

The Healthcare AI Playbook Inside VITAL Case Studio

book
io
nce.

Gemini Notebook

A convincing output can still be wrong.

Perfect formatting does not equal clinical fidelity. One reversed statement changes the meaning of an entire note. We must look beneath the surface to find hidden, fabricated facts.




A controlled environment to know what breaks.

Live AI in the Wild


- No live patient connections.
- No autonomous chart filing.
- No microphone or transcription service.
- No generative model hallucinations.

The VITAL Sandbox

- An interactive education environment.
- Deterministic engines.
- Fixed synthetic fixtures.
- A space to explore, challenge, and plan.



The 5-Level Evidence Laboratory.




Level 1: Assemble the Case.

Start with the work you want to improve. Select a clinical discipline and establish the baseline.



18 workflows across the healthcare spectrum.




Selecting the play: Ambient Clinical Notes

A clinician has finished a fictional visit. The source includes symptoms, explicit negative history, an unconfirmed assessment, and a follow-up plan.



The boundaries of a successful draft.



Every included statement has a source.

Negative history is preserved.

The clinician remains responsible for signing.


Accountable Sponsor:
CMIO + Clinical Service Lead

Level 2: Explore the Workflow.

Open the black box.
Watch how the deterministic
engine processes clinical data
into a drafted note.



The deterministic engine at work.



1. Extract the facts

2. Compose the template

3. Check against the source

Six sources. One drafted note.




Input: 6 synthetic
encounter sources

Transform:
Deterministic engine

Output: Draft ready
for human review

The draft is ready. The evidence is not.

The platform has generated the output. Now, we must prove its fidelity before it reaches a clinician.



Level 3: Change the Scenario.

A system that works perfectly under routine conditions may fail when the facts change. We must test the boundaries.



Testing across three conditions.

Routine Case	Changed Conditions	Control Failure
Goal: Establish reference.	Goal: Explore a trade-off.	Goal: Catch a defect.

Changing one variable reveals the true resilience of the workflow.

The Baseline: Routine execution.

The source meaning is preserved. The draft successfully carries forward explicit negative history without hallucination.



Flipping the scenario: Triggering a failure.



We adjust the configuration to insert an affirmative diagnosis that contradicts an explicit denial in the source.

Catching the unsupported diagnosis.

The scenario injection successfully forced a control failure.
An affirmative diagnosis was fabricated against the
original transcript.




Level 4: Verify the Source.

Human accountability at every decision.
When the system proposes a fact,
a clinician must be able to trace it.



Making the reasoning visible.

A careful reviewer checks negative statements, uncertainty, and the follow-up plan against the raw encounter facts.



Tracing the defect to a halt.

The accountable reviewer manually traces the output. Because the fabricated fact cannot be sourced, the draft is rejected.



The Stop Rule.



Pause the workflow if:

1. A material fabricated fact reaches a signed note.
2. Consent controls fail.
3. Patient identity is mismatched.

The clinician remains responsible.


AI is an engine for extraction and composition. It is not an autonomous actor.

A clinician verifies the draft against the encounter before signing.



Level 5: Report the Evidence.

Translating synthetic lessons into real-world pilot actions, value models, and governance gates.




Carry the reasoning forward.

Leave with a printable case report documenting your exact settings, the review rationale, and explicit stop conditions for the pilot.



The Value Machine: Calculating Net Capacity.




Defining the break-even threshold.

Year-one net value:

\$12,280

Break-even point:

8.0 months




Assumptions are illustrative.
Added work for negative net capacity is fully
costed at 100% applied realization.

The Evidence Gate Record: Preparation to Validation.



The Evidence Gate Record: Pilot to Deployment.



Level 5 Preparation & Validation

- Define Intended Use
- Establish baseline
- Clear data route
- Set stop criteria
- Validate local cases
- Rehearse failure

Level 4 Pilot & Execution

- Prepare people
- Approve controlled pilot
- Measure unintended effects
- Scale or stop decision
- Operationalize monitoring
- Verify value



Before you deploy, know what breaks.

An interactive education environment to run the case, challenge the result, and define what your next pilot needs to prove.



One discipline. Make the reasoning visible.



Explore AI workflows
in a safe sandbox.



Inspect the evidence behind
every generated claim.



Keep humans accountable
at every clinical decision.

VITAL Case Studio. The Healthcare AI Playbook.