


— REFERENCE ARCHITECTURE

Trustworthy Healthcare AI


A reference architecture for agentic AI in clinical workflows — grounded, verified, and auditable.
Proven in HealthQCopilot.

Swipe through the pipeline →




Grounded. Verified. Auditable.

Every encounter flows through specialized agents, real-data grounding, a verification gate, and a clinician. Governance throughout: consent, PHI redaction, tool RBAC, audit (HIPAA-ready).




Before any model runs

A patient encounter is captured; consent is checked and PHI is redacted before a single prompt leaves the boundary.




A planner, not one big model

The planner routes each task to a narrow agent — triage, clinical docs, prior-auth, care gaps — each easier to predict and verify.




Anchored in real data

Agents ground in the patient's EHR/FHIR record and clinical guidelines via RAG. Every claim must cite a source it retrieved.



Cite-or-refuse

A grounding guard checks each claim against the record; a clinical critic re-reads the reasoning. Unsupported output is dropped, not shown.



Reviewed, then delivered

Grounded, verified output is composed and logged to the audit trail — and a clinician reviews before anything is acted on.

— YOUR TURN

Building healthcare AI you can trust?

This is the pattern behind HealthQCopilot — narrow agents, mandatory grounding, a verification gate, and HIPAA-ready governance, on Azure.

DM me “assessment”

Atiqul Islam

Fractional Cloud & AI Architect for SMBs · atiqulislam.dev