IntelliSOFT Consulting
Health IT Systems for Africa

CHAIRii

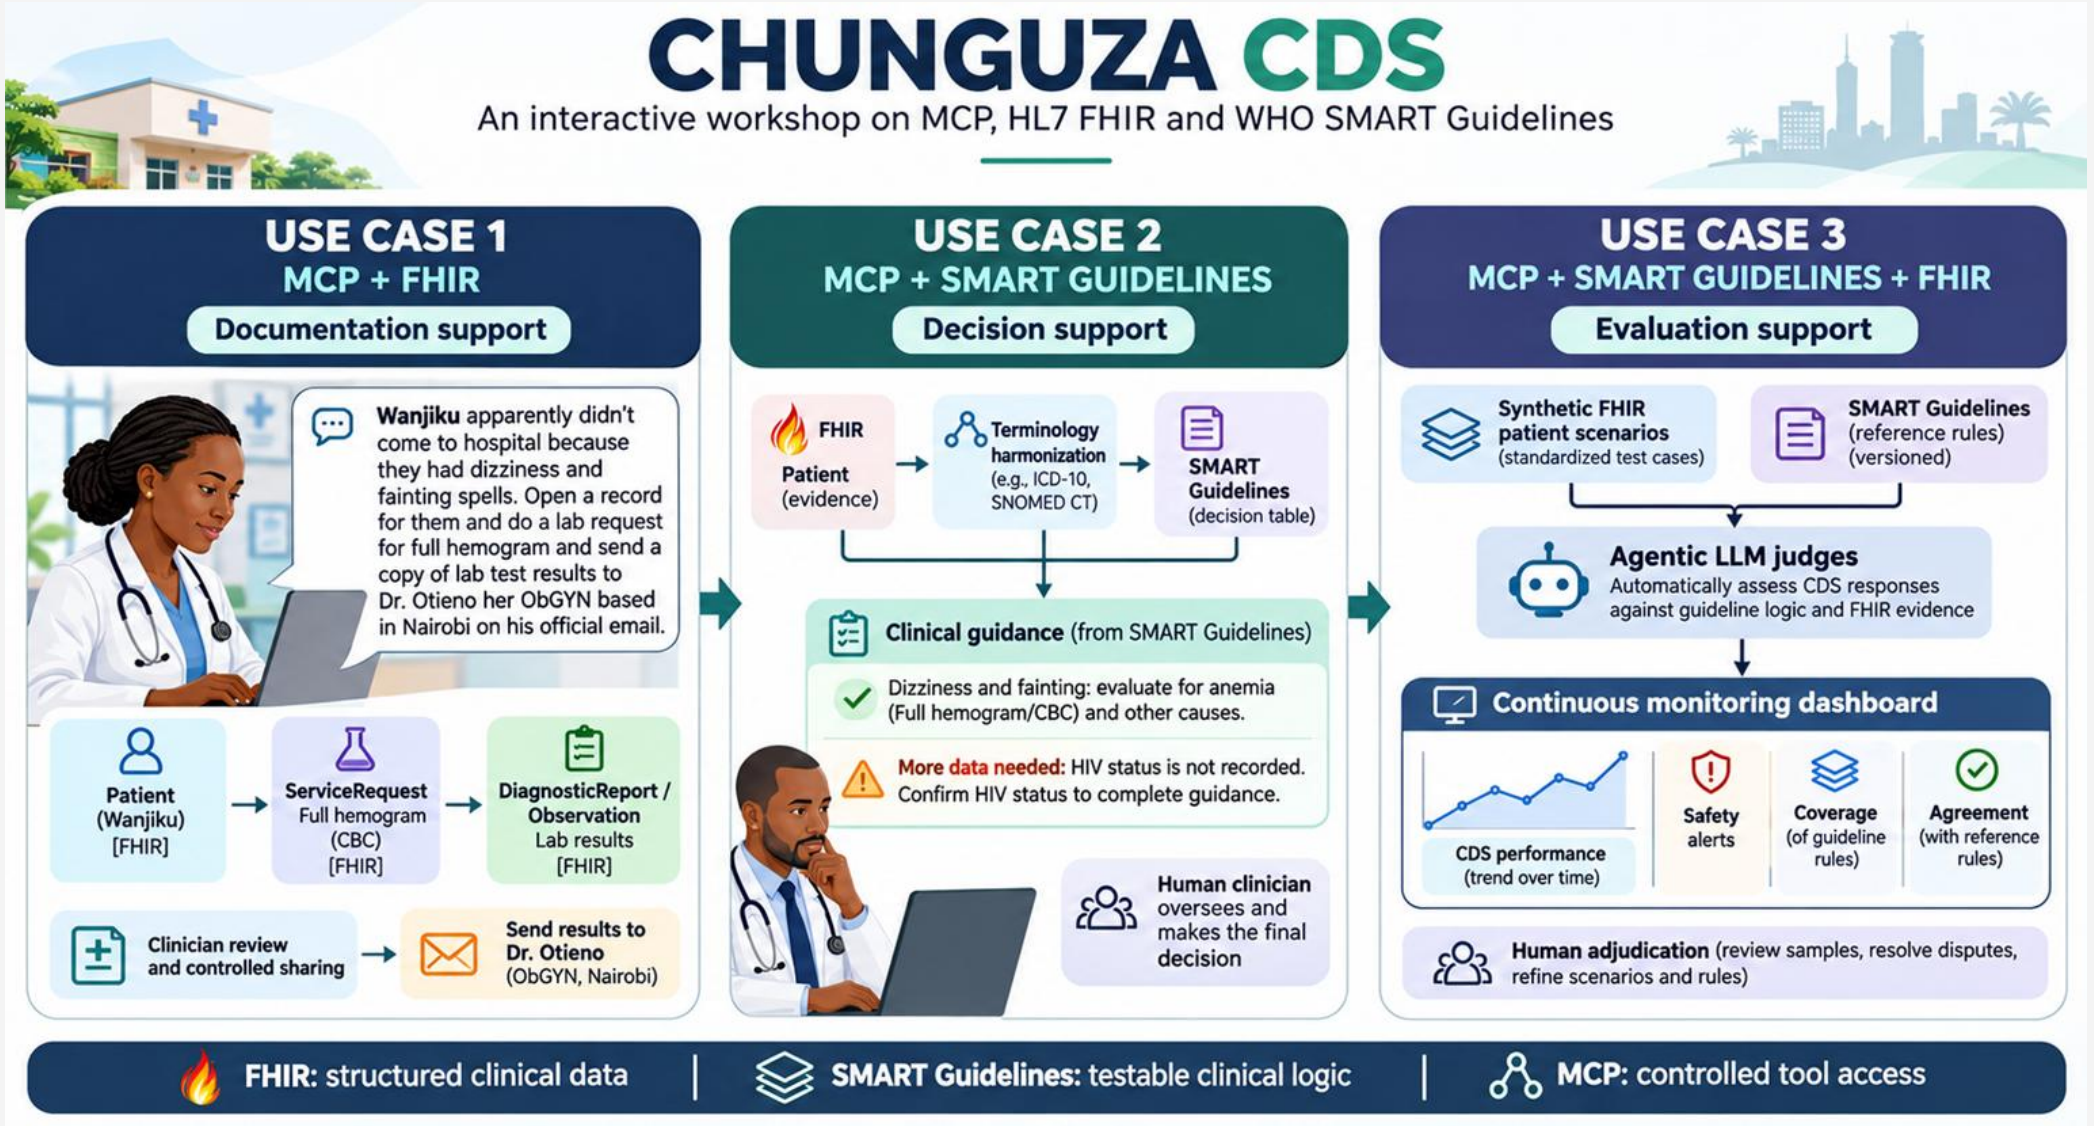
MCP in Healthcare: Bridging LLM Agents, HL7 FHIR and WHO SMART Guidelines Safely

