


HEALTHCARE

AI in Gulf healthcare: evidence at the point of care

Clinical and administrative workflows, responsibility and the boundaries of deployment.



ORIGINAL MAYDEN GRAPHIC / OPERATING QUESTIONS

01 / THE EXECUTIVE VIEW

Operating throughput, retrospective evaluation and clinical authority are separate lines of evidence.

Clinical and administrative workflows, responsibility and the boundaries of deployment.

10

selected entity records

32

recorded claims

27

qualifying claim records

9

candidate research themes

01 / THE OPERATING QUESTION

Operating throughput, retrospective evaluation and clinical authority are separate lines of evidence.

02 / THE OPERATING QUESTION

The AIRIS-TB paper evaluates 1.04 million chest X-rays from January 2022 to January 2024 in an observational retrospective analysis in Abu Dhabi. Projected reporting reduction is calculated from classifications, not demonstrated prospective staffing savings. External-dataset performance and restricted sputum-confirmation pathways limit generalisation. M42/CHSC affiliations and model-developer involvement remain relevant. **Population-scale cross-sectional observational study for AI-powered TB screening on one million CXRs.**

03 / THE OPERATING QUESTION

Clinical, operational and technical reviewers should agree intended use, comparator, period and outcome definitions before collection. Include verification burden and subgroup performance where appropriate. Preserve disputed outputs and corrections.

AI-assisted, bounded public-source research. Purposefully selected entity scopes; no company participation, endorsement, independent assurance or population inference implied. Confidential assessment is in development.

02 / THE ARGUMENT

Four questions to carry into the business.

Read each step as a test of the operating claim. This is an analytical framework, not a measured maturity ladder.

- 1 What activity is reported?**
Documentation, image interpretation and screening coordination change different work. Specify provider, service, patient population, professional role and period.
- 2 What exactly was evaluated?**
The AIRIS-TB paper evaluates 1.04 million chest X-rays from January 2022 to January 2024 in an observational retrospective analysis in Abu Dhabi. Projected reporting reduction is calculated from classifications, not demonstrated prospective staffing savings.
- 3 Who retains clinical authority?**
Final authority should be observable in an accepted or rejected output, correction and follow-up. For documentation, ask who verifies the note and how errors are corrected.
- 4 What would justify a new setting or population?**
Clinical, operational and technical reviewers should agree intended use, comparator, period and outcome definitions before collection. Include verification burden and subgroup performance where appropriate.

How to use this report

Follow the source accounts and their limits, use the diagrams to frame an evidence request, then choose one recurring workflow for deeper examination. A company name identifies an attributed account; it does not imply participation, endorsement or an enterprise rating.

03 / THE ANALYSIS

Operating throughput, retrospective evaluation and clinical authority are separate lines of evidence.

Healthcare AI needs distinct evidence layers: a system may process examinations, perform well retrospectively and still need separate evidence of benefit in a live pathway. Mayden follows the workflow and accountable clinical decision; this is not a clinical-quality rating.

Identify the clinical and organisational question

Documentation, image interpretation and screening coordination change different work. Specify provider, service, patient population, professional role and period. A company-wide statement cannot establish suitability for every clinical use.

Throughput describes activity. Retrospective performance describes evaluation under a study design. Prospective effects on care, workload and safety need further observations. A technical result should not carry a broader claim than its evidence supports.

Read the study in its actual setting

The AIRIS-TB paper evaluates 1.04 million chest X-rays from January 2022 to January 2024 in an observational retrospective analysis in Abu Dhabi. Projected reporting reduction is calculated from classifications, not demonstrated prospective staffing savings. External-dataset performance and restricted sputum-confirmation pathways limit generalisation. M42/CHSC affiliations and model-developer involvement remain relevant. **Population-scale cross-sectional observational study for AI-powered TB screening on one million CXRs.**

This is locally relevant research, not an enterprise-capability measure or proof of a new deployment's effect in another population. The

accompanying cases concern distinct evidence layers and must not be pooled into a provider comparison.

Attributed operating examples

These separate cases were selected by the operating question, not company standing. The diagram does not imply one connected system.

Operating activity

M42 reports AIRIS-TB conducting 2,000 AI chest-X-ray examinations per day. Issuer-reported throughput; staffing/comparator details incomplete. The next research question is: How many eligible examinations reach a reviewed clinical decision?

Source: **M42 and GE HealthCare join forces to advance AI-Enabled Patient Care.**

Retrospective evaluation

The original AIRIS-TB study evaluates 1.04 million chest X-rays from January 2022 to January 2024 in an observational retrospective analysis. M42/CHSC-affiliated research; model evaluation is not prospective patient benefit. The next research question is: Does performance hold in the intended workflow, population and prevalence?

