QMEDIC-QMLIR and
Quantum Longitudinal Diagnostic Intelligence

A Fidelity-Preserving Reconstruction Framework for Low-Dose CT, Patient State Space Modeling, and Quantum-Enabled Clinical Intelligence

Technical White Paper v1.0 - Zenodo Upload Draft

Prepared by SQK Inc.
Seattle, Washington, USA

2026

Field	Content
Company	SQK Inc.
Product / Platform	QMEDIC, QMLIR, QPINN, QPR, QLDI
Document Type	Technical white paper / foundational vision paper for Zenodo upload
Primary Audience	Investors, Mayo Clinic Platform, IBM Quantum, NIH SBIR/STTR reviewers, clinical research partners, strategic OEM partners
Regulatory Status	Research-use and development-stage concept. Not for clinical diagnosis unless and until cleared or authorized by applicable regulators.
Suggested License	CC BY 4.0 or company-approved open-access license, subject to SQK legal review.

Note: This document is intended as a foundational technical white paper and commercialization narrative. Clinical, regulatory, and market claims should be reviewed by qualified clinical, regulatory, and legal advisors before public submission or investor distribution.

Abstract

Low-Dose Computed Tomography (LDCT) is clinically valuable because it can reduce patient radiation exposure, but reduced dose often introduces noise, artifacts, low-contrast degradation, and loss of subtle anatomical structures. These limitations can reduce diagnostic confidence and constrain broad LDCT adoption in screening, follow-up, and oncology workflows.

QMEDIC introduces QMLIR (Quantum Medical Low-dose Image Reconstruction), a customizable reconstruction-first framework designed to operate before conventional DICOM post-processing. QMLIR is positioned to preserve diagnostically meaningful structure, improve low-contrast detectability, and provide higher-fidelity images for radiologists and downstream AI systems. Unlike conventional AI post-processing tools that enhance already reconstructed DICOM images, QMLIR targets the reconstruction stage using sinogram/projection-domain information when available.

The platform further incorporates QPINN as a physics-informed trust layer, QPR as a quantum patch refinement and pattern-recognition engine, and QLDI (Quantum Longitudinal Diagnostic Intelligence) as a future diagnostic intelligence framework. QLDI expands QMEDIC from image reconstruction into patient state-space modeling, retrospective reassessment of prior CT findings, multimodal convergence, and future disease progression forecasting.

This white paper presents the QMEDIC-QMLIR architecture, key technical differentiation, clinical validation direction, preliminary reported SNUH thyroid cancer CT artifact-reduction result, expected benefits for patients and physicians, and a staged commercialization strategy that begins with research-use services during FDA 510(k) and clinical validation preparation, followed by reference-site expansion and enterprise/OEM deployment.

Keywords

Low-Dose CT; LDCT; QMLIR; QMEDIC; QPINN; QPR; Quantum Pattern Recognition; Patient State Space; Quantum Longitudinal Diagnostic Intelligence; QLDI; Medical Imaging AI; FDA 510(k); SaMD; Quantum-HPC; DICOM; Sinogram Reconstruction

Table of Contents

1. Executive Summary

QMEDIC is designed to transform CT image reconstruction from a visual enhancement problem into a reliability-centered medical intelligence platform. The near-term technical wedge is QMLIR, a customizable LDCT reconstruction framework intended to improve image fidelity at the reconstruction stage. The long-term platform vision is QLDI, a longitudinal diagnostic intelligence layer that models patient disease state across time and modalities.

The core strategic distinction is reconstruction-first intelligence. Conventional medical imaging AI tools typically operate after DICOM image generation. They may denoise, segment, annotate, quantify, or detect structures from already reconstructed images. QMEDIC aims to work earlier in the imaging pipeline by improving the foundation image itself before downstream AI or radiologist review.

1.1 Investment and Strategic Thesis

- **Clinical wedge:** Make LDCT more clinically practical by improving image fidelity, low-contrast detectability, and subtle-structure preservation.
- **Technical wedge:** Differentiate through QMLIR reconstruction, QPINN physical consistency, QPR patch refinement, and QLDI patient-state modeling.
- **Commercial wedge:** Begin with research-use sales and service during clinical validation and regulatory preparation, then scale through reference sites, enterprise healthcare systems, and OEM partnerships.
- **Quantum wedge:** Use quantum and quantum-inspired computation selectively for high-dimensional feature learning, patient state-space exploration, and future Quantum-HPC clinical intelligence.

1.2 Positioning Statement

QMEDIC is a reconstruction-first, trust-centered medical imaging platform that combines fidelity-preserving LDCT reconstruction with quantum-enhanced pattern recognition and longitudinal diagnostic intelligence. Its goal is to support standard-dose CT-equivalent diagnostic image quality from LDCT data while preserving clinically meaningful structures and enabling future patient-state intelligence.

2. Clinical Problem and Unmet Need

The clinical value of LDCT is clear: lower radiation exposure can support screening, follow-up, and repeat imaging. However, reducing dose often reduces photon counts and increases uncertainty in the image formation process. The resulting images may contain noise, streak artifacts, low-contrast degradation, blurred boundaries, and suppressed subtle structures. These effects can be especially problematic in oncology, lung screening, thyroid follow-up, liver metastasis surveillance, and longitudinal disease monitoring.

2.1 Why Existing Pipelines Are Not Enough

Vendor DLIR systems can provide robust baseline reconstruction, and conventional AI post-processing can improve denoising, segmentation, CAD, annotation, or quantification. However, most downstream AI systems assume that the reconstructed DICOM image already contains the relevant diagnostic signal. If the foundation image has lost subtle lesion information, post-processing may not reliably recover it.

- Image enhancement can improve appearance without guaranteeing preservation of diagnostic information.
- Over-smoothing can reduce visibility of low-contrast lesions or small structures.
- Downstream AI performance depends on input fidelity.
- Clinical users need confidence, traceability, and error-aware outputs rather than cosmetic enhancement alone.