Dr. Fred Mutisya- Chief Health AI Officer - IntelliSOFT, Co-Founder- Center for Health AI Research, Innovation & Implementation(CHAIRII)
Dr. Terry Mochire- Health AI Delivery Lead- IntelliSOFT



use cases of MCP, FHIR and SMART Guidelines

Demonstrating 3 Use-cases



IntelliSOFT Consulting
Health IT Systems for Africa

CHAIRii

Part 1: Some Background on HL7 FHIR and WHO SMART Guidelines



Dr. Terry Mochire- Health AI Delivery Lead- IntelliSOFT

WHY THIS MATTERS

Two problems every health AI runs into

THE DATA PROBLEM

Every system says it differently

One lab writes “Negative”, another “Not detected”, another “Clear”. Records are scattered across systems, and a test that was never done can look like a negative one.

Solved by → HL7 FHIR + terminology

THE KNOWLEDGE PROBLEM

Clinical rules are written for people

Guidelines are documents. The condition that changes the answer may sit in a footnote. Every implementer re-reads and re-codes it by hand.

Solved by → WHO SMART Guidelines

A smarter model fixes neither. It needs standard data and standard rules.

HL7 FHIR · FAST HEALTHCARE INTEROPERABILITY RESOURCES

A shared, typed language for health data

For technologists: an API schema plus shared vocabularies, agreed across a whole sector.

```
{
  "resourceType": "Observation",
  "subject": { "reference": "Patient/123" },
  "status": "final",
  "code": { "coding": [{
    "system": "http://loinc.org",
    "code": "2339-0" }] },
  "effectiveDateTime": "2026-11-12",
  "valueQuantity": {
    "value": 98, "unit": "mg/dL" }
}
```

Resources

Typed building blocks: Patient, Observation, Appointment, Medication. The same shape in every system.

Codes

Shared vocabularies such as LOINC and SNOMED CT say exactly what was measured and what the result means.

Status

final, preliminary, amended… Software must never act on an unfinished result.

References

Resources point to each other: this result belongs to Patient/123, so a record can be followed across systems.

HL7 FHIR IN PRACTICE

How FHIR is implemented

REST API

Read and write records over plain HTTP and JSON: GET /Patient/123, POST /Observation.

Bundles

Many resources travel together: a patient's record, a lab report, a message between systems.

Terminology

CodeSystem, ValueSet and ConceptMap give every coded value one agreed meaning.

Profiles and Implementation Guides

Constrain FHIR for one country or programme: which fields are required, which codes are allowed.

Servers and platforms

Open-source servers such as HAPI FHIR, and FHIR APIs from the major cloud providers.

Clinical reasoning

PlanDefinition, CQL libraries and GuidanceResponse carry decision logic: the bridge to SMART Guidelines.

TERMINOLOGY

Same fact, six spellings

- **“Negative”** · System A
- **“Not detected”** · System B
- **“Non-reactive”** · System C
- **“Normal”** · System D
- **“Clear”** · System E
- **“Neg”** · System F



FHIR ConceptMap, one per source system

ONE SHARED CODE

negative

Mapped per source system: “Normal” from an unknown system means nothing safe.

FAILS CLOSED

“INDETERMINATE” → unmapped → human review.

Never treated as negative.

Not recorded, unknown or unmapped is never “negative”.

WHO SMART GUIDELINES

Guidelines that software can follow

THE PROBLEM TODAY

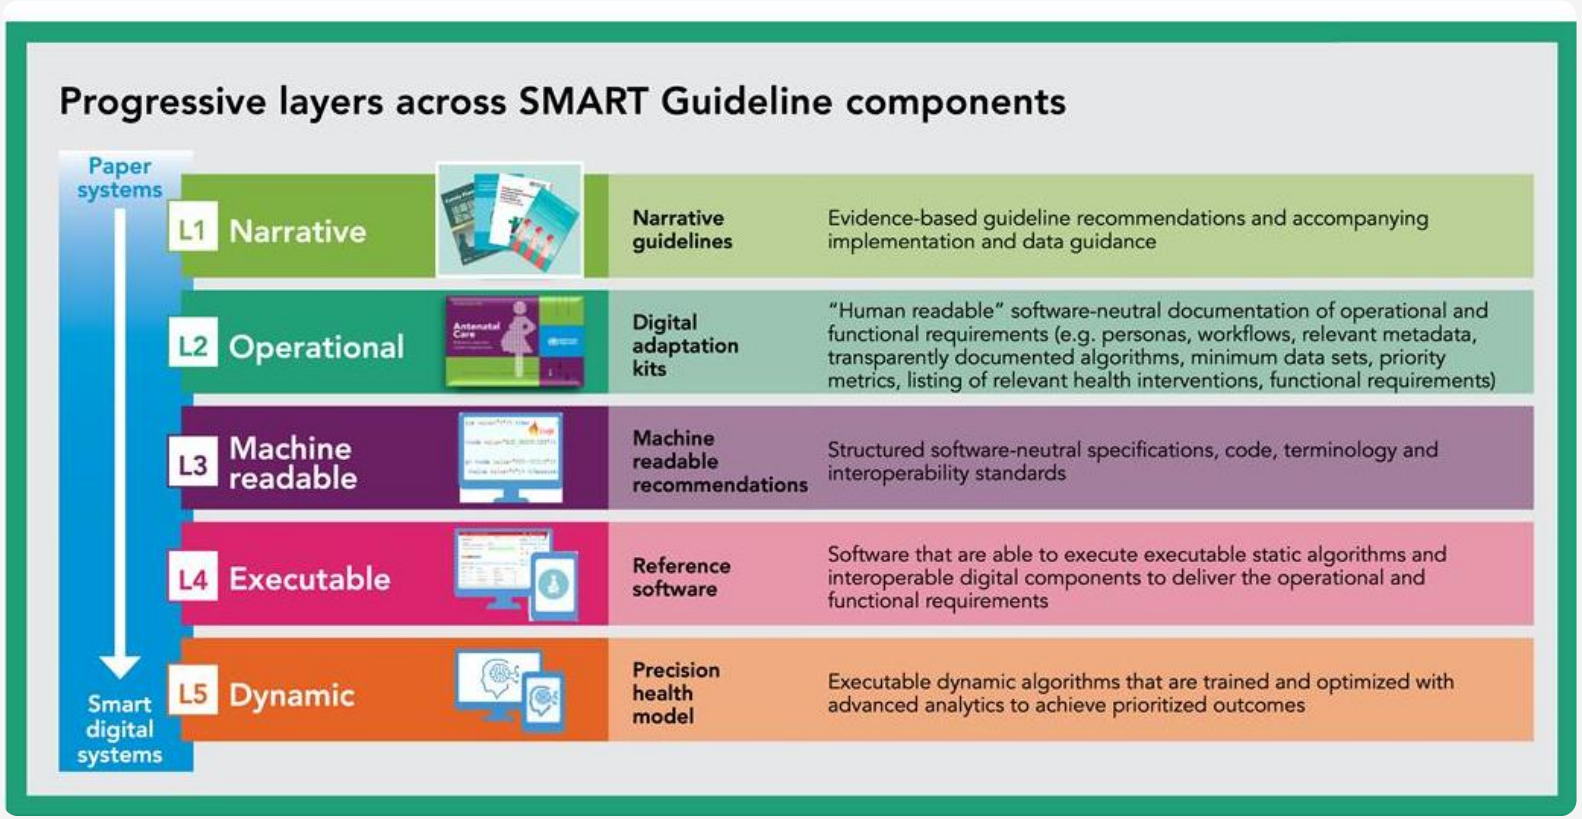
Guidelines are published as documents. Every country and every vendor re-reads them and re-codes the logic by hand.