Source: **M42 announces breakthrough results for its AI-powered tuberculosis screening.**

Clinical authority

HMG says its Responsible AI Framework retains clinician final authority and includes fairness monitoring. A different provider and evidence layer; no assurance result supplied. The next research question is: Can an actual disagreement show how review, intervention and follow-up operate?

Source: **Annual Report 2025.**

EXHIBIT 01 / ATTRIBUTED ACCOUNT

Do not put throughput and evaluation size on the same axis

Activity, model evaluation and clinical authority answer different questions. Cases are not pooled into a clinical-quality score.

01	Operating activity Screening examinations	M42 reports AIRIS-TB conducting 2,000 AI chest-X-ray examinations per day. Boundary: Issuer-reported throughput; staffing/comparator details incomplete. Ask: How many eligible examinations reach a reviewed clinical decision?
02	Retrospective evaluation Radiographs assessed in a study	The original AIRIS-TB study evaluates 1.04 million chest X-rays from January 2022 to January 2024 in an observational retrospective analysis. Boundary: M42/CHSC-affiliated research; model evaluation is not prospective patient benefit. Ask: Does performance hold in the intended workflow, population and prevalence?
03	Clinical authority A clinician’s accepted or rejected output	HMG says its Responsible AI Framework retains clinician final authority and includes fairness monitoring. Boundary: A different provider and evidence layer; no assurance result supplied. Ask: Can an actual disagreement show how review, intervention and follow-up operate?

Company and supplier accounts are attributed descriptions, not independent assurance. Each row describes a different operating task; no company comparison or rating is intended.

Unit: Separate source-described operating cases. Research cutoff: 2026-10-10.

Sources: [HLT-012](#) / [HLT-013](#) / [L34](#) / [HLT-002](#)

03 / THE ANALYSIS

Evaluate the combined workflow

Final authority should be observable in an accepted or rejected output, correction and follow-up. For documentation, ask who verifies the note and how errors are corrected. For screening, trace the output to a reviewed decision and identify eligible examinations outside the workflow. These are research questions, not treatment recommendations.

A randomised study of 50 physicians found no significant improvement in its primary diagnostic-reasoning score from chatbot access. It used vignettes, not patient outcomes. This illustrates why the combined workflow needs evaluation; it does not settle the value of other systems. [Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial](#).

Keep the next study accountable

Clinical, operational and technical reviewers should agree intended use, comparator, period and outcome definitions before collection. Include verification burden and subgroup performance where appropriate. Preserve disputed outputs and corrections.

The future company diagnostic requires qualified human review and permissioned evidence. It must not award clinical assurance, infer patient benefit

from disclosure volume or interpret limited public detail as unsafe care.

Leadership agenda

Clinical owner: define intended use, population and final authority.

Operations: measure eligible work, reviewed outputs and correction effort.

Research reviewer: separate retrospective performance, projected effects and observed workflow outcomes.

Scope and reading boundaries


The source corpus contains 80 purposefully selected entity records across eight sectors, ten per sector. Headquarters markets are Saudi Arabia, the UAE, Qatar, Kuwait and Bahrain; Oman is outside this edition. Groups, subsidiaries and specific operations have different scopes. This is not a representative sample. Source cutoff: 10 October 2026; publication edition: 11 October 2026.

Source extraction and challenge were AI-assisted, not independent human assurance or verification of current operations. Source visibility is not adoption prevalence. Unreported details are unresolved questions, not absent capabilities. Named examples imply neither participation, endorsement, confidential access nor a Mayden client relationship.

EXHIBIT 02 / CONCEPTUAL TOOL

A study's operating boundary matters as much as its sample size

Mayden analytical framework / questions to test, not measured maturity stages



Read each study with its original setting, reference standard, observation period and limits. Missing details remain unknown. This is a conceptual appraisal framework, not an assurance of clinical safety.

Unit: Conceptual decision or evidence request. Research cutoff: 2026-10-10.

EXHIBIT 03 / CONCEPTUAL TOOL

Keep the clinical decision visible in the workflow

Mayden analytical framework / questions to test, not measured maturity stages



Mayden conceptual framework. These questions specify evidence to request; they do not assert that a named organisation has or lacks the practice. This is not a validated instrument or an observed causal sequence.

Unit: Conceptual decision or evidence request. Research cutoff: 2026-10-10.

04 / SOURCE ACCOUNT

Operating activity

Operating object: Screening examinations

HLT-012 / ATTRIBUTED ACCOUNT

M42 reports AIRIS-TB conducting 2,000 AI chest X-ray examinations daily.

