CLEARANCE

AI Governance Diagnostic Report

WALES HEALTH

Healthcare_US · Engagement ref 637cd13c-3720-4e9e-9 · August 05, 2026

HIGH RISK

Engine v1.6 · Library v1.9.6

CLEARANCE

AI Governance Diagnostic Report

Engagement ref 637cd13c-3720-4e9e-9
August 05, 2026
Engine v1.6 · Library v1.9.6
Confidential · Not for distribution

Engagement ref	637cd13c-3720-4e9e-9	**Date issued**	August 05, 2026
Organisation	WALES HEALTH	**Engine version**	v1.6 · Library v1.9.6
Sector	Healthcare_US	**Active layers**	core, us_healthcare

HIGH RISK: 8 governance gaps identified. 0 Critical · 5 High · 3 Medium

CRITICAL	HIGH	MEDIUM	DEPLOYMENT BLOCKERS
0	5	3	0

This report is the output of a fixed-rules diagnostic evaluation. Every gap was triggered by the intake responses against a fixed gap library. The same inputs produce the same findings every time. This is not a consultant's view. It is a structured result. It does not constitute legal or regulatory advice. Specific regulatory obligations should be confirmed with your legal or compliance adviser.

EXECUTIVE SUMMARY

This diagnostic identified 8 governance gaps across WALES HEALTH's AI deployment in the Healthcare_US sector, spanning decision governance and audit capability and HIPAA compliance and patient-data governance. The full findings, regulatory consequences, and remediation actions are set out in this report.

The gaps identified do not represent isolated process failures. They represent the absence of structural governance controls that regulators, courts, and insurers increasingly treat as baseline obligations. Where multiple gaps are present simultaneously, the risk is not additive. It is compounding. A single incident touching an ungoverned AI system will expose every gap at once.

CLEARANCE · AI GOVERNANCE DIAGNOSTIC

Gap Findings

WALES HEALTH · Healthcare_US · August 05, 2026

Every gap below was triggered by intake signals. Each states when it becomes visible and what closes it. Gaps are ordered by severity.

HIGH

G01-CORE-BOUNDARY

No decision boundary enforcement

An incident investigation asks a simple question: what governed this decision before it executed? For this system, there is no answer. The output became the action with nothing in between.

MEDIUM

G02-CORE-PROVENANCE

Provenance weakness

A regulator or auditor asks where the data came from that drove a specific decision. Tracing it back takes days, not minutes, and for some inputs there is no trail at all.

HIGH

G03-CORE-AUDIT

Audit void

A client disputes a decision and asks to see the record. The decision happened. The record of it did not.

HIGH

US-HC-G01-HIPAA-BASIS

No documented legal basis for AI processing of protected health information

If OCR audits your AI deployment, the first question is: what is your legal basis for processing this PHI? Without documented HIPAA authorisation or a recognised exception, there is no answer.

Regulatory obligations: HIPAA Privacy Rule §164.502 · OCR AI Guidance 2025

MEDIUM

US-HC-G02-PURPOSE-LIMITATION

PHI used by AI beyond its original collection purpose without state law compliance

PHI collected for one purpose - say, billing - is being processed by AI for another. Without state law compliance, that is a HIPAA purpose limitation violation.

Regulatory obligations: HIPAA Privacy Rule §164.502(b) · CO SB 26-189 · TX TRAIGA

MEDIUM

US-HC-G03-RETENTION

No retention policy for AI-processed patient data

AI systems create new categories of patient data records. Without a retention policy covering AI-processed PHI, you are accumulating liability with no defined endpoint.

Regulatory obligations: HIPAA Security Rule §164.316 · OCR Guidance · CO SB 26-189

HIGH

US-HC-G04-PATIENT-RIGHTS

No process for patient rights in AI-influenced healthcare decisions

Patients have HIPAA rights over data used in decisions about them. If AI influences those decisions and patients have no mechanism to exercise those rights, that is a structural gap - not a policy gap.

Regulatory obligations: HIPAA Privacy Rule §164.524 · TX TRAIGA §4(b) · CA AB 3030

HIGH

US-HC-G08-AUDIT-LOGGING

No audit logging for AI access to or use of PHI

During an OCR investigation, nobody can produce logs showing what patient data the AI accessed, when, and why.

Regulatory obligations: HIPAA Security Rule 45 CFR 164.312(b), audit controls (required) · HIPAA Security Rule 45 CFR 164.308(a)(1)(ii)(D), information system activity review · OCR Audit Protocol, Technical Safeguards §164.312(b) · NIST SP 800-66r2 §3.13, audit controls implementation guidance

CLEARANCE · AI GOVERNANCE DIAGNOSTIC

Regulatory Obligations

WALES HEALTH · Healthcare_US · August 05, 2026

The obligations below are engaged by the gaps identified in this report. GAP: a triggered finding directly engages this obligation. REVIEW: this obligation is relevant to your sector and worth examining. NOT COVERED: this assessment did not include questions for this obligation.