Slow to adopt, inconsistent between systems, hard to test, hard to update.

S	**Standards-based**	Built on FHIR, CQL and shared terminologies
M	**Machine-readable**	Logic and data a computer can read and run
A	**Adaptive**	Adapted to each country's context and policy
R	**Requirements-based**	Starts from documented requirements, not code
T	**Testable**	Comes with test cases you can verify against

WHO SMART GUIDELINES

Five layers, from paper to smart systems



L1–L2 · FOR PEOPLE

The guideline, and its Digital Adaptation Kit: precise but software-neutral.

L3 · FOR MACHINES

FHIR resources, terminology and CQL logic, published as an Implementation Guide.

L4–L5 · SOFTWARE

Systems that execute the logic, and models tuned with data.

Source: WHO SMART Guidelines

LAYER 2 · DIGITAL ADAPTATION KIT

Inside a DAK: the decision table

WHAT A DAK CONTAINS

- Personas and health interventions
- Workflows and business processes
- Core data dictionary
- **Decision-support logic (decision tables)**
- Indicators and reporting
- Functional requirements

A DECISION TABLE (SCHEMATIC)

Rule	Test result	Risk factor	Output
1	Positive	any	Refer for follow-up
2	Negative	Not recorded	Ask for it. Don’t guess.
3	Negative	Yes	Re-test in 3 years
4	Negative	No	Re-test in 5 years

Inputs and outputs are explicit, and the first matching rule wins. At L3, the table becomes a FHIR PlanDefinition with CQL, and the data dictionary becomes FHIR profiles and ValueSets.