Scope. Abu Dhabi-headquartered international healthcare group; cases identify Capital Health Screening Centre or other operating subsidiaries.

Source date. 2025-07-08. Publication date does not confirm continuing operation.

Locator. Final AIRIS-TB paragraph, line69

M42 and GE HealthCare join forces to advance AI-Enabled Patient Care

Boundary and next question

Issuer-reported throughput; staffing/comparator details incomplete.

How many eligible examinations reach a reviewed clinical decision?

04 / SOURCE ACCOUNT

Retrospective evaluation

Operating object: Radiographs assessed in a study

HLT-013 / ATTRIBUTED ACCOUNT

M42 reports AIRIS-TB validation using more than 1 million screening radiographs and demographic performance checks.

Scope. Abu Dhabi-headquartered international healthcare group; cases identify Capital Health Screening Centre or other operating subsidiaries.

Source date. 2025-08-06. Publication date does not confirm continuing operation.

Locator. Study design, results and demographic consistency paragraphs

M42 announces breakthrough results for its AI-powered tuberculosis screening

Boundary and next question

M42/CHSC-affiliated research; model evaluation is not prospective patient benefit.

Does performance hold in the intended workflow, population and prevalence?

04 / SOURCE ACCOUNT

Clinical authority

Operating object: A clinician's accepted or rejected output

HLT-002 / ATTRIBUTED ACCOUNT

HMG reports a Responsible AI Framework with clinician final authority, security requirements and fairness monitoring.

Scope. Listed healthcare group; HMG hospital workflows distinguished from Cloud Solutions sales to third parties.

Source date. Undated. Publication date does not confirm continuing operation.

Locator. PDF39/printed 75-76, Applying Responsible AI

Annual Report 2025

Boundary and next question

A different provider and evidence layer; no assurance result supplied.

Can an actual disagreement show how review, intervention and follow-up operate?

05 / THE LEADERSHIP AGENDA

Ask for evidence that changes the decision.

01 / Clinical owner

Evidence to request: define intended use, population and final authority.

02 / Operations

Evidence to request: measure eligible work, reviewed outputs and correction effort.

03 / Research reviewer

Evidence to request: separate retrospective performance, projected effects and observed workflow outcomes.

Method and boundaries

This is an AI-assisted, bounded public-source review of a purposefully selected corpus. It publishes sector research and attributed operating examples, without current individual-company scores, ranks or coded profiles. Source accounts establish what a publisher reports; they do not independently assure operations. The working definition and nine candidate themes require empirical validation.

Recorded source dates run from 2024-12-30 to 2025-08-06. 22 of 27 qualifying claim records have no recorded publication date. This is not a common observation window or confirmation of continuing use.

The confidential opt-in diagnostic remains in development. Independent human review, instrument testing, participant data collection and comparable benchmark norms are future gates. Prior named public-evidence comparisons remain explicitly historical; changes in publication format do not establish changes inside companies.

[Inspect the method](#) / [Explore sector research](#)

06 / SOURCES & PROVENANCE

Follow the argument back to its source.

The following sources support the cited analysis or attributed examples. Inclusion is not independent verification, company endorsement or assurance. Research cutoff: 10 October 2026. Publication revision: 11 October 2026.

S01 / M42 and GE HealthCare join forces to advance AI-Enabled Patient Care

M42 / 2025-07-08 / official-implementation-account

Locator: Final AIRIS-TB paragraph, line69

[Open source](#)

S02 / M42 announces breakthrough results for its AI-powered tuberculosis screening

M42 / 2025-08-06 / official-implementation-account

Locator: Study design, results and demographic consistency paragraphs

[Open source](#)

S03 / Annual Report 2025

HMG / Undated / annual-report

Locator: PDF39/printed 75-76, Applying Responsible AI

[Open source](#)

S04 / Population-scale cross-sectional observational study for AI-powered TB screening on one million CXRs

npj Digital Medicine; full text in PMC / 2025-07-09 / original-observational-model-evaluation

Locator: Retrospective observational evaluation; authors also describe a quasi-prospective approach.

[Open source](#)

S05 / Large Language Model Influence on Diagnostic Reasoning: A Randomized Clinical Trial

JAMA Network Open / 2024-10-28 / peer-reviewed-randomized-trial

Locator: Stratified randomised, blinded-rating trial

[Open source](#)

About Mayden

Mayden connects senior judgement, research and engineering to help organisations turn strategic choices into working capabilities. The AI Native Index is a research programme on recurring work, decisions and accountability.

[Explore the Index](#) / [Submit a correction](#)