2.2 The Deeper Problem: Trust in Reconstructed Medical Images

The central problem is not noise alone. The deeper issue is whether a reconstructed image remains physically consistent with source data and clinically reliable for human and AI interpretation. In high-trust imaging workflows, the key question is not only whether an image looks better, but whether diagnostically meaningful information has been preserved.

This creates the need for a reconstruction-first trust layer: a system that improves LDCT image quality while preserving anatomical fidelity, quantifying uncertainty, and supporting downstream AI workflows.

3. QMEDIC Reconstruction-First Workflow

Figure 1 illustrates the QMEDIC + conventional AI post-processing workflow. The diagram emphasizes that QMLIR is positioned at the reconstruction stage, while conventional AI operates after DICOM CT image generation. This is a strategic distinction: QMEDIC is not simply another post-processing tool, but a reconstruction-enabling technology designed to improve the input quality for both radiologists and AI systems.

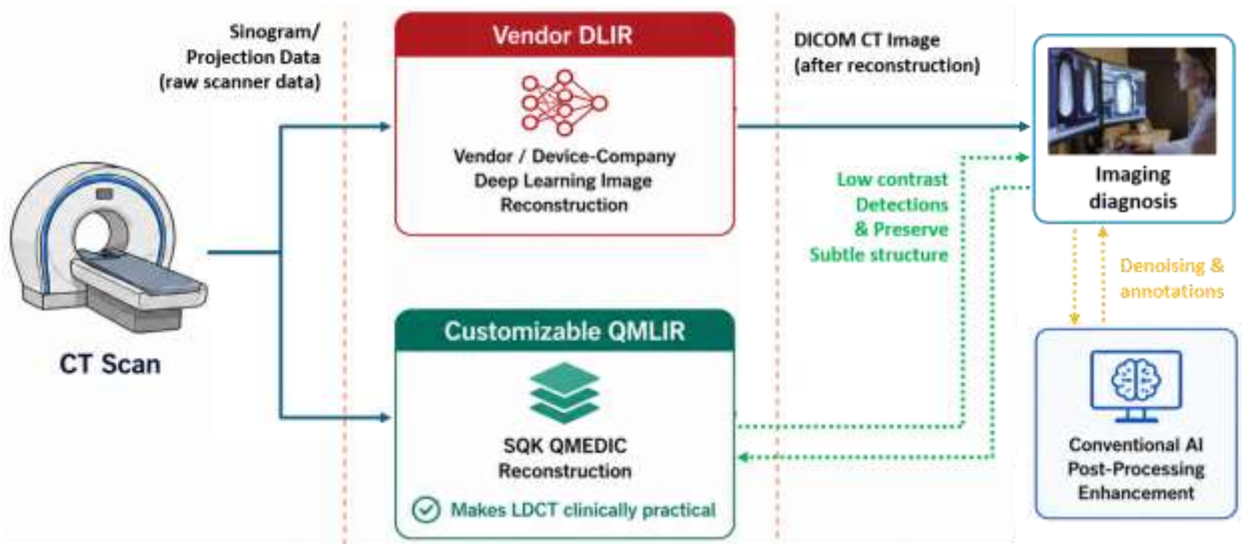

Figure 1. QMEDIC reconstruction-first clinical AI workflow. QMLIR is positioned before conventional AI post-processing to preserve subtle structures and improve low-contrast detectability.

3.1 Workflow Interpretation

- **Raw acquisition:** CT scan data can be represented through sinogram/projection data before DICOM reconstruction is finalized.
- **Vendor DLIR:** Vendor reconstruction provides a fast and stable baseline path.
- **Customizable QMLIR:** QMEDIC reconstruction focuses on LDCT fidelity, low-contrast detectability, and subtle-structure preservation.
- **Conventional AI post-processing:** Downstream systems can perform denoising, segmentation, CAD, quantification, annotation, and clinical decision support on improved inputs.
- **Radiologist-in-the-loop:** The workflow supports expert review and clinical interpretation rather than autonomous diagnosis.

3.2 Strategic Meaning

QMEDIC is best positioned as an enabling layer for LDCT commercialization. It can coexist with vendor DLIR and conventional post-processing AI by improving the fidelity of the reconstructed image and supplying a stronger foundation for downstream analysis.

4. QMLIR: Core Technology and Differentiation

QMLIR stands for Quantum Medical Low-dose Image Reconstruction. It is the central near-term technology of QMEDIC. The objective is to reconstruct LDCT images in a way that preserves clinically meaningful structures while reducing artifact burden and improving image quality under low-dose acquisition conditions.

4.1 QMLIR Design Principles

- **Fidelity over appearance:** Optimization should prioritize diagnostic information preservation rather than visual smoothness alone.
- **Projection-domain awareness:** When available, sinogram/projection-domain information can provide stronger reconstruction constraints than DICOM-only post-processing.
- **Low-contrast lesion support:** QMLIR is designed for conditions where small or low-contrast findings may otherwise be obscured.
- **AI-enabling output:** QMLIR is intended to generate higher-quality DICOM images for radiologists and downstream AI systems.
- **Customizable deployment:** QMLIR can be configured for research, clinical validation, vendor integration, or cloud/API environments.

4.2 Technical Differentiation

Approach	Pipeline Position	Primary Role	Strategic Limit
Conventional AI post-processing	After DICOM image reconstruction	Denoising, super-resolution, segmentation, CAD, quantification	Dependent on existing reconstructed image quality
Vendor DLIR	Scanner/vendor reconstruction stage	Fast baseline reconstruction and vendor-optimized image quality	Often scanner- and vendor-specific
QMEDIC QMLIR	Earlier reconstruction / reconstruction-enabling stage	LDCT fidelity, low-contrast detectability, subtle-structure preservation, downstream AI enablement	Vendor-agnostic strategy with future Quantum-HPC extension

4.3 Innovation Area: LDCT Commercialization Enablement

QMLIR can be framed as a practical commercialization layer for LDCT. The clinical value of LDCT is constrained when images are too noisy or uncertain for reliable interpretation. By improving the reconstruction stage and preserving subtle structures, QMLIR can help expand the practical use of LDCT in screening and follow-up workflows.

