/value), build **transparency packets**, deploy **verification paths**, and run **red-team/evals** before scale. [openai.com], [healthit.gov]
- Implement **monitoring dashboards** for drift and fairness; create **rollback plans**. [fda.gov]

Phase 3 (Months 6–12): Standardize & scale

- Convert successful pilots into **standard work** with change budgets and operator recertification cycles. [accu-health.com]

- Formalize **PCCP-style update plans** for high-impact tools; tighten vendor SLAs to include transparency, audits, and model change notices. [fda.gov]
- Publish the first **AI Transparency & Impact Summary** to clinicians and governance bodies; include DSI/HTI-1 content where applicable. [healthit.gov]

Phase 4 (Year 2): Institutionalization

- Embed AI fluency into **job descriptions, credentialing/OPPE** where appropriate, and leadership development.
- Join **external benchmarking networks**; periodically engage independent red-team reviews. [openai.com]

Frequently Asked Executive Questions

How does this align with HIPAA?
AI deployments remain subject to HIPAA Privacy/Security/Breach rules. Practical implications include vendor BAAs, restrictions on training with PHI, tracking technology risks, and explainability obligations for patient-impacting decisions—areas highlighted in current HHS guidance and legal analyses. [hhs.gov], [hipaajournal.com]

What about EHR-embedded AI?
If a tool meets **HTI-1 DSI** criteria within certified health IT, your EHR vendor must supply transparency artifacts; you remain responsible for safe use, local validation, and clinician enablement. [healthit.gov]

How do we handle fairness?
Adopt a **risk-based equity review** at approval and during monitoring; document "reasonable efforts" per **Section 1557** and adjust workflows, thresholds, or datasets when disparities appear. [nycdentalsociety.org]

What skills do clinicians need?
Beyond basics, the **4Ds** give a common language for safe human-AI collaboration across documentation, patient messaging, and decision support—paired with verification and escalation paths. [anthropic.com]

Case-Oriented Scenarios (Illustrative)

1. **Ambulatory documentation assistance**

- *Risk tier:* Medium (licensed clinician review).
- *Controls:* Prompt library, required edit and sign-off, HTI-1 transparency to users if integrated via certified health IT module.
- *KPIs:* Time saved per note; correction rate; clinician trust; error/near-miss counts; equity check on language complexity for patient-facing outputs. [healthit.gov]

1. **Care navigation messaging**

- *Risk tier:* Medium (patient-facing).
- *Controls:* Mandatory labeling where required, opt-out channels, sensitive topic filters, licensed review for clinical advice, HIPAA tracking tech restrictions. [hhs.gov]

1. **Imaging triage** (regulated SaMD)

- *Risk tier:* High.
- *Controls:* FDA lifecycle alignment, GMLP documentation, post-market surveillance; PCCP-style updates; bias monitoring by modality/site. [fda.gov]

What "Good" Looks Like at Scale

- **Trust first:** Clear statements of intended use, human oversight points, and limits—visible to clinicians and, when relevant, patients—per **HTI-1 transparency** ideals. [healthit.gov]
- **Lifecycle discipline:** Continuous monitoring, bias checks, and secure update pipelines; not one-and-done validations—consistent with **FDA TPLC** thinking. [fda.gov]
- **Operator ownership:** Named workflow owners and AI-fluent operators with authority to adjust standard work; training-to-competency, not training-to-completion. [accu-health.com]
- **Portfolio clarity:** Executives see value and risk across a unified dashboard aligned to **NIST AI RMF** functions. [nist.gov]

Conclusion & Call to Action

AI's trajectory in healthcare is accelerating, but sustainable advantage will come from **fluency**, not feature lists. Leaders who treat AI fluency as **strategic infrastructure**—governed, operator-owned, and measured—will scale responsibly, safeguard equity and safety, and convert innovation into durable clinical and financial outcomes.

Immediate next steps for the C-suite:

1. **Name an executive sponsor** and form your **AI Executive Council** this month. [nist.gov]
2. **Publish a 90-day plan** to:
    - Complete an **AI inventory** (flag DSIs),
    - Stand up the **Algorithm Governance Committee**,
    - Launch **operator foundations** training (4Ds), and
    - Implement **Section 1557** identification/mitigation procedures. [healthit.gov], [anthropic.com], [nycdentalsociety.org]
3. **Choose 3 workflows** for governed pilots with clear KPIs, transparency artifacts, and rollback plans—then scale what works. [openai.com], [fda.gov]

By building this capability now, your organization will be ready for the next wave—safer, faster, and with trust intact.

Acknowledgments & Conceptual Inspiration (no copying)

- Accu-Health's healthcare-specific fluency and operator emphasis informed the "people + process" lens and the maturity signals. [accu-health.com]
- Anthropic's AI Fluency course and 4Ds model helped shape operator competencies. [anthropic.com], [anthropic....illjar.com]
- The AI Fluency community perspectives (AIM: Allocate, Iterate, Mediate) reinforced cross-functional collaboration patterns. [aifluency.org]
- NIST AI RMF, ONC HTI-1, HHS OCR Section 1557, and FDA GMLP provided governance, transparency, nondiscrimination, and lifecycle anchors. [nist.gov], [healthit.gov], [nycdentalsociety.org], [fda.gov]

Appendix: Select Regulatory & Industry References

- **NIST AI RMF 1.0 & GenAI Profile**—trustworthiness and lifecycle risk management. [nist.gov]
- **ONC HTI-1 Final Rule**—transparency for AI-enabled decision support in certified health IT. [healthit.gov]
- **HHS OCR Section 1557**—nondiscrimination duties for patient care decision support tools. [nycdentalsociety.org]
- **FDA GMLP & AI-SaMD lifecycle**—total product lifecycle oversight and PCCP principles. [fda.gov]
- **Responsible AI practices (industry commitments)**—internal/external red-teaming and transparency. [openai.com]
- **HIPAA guidance (HHS/OCR)**—privacy, security, breach rules and tracking tech bulletins relevant to AI-mediated communications. [hhs.gov]