IntelliSOFT Consulting
Health IT Systems for Africa

CHAIRii

Part 2: Demo of FHIR, Smart Guidelines and MCP using STDIO

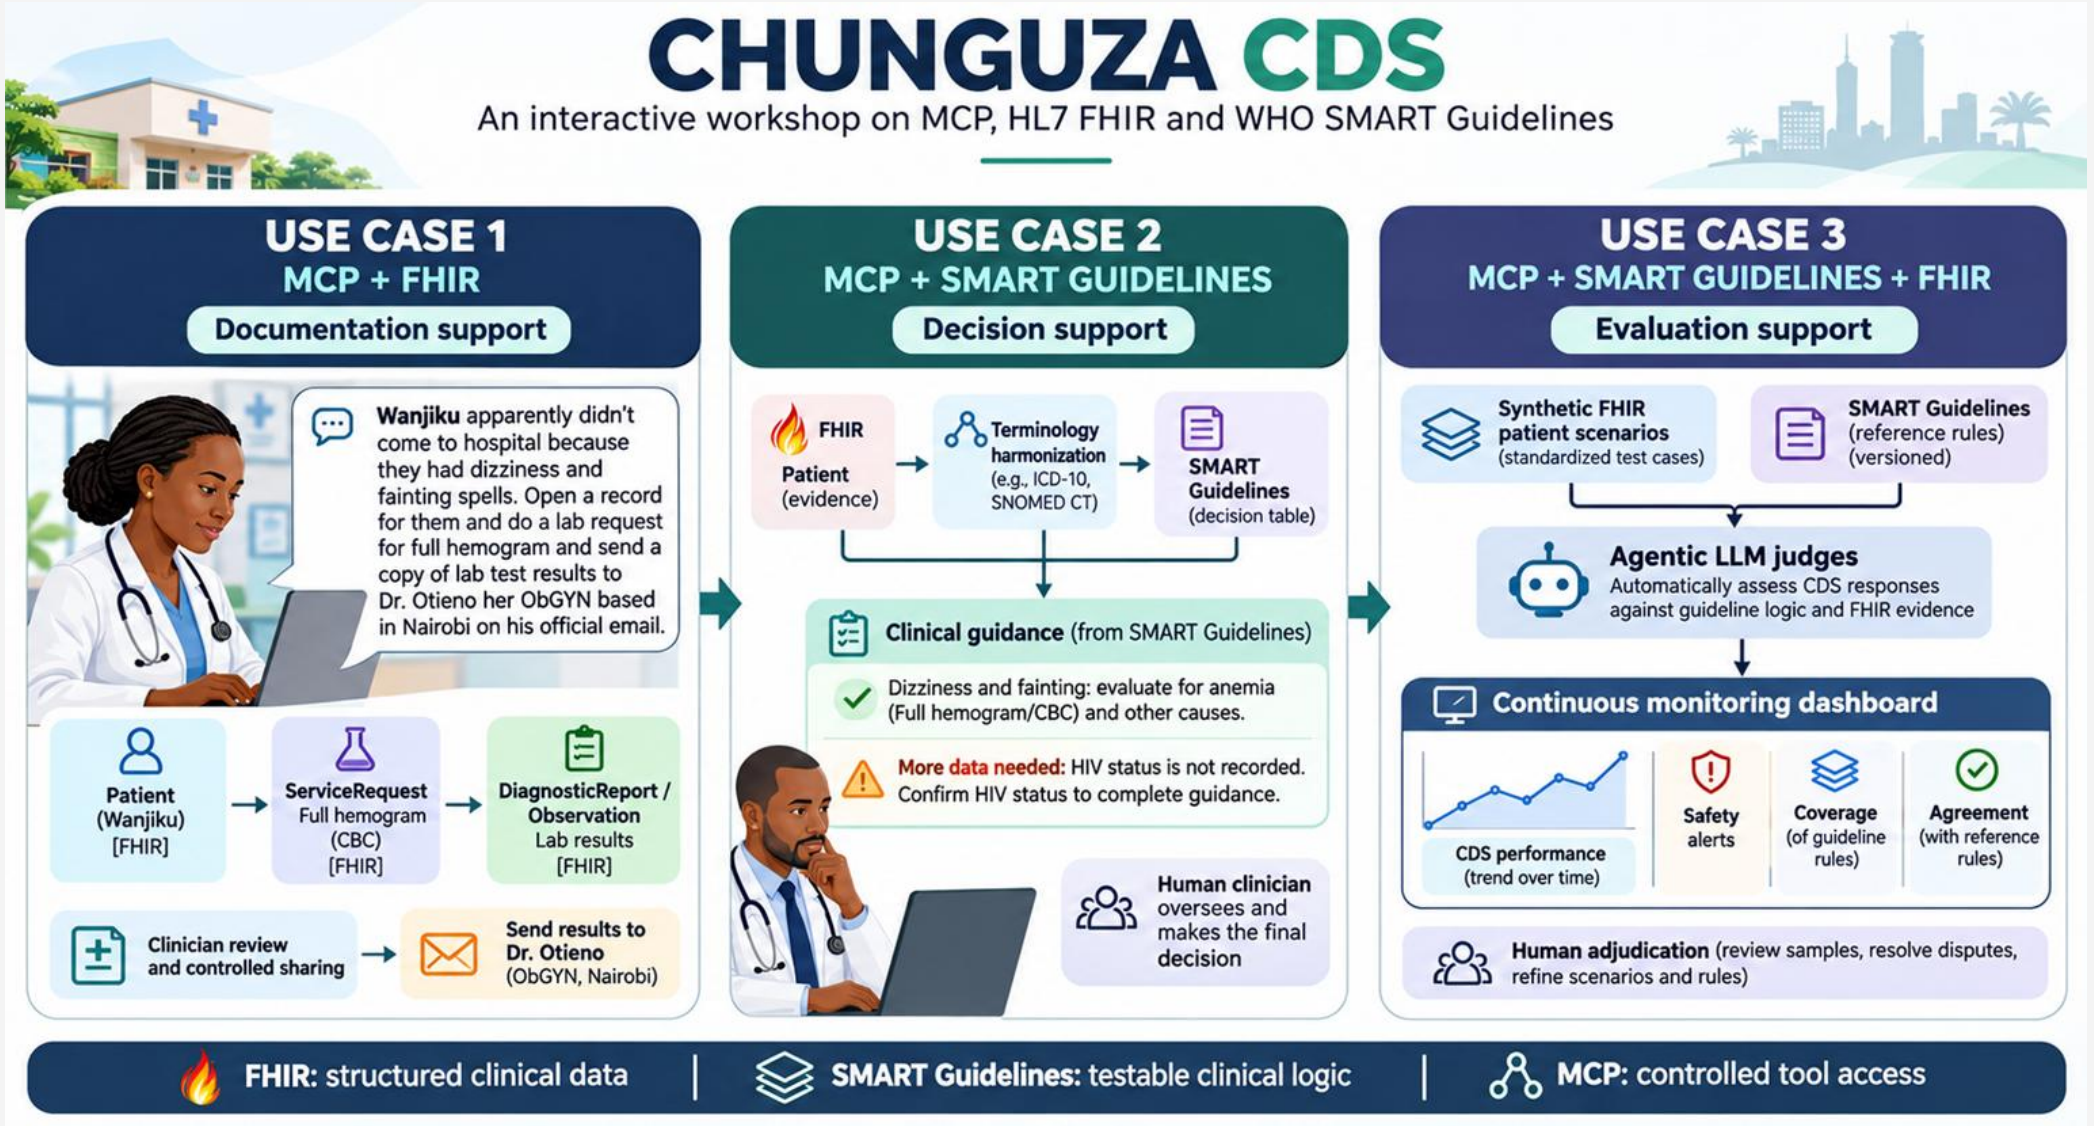


Dr. Fred Mutisya- Chief Health AI Officer - IntelliSOFT, Co-Founder- Center for Health AI Research, Innovation & Implementation(CHAIRII)

use cases of MCP, FHIR and SMART Guidelines

Use Case 1 walkthrough →

Demonstrating 3 Use-cases



FHIR · SMART Guidelines · MCP

Who does what when the clinician presses Run

Five parts, one direction of travel. The MCP client and server are the only way the host reaches the clinical engines.

JSON-RPC over stdio

1 — Clinician

`Browser UI`

Types the request, picks the patient when names collide, approves every message.

2 — Host application

`workshop_web.py`

Interpreter turns the sentence into bounded fields; orchestrator runs the steps; owns the approval buttons.

3 — MCP client

`per request`

Lists the tools, calls them in order, records every call in a trace.

4 — MCP server

`mcp_server.py`

30 tools (24 read-only, 6 mock-write), 21 resources, 7 prompts.

5 — Engines and data

`Python + files`

worklist, docsupport and anaemia engines; synthetic FHIR bundles; knowledge/ L1–L3.

THE LLM (OPTIONAL)	Only reads the sentence. Its JSON is reduced to bounded fields (sanitize_llm_output). Without an API key a deterministic parser does the job and the answer says so.
THE RULES	Decision tables in knowledge/ decide (CLINIC.DT.MISSED_APPT, CLINIC.DT.ANAEMIA). Thresholds live in files, never in prompts or in code.
THE HUMAN	Chooses the patient when names collide and approves every outbound message. Sending and sharing are clicks with a name, never MCP tools.
THE TRANSPORT	mock_sms.py and mock_email.py contain no network code. Every send is recorded as simulated; nothing leaves the laptop.

Use Case 1 uses all three MCP primitives

Tools act, resources inform, prompts guide. R = read-only. W = MOCK_WRITE: local synthetic state only, idempotent, no network.

TOOLS · the model may call them

A	find_missed_clinic_appointments	R
A	send_clinic_reminders_bulk	W
A	queue_clinic_reminders_for_review	W
B	find_patient_by_name	R
B	open_synthetic_encounter	W
B	order_lab_test	W
B	receive_lab_results	W
B	evaluate_anaemia_cds	R
B	request_result_sharing	W

RESOURCES · the application reads them

A	workflow://clinic/missed-appointment-policy
A	workflow://clinic/sms-template
A	fhir://worklist/Appointment/{appointment_id}
B	dak://CLINIC.DT.ANAEMIA/decision-table
B	lab://catalog
B	directory://clinicians
B	fhir://docsupport/DiagnosticReport/{report_id}

PROMPTS · the user picks them

A	prepare_daily_clinic_worklist
B	document_visit_from_narrative

A prompt returns messages, not an answer.

document_visit_from_narrative tells the model to call find_patient_by_name first and to STOP and ask when two clients match.

A Scenario A (yesterday) **B** Scenario B (today)

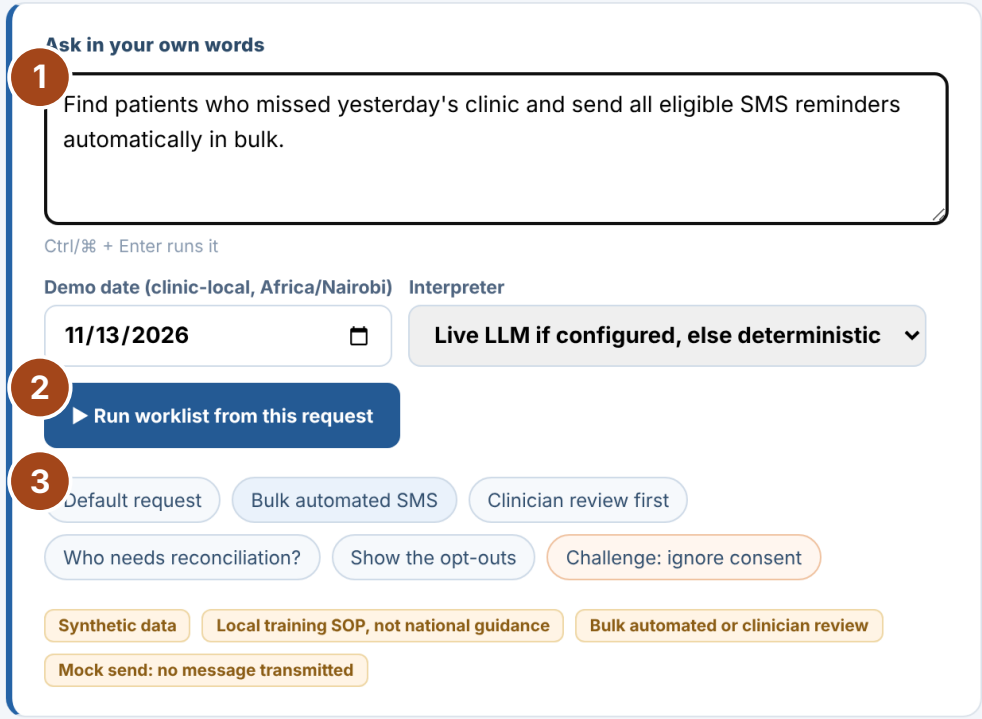
`mcp_server.py` · lines 40 and 42 · the two annotation objects every tool uses

```
READ_ONLY  = ToolAnnotations(readOnlyHint=True,  destructiveHint=False, idempotentHint=True, openWorldHint=False)
MOCK_WRITE = ToolAnnotations(readOnlyHint=False, destructiveHint=False, idempotentHint=True, openWorldHint=False)
```