The preferred benchmark language for QMEDIC is standard-dose CT-equivalent diagnostic image quality from LDCT data. This does not imply automatic equivalence across all indications; rather, it defines the target clinical-quality threshold to be validated through structured reader studies, image-quality metrics, and task-specific outcomes.

5. QPINN: Physics-Informed Trust Layer

QPINN is positioned as QMEDIC's physics-informed trust layer. Its purpose is to reduce the risk that image improvement creates visually plausible but physically inconsistent outputs. In medical imaging, hallucination risk can include false texture, distorted boundaries, masked pathology, or over-smoothed anatomy.

5.1 Core Functions

- **Physics consistency:** Evaluate whether reconstructed output remains aligned with source-domain constraints such as CT projection or sinogram consistency.
- **Uncertainty estimation:** Identify ambiguous or unstable regions through confidence maps, entropy, calibration error, or ensemble variance.
- **Artifact-risk signaling:** Flag regions where residual error, uncertainty, or physical inconsistency exceeds an acceptable threshold.

- **Conservative fallback:** Route low-confidence cases to expert review or conservative reconstruction rather than forcing aggressive enhancement.

5.2 Inference-Phase Guardrail

A key design principle is that quality control should not occur only during model training. QPINN can apply inference-phase checks after reconstruction or refinement. This allows each output image or region of interest to be evaluated for probabilistic uncertainty and physical forward consistency before being used downstream.

Simplified conceptual residual check:

Res = F(y_hat)

where y_hat is the reconstructed output and F is a modality-specific forward-consistency operator. If the residual exceeds a defined tolerance threshold, the region can be flagged for review.

5.3 Clinical Outputs

Output	Meaning	Clinical / Operational Use
Confidence map	Shows regions of higher or lower model confidence	Radiologist prioritization and QC
Residual map	Shows physical consistency deviations	Hallucination/artifact-risk detection
Artifact-risk flag	Marks suspicious reconstructed regions	Conservative review and fallback
QC report	Summarizes settings, metrics, and audit metadata	Regulatory and clinical validation readiness

6. QPR: Quantum Patch Refinement and Pattern Recognition

QPR is the QMEDIC engine for targeted refinement and pattern recognition. In the near-term reconstruction context, QPR can support patch-level refinement, subtle-lesion preservation, artifact reduction, and edge restoration. In the future QLDI context, QPR can serve as a longitudinal pattern-recognition engine for patient state-space analysis.

6.1 Why Patch-Level Analysis Matters

Whole-volume CT processing can be computationally expensive and clinically inefficient. A more practical strategy is ROI-focused inference: use present images, AI detections, radiologist annotations, PET findings, or lab signals to identify suspicious regions and then focus QPR/ML/simulation on those regions. This can reduce computation and improve sensitivity to hidden abnormalities.

- Focuses computation on suspicious ROIs rather than the entire CT volume.
- Supports retrospective review of prior CT at the same anatomical location.
- Improves the story for quantum use by framing the problem as high-dimensional ROI state-space inference rather than generic image filtering.

6.2 Quantum AI and QUKKOS Research Foundation

Prior QMEDIC research describes a hybrid quantum-classical AI for photon-counting CT super-resolution using a variational quantum circuit as a generator in the latent space of a classical autoencoder. The same research direction also connects the model to QUKKOS, a measurement-based quantum compilation system designed to improve efficiency, resource requirements, and noise robustness.

This research foundation supports the broader QPR thesis: quantum-enhanced generative learning and high-dimensional feature extraction may be useful for medical image refinement tasks where small structures and subtle distributions are clinically meaningful.

6.3 Longitudinal Pattern Recognition

For QLDI, QPR can move beyond image patch refinement into longitudinal pattern recognition. The system can identify subtle patterns across past CT, current CT, PET, biomarkers, pathology, and clinical history that may indicate hidden lesions, early metastasis, or future progression before they become obvious in a single image.

7. QMEDIC Target Use Cases

QMEDIC is designed as a reconstruction-first medical imaging platform that can support multiple research, validation, and commercialization pathways. The initial focus is low-dose CT (LDCT), where improved reconstruction fidelity, artifact suppression, subtle-structure preservation, and uncertainty-aware validation can directly support clinical research and downstream medical AI workflows.

7.1 LDCT Micro-Nodule Detection Support

The first target use case is LDCT-based pulmonary micro-nodule detection support. SQK is preparing a research direction with Mayo Clinic focused on improving the detection and visibility of small lung nodules, particularly micro-nodules at or below 4 mm.

In this use case, QMEDIC uses LDCT DICOM data, and when available, reconstruction-stage information, to improve image fidelity, reduce noise and artifacts, and preserve subtle low-contrast structures. The goal is to provide higher-quality imaging inputs for radiologists and downstream AI systems that support lung screening, follow-up, and early oncology workflows.

This use case positions QMEDIC as an enabling technology for clinically practical LDCT, especially in scenarios where small lesions may be obscured by low-dose noise, reconstruction artifacts, or over-smoothing.

7.2 Longitudinal Diagnostic Intelligence for Prior CT Reassessment

The second target use case is Longitudinal Diagnostic Intelligence (LDI), being prepared as a research direction with the University of Chicago. This use case focuses on patients who currently have a detected lesion of 5 mm or larger. QMEDIC would retrospectively reconstruct and reassess prior CT studies at the same anatomical location where the lesion was not previously detected.

The objective is to determine whether subtle features, early signals, or hidden phenotypes can be measured from prior CT data after QMEDIC-based reconstruction or refinement. These reconstructed features may then be used to support longitudinal analysis, disease progression estimation, and future risk prediction.

This use case expands QMEDIC beyond single-scan image improvement into patient-state modeling, retrospective reassessment, and future progression forecasting. It aligns with the QLDI vision of modeling disease state over time using imaging, lesion features, clinical context, and multimodal evidence.

7.3 Artifact Reduction with Hallucination-Suppressed Reconstruction

A third use case is artifact reduction with hallucination-suppressed reconstruction. SQK's previous research direction focused on improving artifact reduction while incorporating safeguards against AI-generated distortion or clinically unsafe enhancement.