OBLIGATION	RELEVANCE	STATUS
DCB0160	Clinical safety case required before deployment of software in clinical settings	REVIEW
NHS DSPT	Data Security and Protection Toolkit, annual submission required	REVIEW
NHS DTAC	Digital Technology Assessment Criteria, required for NHS procurement	REVIEW

NEXT STEP

CLEARANCE identifies gaps. It does not remediate them. Remediation is carried out by a qualified partner. Contact whoever provided this report to discuss next steps.

reports@clearance-diagnostic.com

CLEARANCE · AI GOVERNANCE DIAGNOSTIC

Remediation Roadmap

WALES HEALTH · Healthcare_US · August 05, 2026

One action per gap. Ordered by timeline. Critical gaps require immediate action. High severity gaps must be addressed within 90 days. Medium severity gaps within 180 days.

TIMELINE	GAP	ACTION REQUIRED	PRIORITY	OWNER
90 DAYS	**G01-CORE-BOUNDARY** No decision boundary enforcement	Introduce a rule evaluation layer between this system's output and the action it takes. Even a minimal documented ruleset is a material improvement on none.	HIGH	your system architect
90 DAYS	**G03-CORE-AUDIT** Audit void	Start logging every decision this system makes today: date, input, output, rule applied, in a write-once store. This is the foundation. Every other governance control depends on it existing.	HIGH	your compliance lead
90 DAYS	**US-HC-G01-HIPAA-BASIS** No documented legal basis for AI processing of protected health information	Document the HIPAA legal basis for every AI use of PHI this week. Assign a named owner. Review with legal counsel within 30 days.	HIGH	your compliance lead
90 DAYS	**US-HC-G04-PATIENT-RIGHTS** No process for patient rights in AI-influenced healthcare decisions	Define and document the process for patients to request access to, or correction of, AI-influenced decision records. Disclose AI use in patient-facing communications.	HIGH	your compliance lead

TIMELINE	GAP	ACTION REQUIRED	PRIORITY	OWNER
90 DAYS	**US-HC-G08-AUDIT-LOGGING** No audit logging for AI access to or use of PHI	Implement audit logging at the AI inference layer before the next deployment cycle.	HIGH	your IT / security lead
180 DAYS	**G02-CORE-PROVENANCE** Provenance weakness	Create a register of every data source feeding this system. Assign an owner and a validation status to each.	MEDIUM	your system architect
180 DAYS	**US-HC-G02-PURPOSE-LIMITATION** PHI used by AI beyond its original collection purpose without state law compliance	Map every AI input to its original collection purpose. Flag any processing that goes beyond that purpose. Check state law requirements for your operating states.	MEDIUM	your compliance lead
180 DAYS	**US-HC-G03-RETENTION** No retention policy for AI-processed patient data	Extend your existing data retention policy to cover AI-processed PHI specifically. Define retention periods by PHI type. Confirm with your privacy officer.	MEDIUM	your compliance lead

CLEARANCE · AI GOVERNANCE DIAGNOSTIC

Evidence Appendix

WALES HEALTH · Healthcare_US · August 05, 2026

ASSESSMENT RESPONSES

QUESTION	RESPONSE	CATEGORY
AI system is live in production or active pilot	Yes	Your AI system
System makes or influences high-impact decisions	Yes	Your AI system
Decisions execute without individual human review	Yes	Your AI system
A written decision boundary document exists	No	Rules and boundaries
That document is version-controlled and enforced at runtime	No	Rules and boundaries
System uses data from outside the organisation	Yes	Data
Purpose-compatibility assessment completed for external data use	No	Data
Every input is traceable to its origin, owner, and last update	Yes	Data
Automated validation runs before data reaches the system	No	Data
Every decision is logged automatically at the time it happens	Yes	Records and accountability
A tamper-evident, provable audit trail exists	No	Records and accountability
A functioning human override mechanism exists and is used	Yes	Records and accountability
A named individual is personally accountable for this system	No	Records and accountability
Documented data subject rights process covers this AI system	No	Records and accountability
An incident, near-miss, or complaint has occurred	Yes	When things go wrong
Documented corrective action was taken following the incident	Yes	When things go wrong
Does the system process or access Protected Health Information (PHI) as defined under HIPAA?	Yes	HIPAA Controls
Are Business Associate Agreements (BAAs) in place covering all AI vendors and their use of PHI?	Yes	HIPAA Controls
What type of PHI does this AI system process?	clinical_notes	HIPAA Controls
Are documented access controls in place governing who or what can access the PHI processed by this AI?	Yes	HIPAA Controls
Are patients informed when an AI system influences decisions about their care, treatment, or coverage?	Yes	HIPAA Controls
List the US states where this AI system operates or makes decisions affecting patients (comma-separated):	['TX']	HIPAA Controls
Is PHI encrypted at rest and in transit through the AI pipeline?	Yes	HIPAA Controls
Does your incident response plan cover AI-specific PHI breach scenarios?	Yes	HIPAA Controls

QUESTION	RESPONSE	CATEGORY
Is PHI transmitted to third-party AI APIs (OpenAI, Azure, Anthropic, etc.)?	**Yes**	HIPAA Controls

ASSESSMENT INTEGRITY

Organisation: WALES HEALTH
Sector: Healthcare_US
Run ID: 637cd13c-3720-4e9e-90f3-1928de611325
Engine version: v1.6 · **Library version:** v1.9.6
Active layers: core, us_healthcare
Findings hash: 637cd13c-3720-4e9e-90f3-1928de611325
Generated: 2026-08-05T18:07:58.501713+00:00