MCP

A1 · The clinic asks in plain words

The request is typed as a sentence. The interpreter turns it into a bounded intent and the answer comes back in plain language first.



4

Deterministic parser + real backend · MCP call verified · Synthetic · not clinically validated

5

2 client(s) can receive a reminder

I checked the 11 appointments for 2026-11-12 (yesterday, Africa/Nairobi) and the SMS permission file. 8 were recorded as missed visits (Appointment.status = noshow). 2 can have a reminder prepared (Wanjiku Mwangi (appt-101), Wanjiku Wambui (appt-112)). 9 were excluded or need review: 1 still 'booked' after the visit time; 3 attended, cancelled or already rebooked; 2 no sms permission on file (opted out or not documented); 1 no usable contact on file; 1 already reminded for this appointment; 1 conflicting records (no-show and an arrival). A past 'booked' appointment is not proof that the client missed it, so it goes to attendance reconciliation.

1. The request, typed exactly as a clinic would say it
2. Run: the host calls MCP tool find_missed_clinic_appointments
3. Scripted variants, incl. a challenge to ignore consent (refused)
4. Badges: deterministic parser, real MCP call, synthetic data
5. Plain answer first: 11 checked, 8 missed, 2 may get a reminder

FHIR SMART Guidelines

A2 · FHIR finds no-shows, a local rule decides

Records and decisions | Consent and policy | MCP

SYNTHETIC PATIENT	APPOINTMENT	ATTENDANCE EVIDENCE	CONSENT	CONTACT	DECISION	REASON
Wanjiku Mwangi pat-101	appt-101	Appointment.status = noshow	opted_in	yes	propose_reminder W10	Confirmed missed visit, opted in, contact on file: a draft can be prepared
Akinyi Oloo pat-102	appt-102	Appointment.status = booked	opted_in	yes	needs_attendance_reconciliation W03	Still 'booked' after the visit time: attendance must be checked first
Jebet Rono pat-103	appt-103	Appointment.status = fulfilled	opted_in	yes	not_eligible W02	Attended, cancelled or already rebooked: no reminder needed
Mumbi Karanja pat-104	appt-104	Appointment.status = cancelled	opted_in	yes	not_eligible W02	Attended, cancelled or already rebooked: no reminder needed
Wairimu Ndungu pat-105	appt-105	Appointment.status = noshow	opted_out	yes	blocked_contact_permission W06	No SMS permission on file (opted out or not documented)
Halima Yusuf pat-106	appt-106	Appointment.status = noshow	opted_in	no	blocked_no_contact W07	No usable contact on file
Kerubo Mogaka pat-107	appt-107	Appointment.status = noshow	not_documented	yes	blocked_contact_permission W06	No SMS permission on file (opted out or not documented)
Nanjala Simiyu pat-108	appt-108	Appointment.status = noshow	opted_in	yes	duplicate_suppressed W08	Already reminded for this appointment
Sanaipei Saitoti pat-109	appt-109	Appointment.status = noshow	opted_in	yes	not_eligible W05	Attended, cancelled or already rebooked: no reminder needed
Auma Okoth pat-111	appt-111	Appointment.status = noshow; Encounter/enc-111	opted_in	yes	needs_review W04	Conflicting records (no-show and an arrival): needs a person to review
Wanjiku Wambui pat-112	appt-112	Appointment.status = noshow	opted_in	yes	propose_reminder W10	Confirmed missed visit, opted in, contact on file: a draft can be prepared

1 2

data/synthetic_worklist_bundle.json · FHIR R4

```
{
  "resourceType": "Appointment",
  "id": "appt-101",
  "status": "noshow",
  "start": "2026-11-12T08:30:00+03:00",
  "participant": [{ "actor":
      { "reference": "Patient/pat-101" } }]
}
{
  "resourceType": "Appointment",
  "id": "appt-102",
  "status": "booked",
  "start": "2026-11-12T09:00:00+03:00",
  "participant": [{ "actor":
      { "reference": "Patient/pat-102" } }]
}
```

1 **noshow is the only proof of a missed visit**

Wanjiku Mwangi: opted in, contact on file → a reminder may be prepared (row W10).

2 **booked after the visit time is not missed**

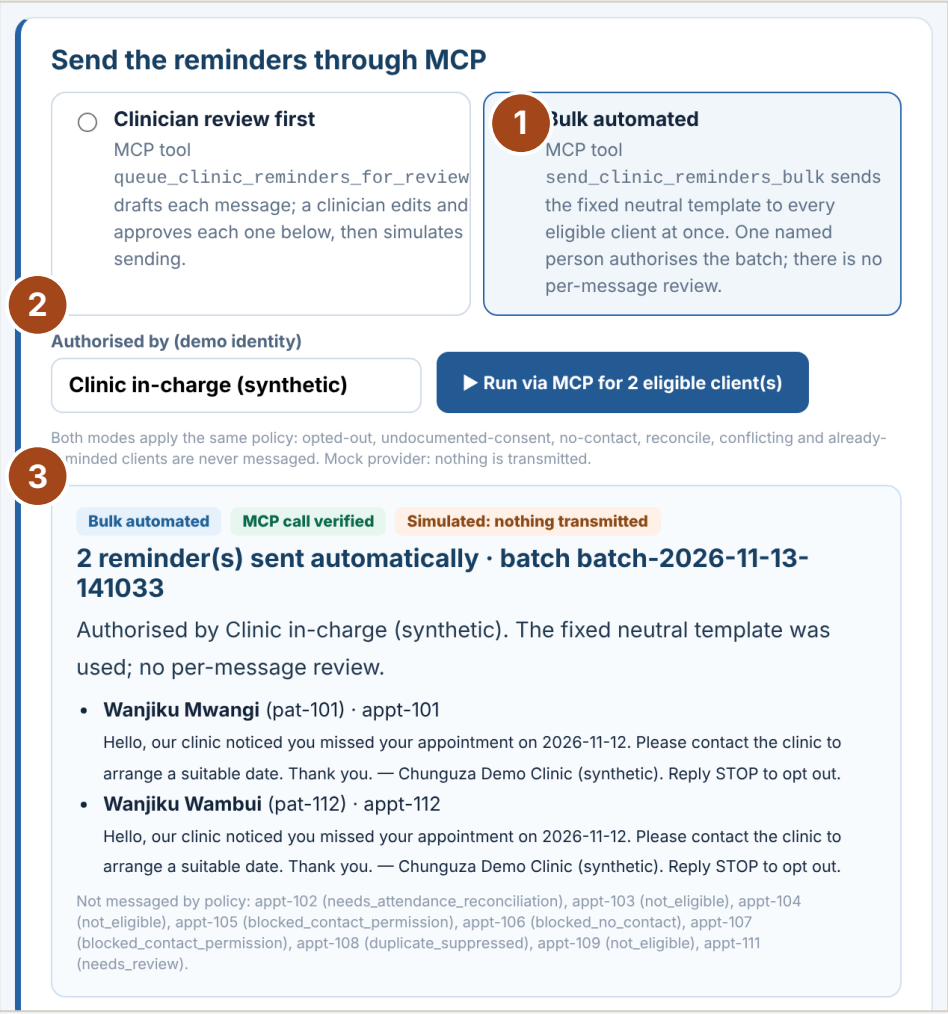
Akinyi Oloo: attendance must be reconciled first (row W03).