In this workflow, QMEDIC is used to reduce reconstruction artifacts while preserving clinically meaningful anatomy and minimizing the risk of false texture, structure loss, or lesion masking. QPINN can support this use case by applying physics-informed consistency checks, uncertainty estimation, residual mapping, and artifact-risk flagging.

This use case is especially relevant for clinical imaging workflows where image enhancement alone is not sufficient. The key requirement is not only to make the image look cleaner, but to ensure that the reconstructed image remains reliable, traceable, and clinically meaningful.

7.4 Downstream Medical AI Enablement

QMEDIC can also be used as an upstream image-quality improvement layer for downstream medical AI systems. Many AI tools for segmentation, CAD, quantification, annotation, and clinical decision support depend heavily on the quality of the input DICOM image. If the foundation image contains noise, artifacts, or lost subtle structures, downstream AI performance may be limited.

By improving image quality at the reconstruction or refinement stage, QMEDIC can provide stronger inputs for downstream AI models. This use case positions QMEDIC as a productivity and reliability layer for medical AI developers, imaging informatics teams, and healthcare AI platforms.

7.5 Customer-Driven Research and Validation Projects

In addition to the above use cases, SQK is prepared to adapt QMEDIC to customer-defined clinical research needs. Medical centers, research hospitals, imaging laboratories, and strategic partners may have specific priorities based on anatomy, modality, disease area, scanner environment, or clinical workflow.

Potential customer-driven projects may include oncology follow-up, lung screening, thyroid CT, liver metastasis surveillance, low-contrast lesion visibility, multimodal imaging convergence, or AI-ready DICOM generation. SQK is prepared to work with clinical and research partners to define study design, validation metrics, deployment mode, and evidence-generation strategy according to each partner's needs.

Overall, QMEDIC's target use cases are designed to support a staged path from research-use validation to clinical workflow integration and future regulated commercialization.

8. Clinical Validation and Evidence Strategy

QMEDIC validation should be presented in layers: technical image-quality validation, clinical task validation, workflow validation, and regulatory-grade documentation. This staged structure is important because image metrics alone do not establish clinical value.

8.1 Preliminary Clinical Signal: SNUH Thyroid Cancer CT

SQK has reported a Seoul National University Hospital thyroid cancer CT validation signal with an Artifact Reduction Score of 0.94. This result should be treated as a preliminary or internal validation claim until fully documented with protocol details, cohort description, comparative baseline, reader-study design, statistical analysis, and independent reproducibility checks.

Item	Description
Clinical domain	Thyroid cancer CT
Institution	Seoul National University Hospital collaboration context
Reported result	Artifact Reduction Score: 0.94
Interpretation	Promising artifact-reduction signal requiring expanded validation and documentation
Next evidence step	Reader study, baseline comparison, task-specific performance, and multi-site reproducibility

8.2 Validation Metrics

- **Image-quality metrics:** PSNR, SSIM, RMSE, CNR, NRMSE, noise power spectrum, local variance.
- **Clinical task metrics:** Lesion visibility, low-contrast detectability, reader preference, ROC/AUC for lesion detection.

- **Structure preservation metrics:** Lesion boundary consistency, anatomical continuity, HU consistency, edge preservation.
- **Trust metrics:** Uncertainty calibration, residual error, artifact-risk flags, out-of-distribution indicators.
- **Workflow metrics:** Runtime, throughput, repeatability, QC-report completeness, integration latency.

8.3 Evidence Generation Plan

- Retrospective technical validation using paired LDCT and standard-dose CT comparators where available.
- Reader study with radiologists focused on low-contrast lesion detectability and artifact burden.
- Multi-site reproducibility using scanner/vendor/protocol diversity.
- Downstream AI impact study evaluating CAD/segmentation/quantification performance before and after QMLIR.
- Regulatory documentation package aligned to intended use, risk controls, cybersecurity, and software lifecycle management.

9. Patient State Space and QLDI

QMLIR is the near-term reconstruction engine. QLDI is the long-term diagnostic intelligence layer. QLDI stands for Quantum Longitudinal Diagnostic Intelligence. It is based on the concept that patient diagnosis is not a single-image problem, but a time-dependent state-space problem.

9.1 Patient State Space Definition

Patient State Space is the multidimensional representation of a patient's condition composed of imaging findings, lesion characteristics, biomarkers, pathology, genomics, treatment history, and clinical observations across time.

At time t, a simplified patient state may be represented as:

$S(t) = \{I(t), L(t), B(t), P(t), G(t), C(t)\}$

where I = imaging, L = lesion features, B = biomarkers, P = pathology, G = genomics, and C = clinical observations.

9.2 Why State Space Becomes the Core Problem

As CT, PET, MRI, pathology, biomarkers, genomics, and clinical history accumulate, the number of possible patient states and disease trajectories grows rapidly. Conventional AI often analyzes one study at a time; QLDI aims to model how the patient state evolves across time.

9.3 QLDI Architecture

QLDI Layer	Function
Layer 1: QMLIR Reconstruction	Generate high-fidelity LDCT outputs and preserve subtle structures.
Layer 2: QPINN Trust Layer	Estimate uncertainty and verify physical consistency.
Layer 3: QPR Pattern Engine	Perform ROI-level refinement and hidden-pattern discovery.
Layer 4: Longitudinal Intelligence	Reassess prior imaging, integrate multimodal evidence, and forecast progression.

9.4 Key Future Functions

- **Retrospective lesion reassessment:** Re-examine prior CT at suspicious locations using new evidence from current imaging or clinical data.
- **Adaptive converged detection:** Update diagnostic confidence as CT, PET, biomarkers, pathology, and clinical history converge.
- **Progression forecasting:** Estimate potential growth, metastasis risk, and follow-up urgency.
- **Metastatic trajectory modeling:** Model disease transitions across organs such as lung and liver, which are common metastatic destinations.

10. System Architecture and Deployment Model

Figure 2 summarizes the current high-level QMEDIC system architecture. It includes a web frontend, FastAPI backend, asynchronous workers, queue-based processing, DICOMWeb/PACS integration, local/S3 artifact storage, model inference, and output management. This architecture supports research deployment, validation workflows, and future regulated commercialization.

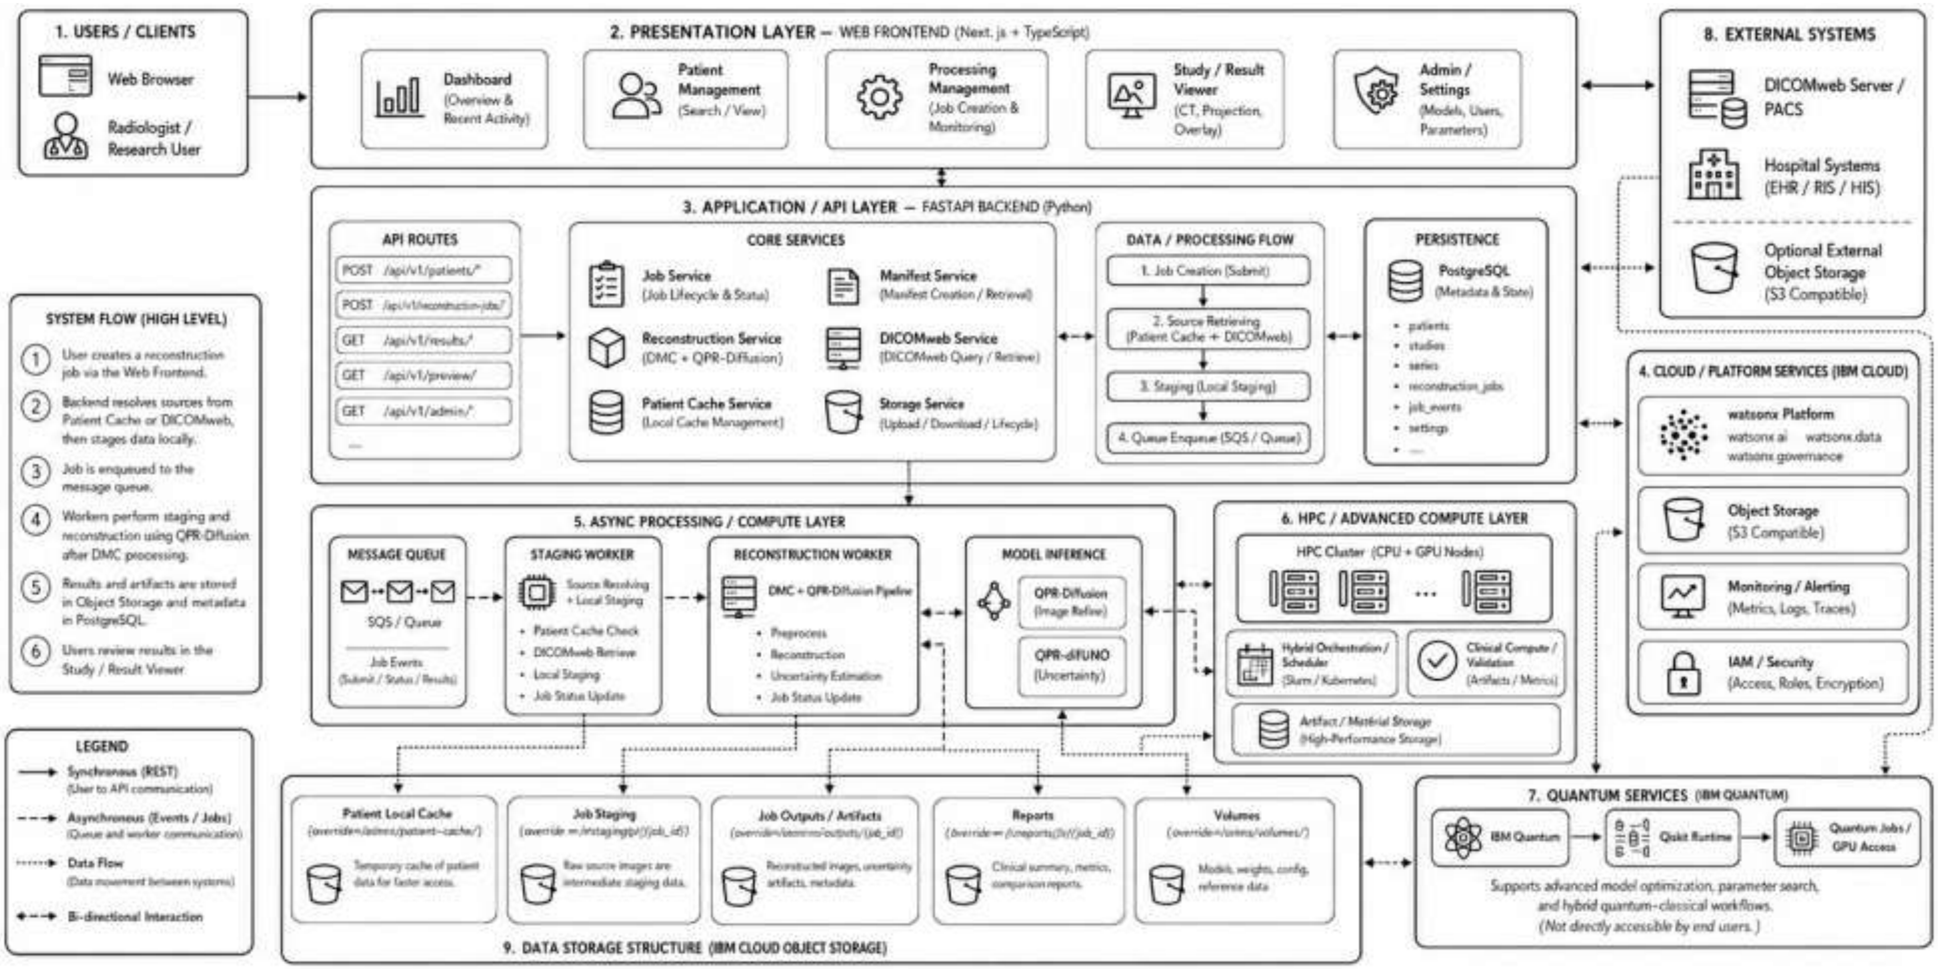

Figure 2. High-level QMEDIC system architecture. The system is designed around web workflow management, FastAPI services, DICOMWeb integration, queue-based workers, QPINN/QPR/Diffusion inference, artifact storage, and clinical review outputs.