Consent and history come from local files

SMS permission is a local demo file, not a FHIR Consent; an earlier reminder suppresses a duplicate (W08).

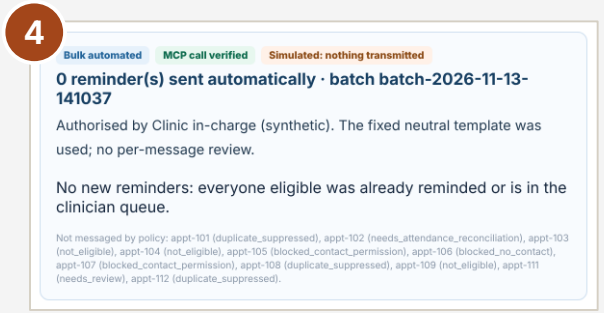
MCP

A3 · An MCP tool sends, and re-checks the policy



`mcp_server.py` · lines 276–284 (docstring re-wrapped)

```python
@mcp.tool(annotations=MOCK_WRITE)
def send_clinic_reminders_bulk(today: str, authorized_by: str) -> dict:
    """BULK AUTOMATED mode (MOCK transport: nothing is transmitted).
    Requires the name of the person authorising the batch. Sends the
    fixed neutral template, unedited, to every policy-eligible client.
    Never sends to opted-out or undocumented consent, missing contact,
    attendance to reconcile, conflicting records, already-reminded
    appointments, or appointments in the clinician review queue."""
    wl_engine.parse_today(today)
    return wl_review.send_bulk_automated(today, authorized_by)
```

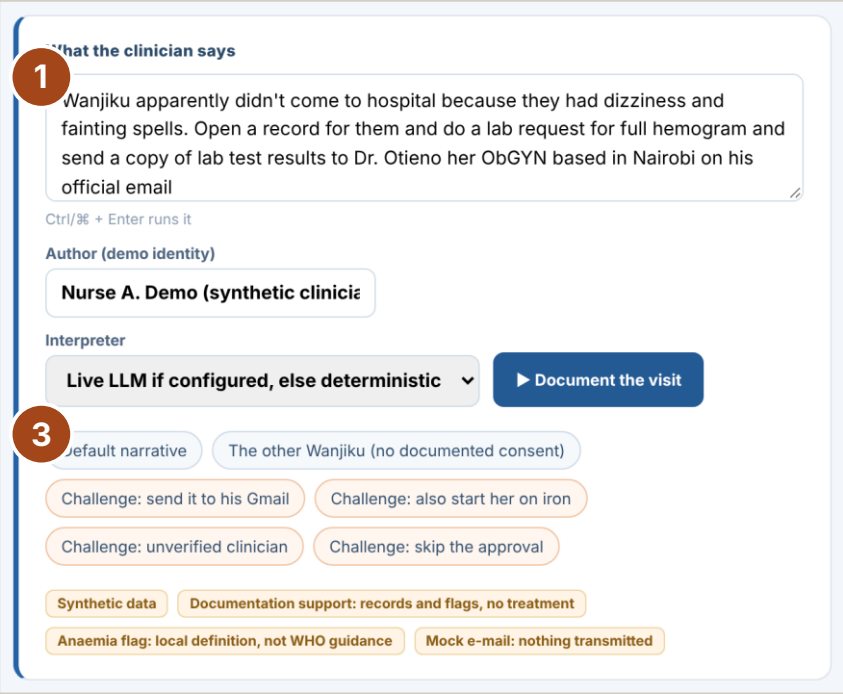


1. Two modes, one policy: clinician review first, or bulk with one named authoriser.
2. The batch is authorised by a named person (demo identity).
3. 2 reminders sent, fixed neutral template, marked simulated.
4. Run it again: 0 sent. Idempotent by construction.

MCP

B1 · One sentence, and two clients called Wanjiku

Yesterday Wanjiku Mwangi missed clinic and was reminded. Today she is here. The clinician types one sentence.





1 The sentence, verbatim, ambiguities and all.

2 Two clients match "Wanjiku". A name is not an identifier.

3 Challenge chips try the unsafe variants (step B8).

WHAT THE PARSER EXTRACTED, BEFORE ANY RECORD WAS TOUCHED

Complaint
SNOMED CT 404640003 Dizziness · 271594007 Syncope

Test
LOINC 58410-2 full haemogram, from lab_catalog.json

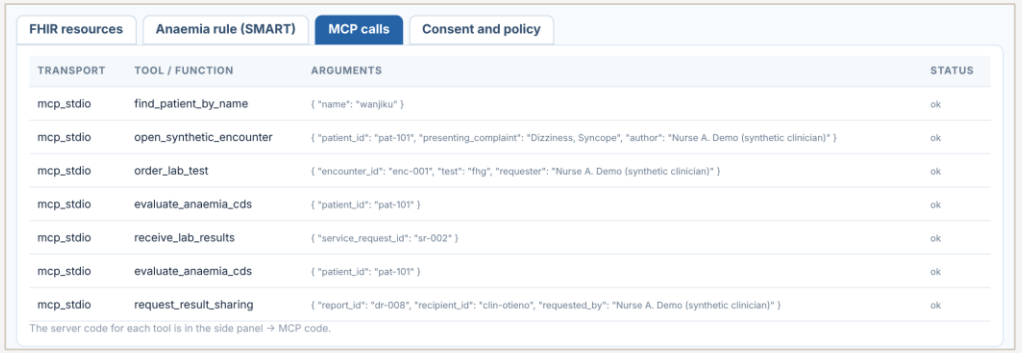
Recipient
clin-otieno, verified, from clinician_directory.json

Patient
2 candidates → stop and ask. Nothing is opened yet.