10.1 Architecture Interpretation

- **Frontend:** Dashboard, patient management, processing management, study/result viewer, and admin/settings.
- **Backend:** FastAPI services for patients, reconstruction jobs, results, preview, DICOMWeb, manifests, and storage.
- **Workers:** Asynchronous processing for staging, source resolving, reconstruction, QMC, QPR, Diffusion, uncertainty, and clinical comparison.
- **Inference:** QPR-Diffusion for image refinement and QPINN/UFNO for uncertainty and trust assessment.
- **Storage:** Patient local cache, job staging, outputs, previews, reports, volumes, and manifests.
- **External integration:** DICOMWeb/PACS, optional S3/object storage, and monitoring/alerting.

10.2 Deployment Modes

Mode	Use Case	Customer Type
Research SaaS/API	Validation, cohort processing, algorithm evaluation	Research hospitals, academic centers, imaging labs
Private cloud / hybrid	Secure institutional deployment with controlled data movement	Clinical research networks, public-health organizations
On-premise appliance	High-control environments with strict data residency	Large hospitals, government healthcare systems
OEM / vendor integration	Embedded reconstruction or downstream AI enablement	Imaging vendors and PACS/AI platforms

11. Quantum-HPC and Medical Quantum Data Center Strategy

The long-term quantum thesis for QMEDIC is not that every CT image should be processed entirely on a quantum computer. The stronger thesis is that quantum and quantum-inspired systems may be useful for selected high-dimensional subproblems: patient state-space exploration, probabilistic inference, quantum pattern recognition, and optimization across complex multimodal trajectories.

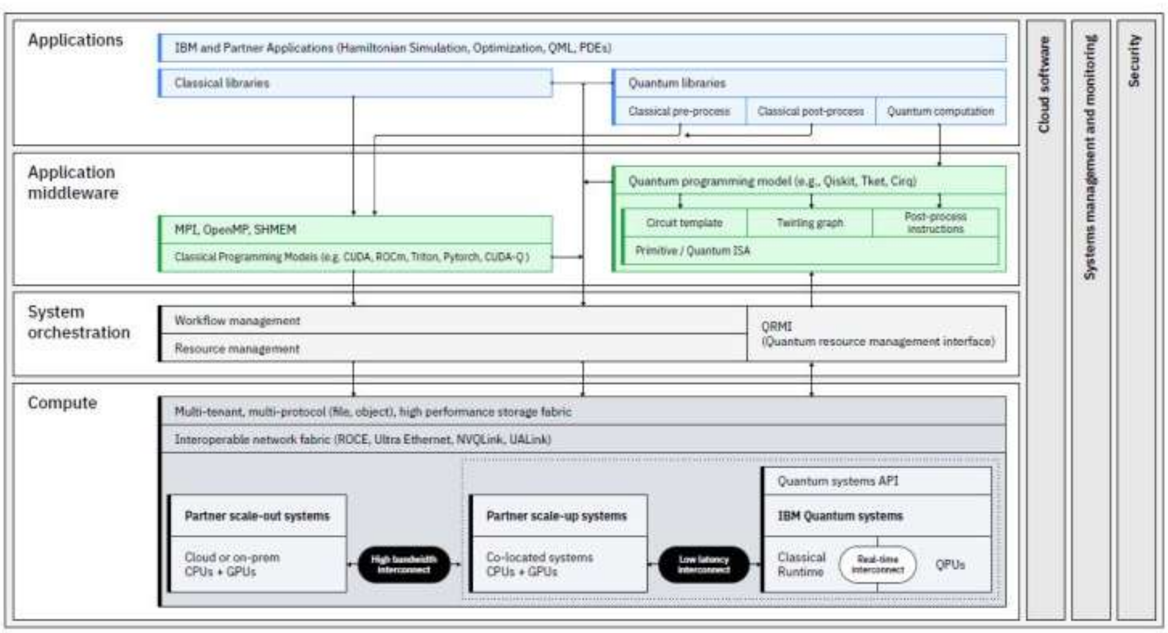

Figure 3. Quantum-centric supercomputing architecture concept. QMEDIC can align with a future Quantum-HPC model that combines classical compute, quantum programming, orchestration, cloud software, and system management.

11.1 Why Quantum Matters for QLDI

The primary computational challenge is not image reconstruction alone. The larger challenge is probabilistic exploration of high-dimensional longitudinal patient state space. As multimodal and longitudinal data accumulate, QLDI must reason over possible hidden states, disease transitions, and future trajectories.

- **High-dimensional feature learning:** QMC/VQC approaches can encode and explore complex distributions in latent space.
- **Probabilistic inference:** Quantum-inspired methods may support efficient sampling and uncertainty-aware reasoning.
- **Optimization:** Quantum-HPC orchestration may support resource allocation, parameter search, and trajectory optimization.
- **Medical quantum data center:** In the future, hospitals or national imaging networks may require specialized infrastructure for longitudinal patient modeling at population scale.

11.2 Near-Term Practicality

The near-term implementation should remain hybrid. Classical GPU/HPC systems handle bulk reconstruction and inference, while quantum or quantum-inspired components are evaluated for targeted modules such as QPR feature extraction, compressed latent-space modeling, and selected state-space inference tasks.

12. Expected Benefits for Patients, Physicians, and Healthcare Systems

12.1 Patient Benefits

- **Lower radiation exposure potential:** Clinically practical LDCT can reduce exposure in screening and follow-up settings.
- **Earlier detection potential:** Improved low-contrast detectability and retrospective reassessment may surface subtle findings earlier.
- **Reduced repeat imaging:** Better confidence in LDCT images may reduce unnecessary repeat scans in some workflows.
- **More personalized monitoring:** QLDI can support time-based disease monitoring and risk updates.

12.2 Physician and Radiologist Benefits

- **Improved image fidelity:** Enhanced confidence in subtle structures and low-contrast findings.
- **Better AI inputs:** CAD, segmentation, quantification, and annotation tools can work on higher-quality reconstructed images.
- **Uncertainty-aware workflow:** Confidence maps, residual maps, and artifact-risk flags support safer review.
- **Longitudinal intelligence:** Prior CT studies can be reassessed when new evidence suggests a hidden lesion or progression pattern.

12.3 Healthcare System Benefits

- **Screening scalability:** Improved LDCT quality can support broader screening programs.
- **Operational efficiency:** API/SaaS deployment and structured QC outputs can reduce workflow friction.
- **Evidence generation:** Validation-ready outputs and audit metadata support clinical research and regulatory preparation.
- **Strategic platform value:** QMLIR can become an enabling layer for downstream medical AI and OEM integration.

13. FDA 510(k) and Regulatory Readiness Strategy

QMEDIC should pursue a staged regulatory-readiness strategy. The initial commercial path should be research-use deployment and service sales while the company builds clinical validation data, quality-system maturity, cybersecurity documentation, software lifecycle controls, and predicate-device analysis for a future FDA 510(k) strategy.

A 510(k) strategy will depend on intended use, risk classification, predicate selection, performance claims, and whether the product is positioned as reconstruction software, image-processing software, CAD support, or a broader diagnostic support system. This section is a planning framework and is not regulatory advice.

13.1 Regulatory Phasing

Timeline	Phase	Major Activities
0-24 months	Research-use service and clinical validation preparation	RUO / non-diagnostic deployments; retrospective validation; reader studies; cybersecurity and QMS preparation
18-36 months	FDA 510(k) preparation and submission planning	Intended-use definition; predicate analysis; performance testing; software documentation; pre-submission as appropriate
24-48 months	Regulatory clearance and reference-site expansion	US/Canada/Asia reference sites; clinical adoption studies; reimbursement and enterprise procurement preparation
Year 5+	Enterprise and OEM scale	VA, IDNs, national imaging networks, and OEM/vendor integration

13.2 FDA Readiness Package

- Intended use and indications-for-use statement.
- Predicate device analysis and substantial equivalence rationale where applicable.
- Software architecture, hazard analysis, cybersecurity plan, and version-control documentation.
- Clinical performance validation with representative datasets and reader studies.
- Image-quality, task-performance, robustness, repeatability, and uncertainty calibration evidence.
- Post-market monitoring plan, update governance, and AI lifecycle-management strategy.

13.3 Claim Discipline

During the validation period, QMEDIC should avoid over-claiming diagnostic autonomy. Safer claims include image-quality support, reconstruction fidelity, low-contrast detectability support, artifact-risk flagging, and radiologist review support. Diagnostic or disease-progression claims should be introduced only when supported by sufficient evidence and regulatory strategy.

14. Commercialization Roadmap

The commercialization plan should sequence risk, evidence, and customer maturity. The strongest path is not immediate broad hospital sales, but staged adoption: research-use revenue first, clinical validation and FDA readiness in parallel, reference sites after clearance, and enterprise/OEM scale once evidence and regulatory position are stronger.

14.1 Phase 1: Research and Clinical Validation Revenue (Years 0-2)

- Sell QMEDIC as research-use software/service to domestic and international research institutes, clinical research groups, imaging laboratories, and innovation-friendly hospitals.
- Focus on LDCT reconstruction, artifact suppression, image-quality evaluation, and AI-enabling workflows.
- Generate revenue while collecting validation evidence for FDA 510(k) preparation.
- Develop core reference datasets and early customer case studies without clinical-diagnostic claims.

14.2 Phase 2: Regulatory Clearance and Reference Build-Up (Years 2-5)

- Prepare and pursue FDA 510(k) clearance according to intended use and predicate strategy.
- Build clinical reference sites in the United States, Canada, and Asia.
- Target leading academic medical centers, imaging research networks, and specialty oncology or screening programs.
- Develop reimbursement, procurement, integration, and cybersecurity readiness for enterprise healthcare customers.

14.3 Phase 3: Enterprise and OEM Scale (Year 5+)

- Target large-scale customers such as the Veterans Health Administration (VA), integrated delivery networks, national screening programs, large imaging-center chains, and enterprise hospital networks.
- Develop strategic OEM and vendor partnerships with medical-device and imaging companies such as Philips, GE HealthCare, Siemens Healthineers, and Canon Medical.
- Offer software licensing, cloud-based reconstruction service, API usage, enterprise subscriptions, and OEM-embedded modules.
- Expand QLDI into longitudinal diagnostic intelligence and Medical Quantum Data Center infrastructure.

14.4 Business Model

Model	Revenue Structure	Timing
Research SaaS/API	Per-study or annual subscription	Years 0-2
Clinical validation service	Paid validation projects and co-development	Years 0-3
Enterprise clinical license	Annual site license or volume-based pricing	Years 2-5
OEM licensing	Embedded module / revenue share	Year 5+
QLDI platform subscription	Longitudinal intelligence and patient-state analytics	Year 5+

15. Risks and Mitigation

QMEDIC's opportunity is large, but the path includes technical, clinical, regulatory, adoption, and market risks. A strong white paper should acknowledge these risks and describe mitigations rather than presenting the platform as risk-free.

Risk	Description	Mitigation
Technical variability	Performance may vary across scanners, protocols, anatomy, and dose levels.	Multi-site validation, protocol stratification, conservative modes, quality thresholds.
Clinical evidence gap	Image metrics may not translate into diagnostic impact.	Reader studies, task-specific outcomes, lesion-detectability metrics, clinical partner review.
Regulatory uncertainty	Claims may shift classification or require more evidence.	Early regulatory counsel, pre-submission planning, claim discipline, QMS preparation.
Workflow adoption	Radiologists may not adopt outputs that disrupt workflow.	PACS/DICOMWeb integration, QC reports, radiologist-in-the-loop design.
Quantum overclaim risk	Reviewers may challenge unsupported quantum-advantage claims.	Position quantum selectively for high-dimensional subproblems and maintain classical fallback.
Data access	Projection/sinogram access can be limited by vendors.	Support multiple modes: projection-enabled, DICOM-enabled, partner/OEM integration.