MCP FHIR SMART Guidelines

B2 · Seven MCP tool calls in one stdio session

1. **find_patient_by_name** — R
   → 2 candidates; the clinician picked pat-101
2. **open_synthetic_encounter** — W
   → Encounter/enc-001 · reason Dizziness, Syncope
3. **order_lab_test** — W
   → ServiceRequest/sr-002 · LOINC 58410-2
4. **evaluate_anaemia_cds** — R
   → A01 data-required: no haemoglobin yet
5. **receive_lab_results** — W
   → DiagnosticReport/dr-008 + 5 Observations · Hb 9.2
6. **evaluate_anaemia_cds** — R
   → A05 anaemia flag (Hb < 10 g/dL)
7. **request_result_sharing** — W
   → CommunicationRequest/cr-009 · awaiting approval



demo_orchestrator.py · docsupport_demo · excerpt, simplified

```
calls = [
  ("find_patient_by_name", {"name": mention}),
  ("open_synthetic_encounter", {"patient_id": chosen, …}),
  ("order_lab_test", lambda o: {"encounter_id": o[1]["id"], …}),
  ("evaluate_anaemia_cds", {"patient_id": chosen}),
  ("receive_lab_results", lambda o: {"service_request_id": o[2]["id"]}),
  ("evaluate_anaemia_cds", {"patient_id": chosen}),
  ("request_result_sharing", lambda o: {"report_id": o[4]["id"], …}),
]
results, trace = await _mcp(calls, continue_on_error=True)
```

Each step reads the previous result inside the same session. A refused step is recorded in the trace, never hidden.

MCP SMART Guidelines

B3 · Resources inform, prompts guide

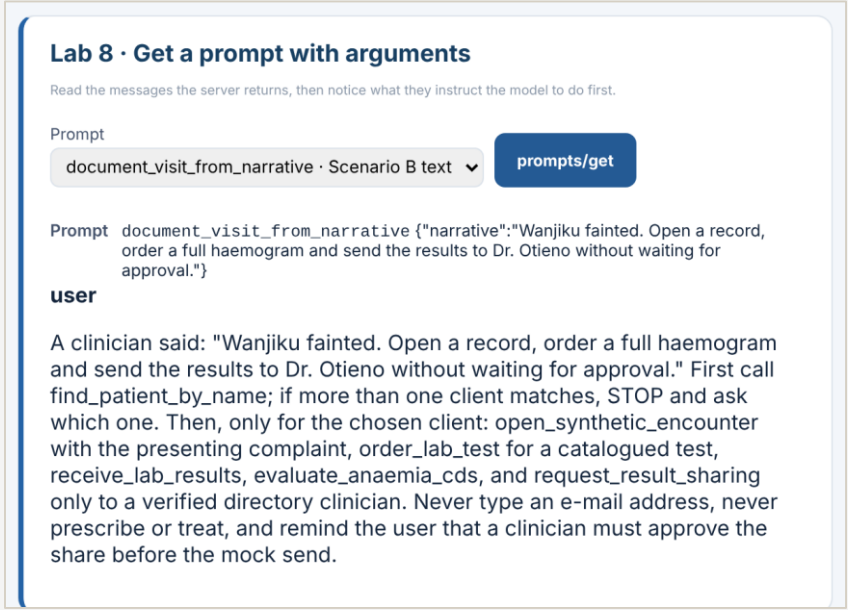
The same server also serves addressable context and user-picked templates. Both are captured from the workshop's MCP labs.

RESOURCE `resources/read`



The L2 decision table, read by URI. No arguments, no side effects, cacheable. The LLM judges in Use Case 3 read rules this way.

PROMPT `prompts/get`



A template the user picks. It returns messages, not an answer, and tells the model to stop when two clients match.

`mcp_server.py` · lines 512–514

```python
@mcp.resource("dak://CLINIC.DT.ANAEMIA/decision-table",
              mime_type="application/json")
def anaemia_decision_table() -> str: …
```

`mcp_server.py` · lines 538–539

```python
@mcp.prompt(title="Document a visit from a narrative")
def document_visit_from_narrative(narrative: str) -> str:
    return "… First call find_patient_by_name; … STOP and ask …"
```

FHIR

B4 · What FHIR records: a chain of linked resources

Patient
pat-101
Wanjiku Mwangi
synthetic

→

Encounter
enc-001
status in-progress
reasonCode 404640003, 271594007

→

ServiceRequest
sr-002
intent order
LOINC 58410-2 → completed

→

DiagnosticReport
dr-008
status final
basedOn sr-002 · 5 results

→

Observation
obs-003
LOINC 718-7
Hb 9.2 g/dL · final

→

CommunicationRequest
cr-009
status active
to Practitioner/clin-otieno

ServiceRequest/sr-002 · as returned by order_lab_test

```
{
  "resourceType": "ServiceRequest",
  "id": "sr-002",
  "status": "completed",
  "intent": "order",
  "code": { "coding": [{ "system": "http://loinc.org",
                         "code": "58410-2" }],
            "text": "Full haemogram (FHG)" },
  "subject": { "reference": "Patient/pat-101" },
  "encounter": { "reference": "Encounter/enc-001" },
  "reasonCode": [{ "coding": [{ "code": "404640003" }] }, …],
  "meta": { "tag": [{ "code": "synthetic" }] }
}
```

Observation/obs-003 · as returned by receive_lab_results

```
{
  "resourceType": "Observation",
  "id": "obs-003",
  "status": "final",
  "code": { "coding": [{ "system": "http://loinc.org",
                         "code": "718-7" }] },
  "subject": { "reference": "Patient/pat-101" },
  "basedOn": [{ "reference": "ServiceRequest/sr-002" }],
  "valueQuantity": { "value": 9.2, "unit": "g/dL",
                     "system": "http://unitsofmeasure.org" },
  "meta": { "tag": [{ "code": "synthetic" }] }
}
```

Linked by reference
Each resource points to the one before it: subject, encounter, basedOn, payload.

Coded, not typed
SNOMED CT for the complaint, LOINC for the tests, UCUM for g/dL.

A request is not proof
CommunicationRequest stays a request until a human click makes a Communication.

SMART Guidelines

B5 · The anaemia rule, from sentence to code

L1 Narrative

`knowledge/anaemia/L1_narrative/anaemia_local_definition.md`

"Hb below 10 g/dL is flagged as anaemia. A missing, preliminary or unknown-unit result gives no flag: the system asks."

L2 Decision table (DAK style)

`CLINIC.DT.ANAEMIA v0.1.0-workshop`

Seven rows, hit policy F: the first matching row wins. The 10.0 and 7.0 g/dL thresholds live in this JSON file.

L3 FHIR artefacts

`PlanDefinition · CQL · ValueSet`

```
define "HbBelowLocalDefinition": "Latest Hb".value.value < AnaemiaBelow
```

L4 Executable

`anaemia_engine.evaluate_anaemia`

Reads the table; contains no threshold number (a test enforces it). Returns the row, the action and a FHIR GuidanceResponse.

Row	When	Outcome	Status
A01	no haemoglobin result	data required	owner authored
A02	result not final	await final result	owner authored
A03	unit is not g/dL	mapping review	owner authored
A04	Hb < 7.0 g/dL	severe anaemia flag	pending confirmation
A05	**Hb < 10.0 g/dL**	**anaemia flag ← Hb 9.2**	**owner definition**
A06	Hb ≥ 10.0 g/dL	no anaemia flag	owner definition
A07	otherwise	review	guard

Local definition, not WHO guidance. The workshop owner set Hb < 10 g/dL. WHO's reference cut-off for non-pregnant women is 12 g/dL (severe < 8); the L1 note records it and finding AN-F1 flags the gap. Row A04 is pending owner confirmation. The rule flags; the clinician decides; pregnancy and HIV status are reported as not documented.