15.1 Conservative Deployment Principle

QMEDIC should maintain conservative operating modes and fallback pathways. If uncertainty, residual error, or artifact-risk indicators exceed thresholds, the system should preserve baseline outputs, flag the case, and route the finding for expert review rather than presenting aggressive reconstruction as definitive.

16. Conclusion

QMEDIC-QMLIR provides a focused near-term product thesis: make LDCT clinically practical by improving image fidelity at the reconstruction stage while preserving subtle anatomical structures and supporting downstream AI post-processing. This reconstruction-first strategy differentiates QMEDIC from conventional post-processing AI and positions it as an enabling layer for LDCT commercialization.

QPINN and QPR extend the platform from image quality improvement into trust and intelligence. QPINN provides physics-informed confidence, uncertainty, and hallucination-risk control. QPR provides targeted patch refinement and quantum pattern recognition for subtle structures and longitudinal findings.

QLDI extends QMEDIC beyond image reconstruction into longitudinal diagnostic intelligence. By modeling Patient State Space over time, QLDI can support retrospective reassessment, multimodal convergence, disease progression forecasting, and future Medical Quantum Data Center infrastructure.

The proposed commercialization path is staged and realistic: research-use revenue and clinical validation in the first two years, FDA 510(k) preparation and reference-site build-up over the next two to three years, and enterprise/OEM deployment from year five onward. This creates a coherent path from technical innovation to regulated healthcare impact.

QMEDIC is therefore not simply a CT image enhancement product. It is a foundation for future patient-centered diagnostic intelligence built upon fidelity-preserving reconstruction, physics-informed trust, and longitudinal patient state-space modeling.

Appendix A. Suggested Zenodo Metadata

Metadata Field	Suggested Entry
Title	QMEDIC-QMLIR and Quantum Longitudinal Diagnostic Intelligence: A Fidelity-Preserving Reconstruction Framework for Low-Dose CT, Patient State Space Modeling, and Quantum-Enabled Clinical Intelligence
Creators	Brandon Y. Kim; Youngje (Lucas) Son; SQK Inc.
Description	Technical white paper describing QMEDIC-QMLIR, QPINN, QPR, Patient State Space, QLDI architecture, clinical validation strategy, FDA 510(k) readiness, and commercialization roadmap.
Keywords	LDCT; CT reconstruction; QMLIR; QMEDIC; QPINN; QPR; QLDI; quantum AI; medical imaging; Patient State Space; SaMD; FDA 510(k)
Version	1.0
Resource type	Publication - Working paper / Technical report
License	CC BY 4.0 or company-approved open-access license
Funding / Acknowledgment	To be completed by SQK before submission.
Related identifiers	Optional: link to QMEDIC PCCT Quantum AI paper, company website, and future arXiv/DOI records.

Appendix B. Validation Metric Framework

Metric Category	Candidate Metrics	Purpose
Reconstruction improvement	RMSE, residual error, PSNR, SSIM, CNR, NRMSE	Measure improvement over baseline reconstruction.
Structure preservation	Anatomical continuity, lesion boundary consistency, HU consistency, edge preservation	Confirm preservation of clinically meaningful structures.
Low-contrast detectability	Reader visibility score, CNR, L-detectability, ROC/AUC	Assess detectability under low-dose conditions.
Uncertainty calibration	Confidence map, entropy score, calibration error, conformal prediction coverage	Identify ambiguous or low-confidence regions.
Artifact-risk detection	OOD score, residual map, artifact flag, hallucination-risk label	Detect AI-generated distortion or unreliable enhancement.
Physics consistency	CT projection/sinogram residual, forward-model residual, acquisition metadata consistency	Check alignment between source data and output image.
Workflow reproducibility	Run-to-run variation, repeatability score, version trace	Measure stability under repeated execution.
Operational performance	Latency, throughput, compute cost, cloud cost per study	Support SaaS/API deployment readiness.
Reporting completeness	QC report, audit metadata, model version, manifest trace	Support clinical validation and regulatory documentation.

Appendix B-1. Minimal Clinical Validation Study Design

- Define target indication and intended use, such as LDCT reconstruction support for selected anatomy and protocol ranges.
- Select baseline comparators: vendor reconstruction, vendor DLIR, and existing post-processing approaches where applicable.
- Use paired LDCT/standard-dose CT or clinically accepted reference images when available.
- Perform blinded reader study focused on lesion visibility, artifact burden, confidence, and diagnostic acceptability.
- Report subgroup performance by scanner, dose level, anatomy, reconstruction kernel, slice thickness, and lesion size.

Appendix C. Reference Sources

[1] SQK Inc. QMEDIC White Paper Draft v0.5, 2026. Internal/uploaded draft used as source material for reliability-centered positioning, cross-domain guardrail framework, metrics, architecture, and commercialization strategy.

[2] Son, Y.; Yu, H.W.; Dang, S.; Kim, B.; Park, J. "QMedic: Photon-Counting CT Image Super-Resolution Using a QuantumGAN with Measurement-Based Compilation." Proceedings of ACDSA 2026. Uploaded manuscript used as source material for QGAN, QUKKOS, PCCT super-resolution, and quantum-classical medical imaging research context.

[3] SQK Inc. QMEDIC System Architecture v11. Uploaded system architecture diagram used as source material for deployment, worker, DICOMWeb, artifact storage, QPR-Diffusion, and QPINN-UFNO workflow description.

[4] Quantum-Centric Supercomputing Architecture diagram. Uploaded architecture figure used as conceptual reference for Quantum-HPC, system orchestration, quantum middleware, and future Medical Quantum Data Center narrative.

[5] U.S. Food and Drug Administration. Premarket Notification 510(k) overview and 510(k) submission resources. Referenced for general regulatory framing; not a substitute for regulatory counsel.

[6] U.S. Food and Drug Administration. Software as a Medical Device and AI-enabled device software lifecycle guidance resources. Referenced for planning-level SaMD and AI lifecycle considerations; not a substitute for regulatory counsel.