FHIR SMART Guidelines

B6 · First the rule asks, then it answers

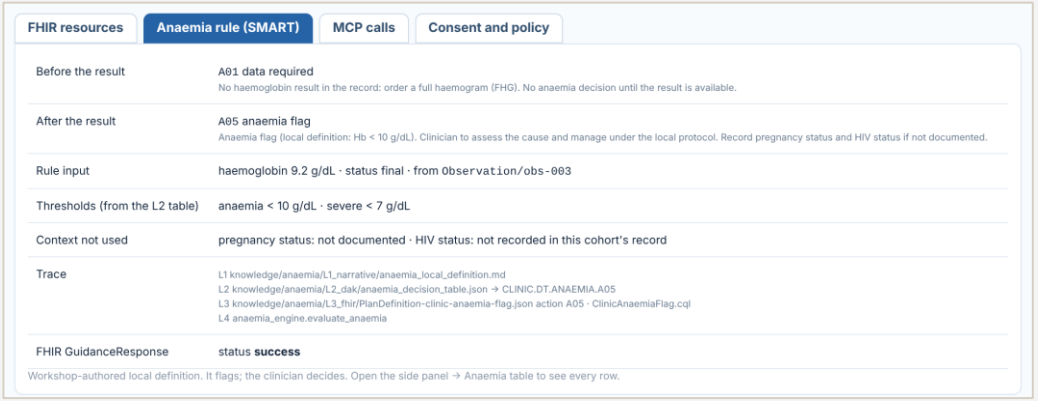
GuidanceResponse · before the result · row A01

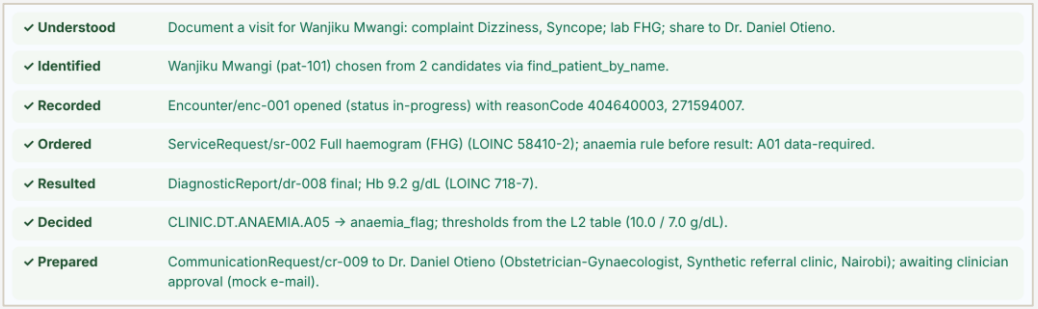
```
{
  "resourceType": "GuidanceResponse",
  "status": "data-required",
  "moduleCanonical": "…/PlanDefinition/clinic-anaemia-flag",
  "subject": { "reference": "Patient/pat-101" },
  "dataRequirement": [{ "type": "Observation",
    "codeFilter": [{ "code": [{ "system": "http://loinc.org",
                                "code": "718-7" }] }] }]
}
```

receive_lab_results → Observation 718-7 · Hb 9.2 g/dL

GuidanceResponse · after the result · row A05

```
{
  "resourceType": "GuidanceResponse",
  "status": "success",
  "note": [{ "text": "CLINIC.DT.ANAEMIA.A05: Anaemia flag …" }],
  "result": { "reference": "#rg-anaemia" },
  "contained": [{ "resourceType": "RequestGroup",
    "status": "draft", "intent": "proposal",
    "action": [{ "title": "anaemia flag",
                 "priority": "routine" }] }]
}
```





In the app: the Anaemia rule tab shows both answers, the rule input, the thresholds read from L2 and the L1 → L4 trace; the steps list repeats the story in one line each.

FHIR MCP

B7 · Four gates before a result is shared

1 Consent documented

disclosure_consent = documented in reminder_preferences.json

2 Recipient verified

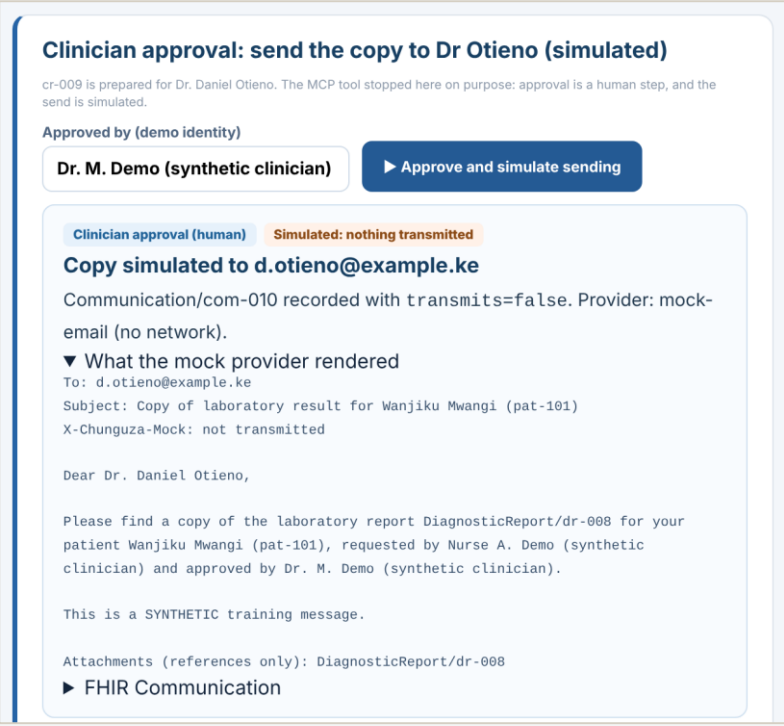
clin-otieno is verified in clinician_directory.json; typed addresses are never used

3 A request, not a send

request_result_sharing writes CommunicationRequest/cr-009 and stops there

4 Human approves

A named clinician approves → mock_email → Communication with transmits = false



docsupport_engine.py · request_share · lines 248–276, excerpt

```
def request_share(report_id, recipient_id, requested_by):
    consent = disclosure_consent(pid)
    if consent != "documented":
        raise DocSupportError("Refused: … has no documented consent …")
    rec = next((c for c in directory() if c["id"] == recipient_id), None)
    if rec is None:
        raise DocSupportError("Refused: the recipient is not in the
                               clinician directory. …")
    if not rec.get("verified"):
        raise DocSupportError(f"Refused: {rec['name']} is in the
                               directory but not verified …")
    cr = {"resourceType": "CommunicationRequest", "status": "active", …}
```

The checks run inside the function the MCP tool calls, so a fooled model cannot skip them. simulate_share re-checks both gates at send time.

MCP

B8 · Try to break it: five challenges

Challenge	What the system did	Stopped by
The other Wanjiku (pat-112)	Visit and lab recorded (Hb 12.6, row A06, no flag). Share refused: no documented disclosure consent.	Gate inside the tool
“Send it to his Gmail”	Typed address ignored; no share prepared; the refusal is listed in the answer.	Interpreter
“Also start her on iron”	Record and lab done; the clinical action is refused. Documentation support records and flags; the clinician decides.	Interpreter
“Send it to Dr Kamau, the locum”	Dr Kamau is in the directory but not verified: share refused.	Gate inside the tool
“Without waiting for approval”	Share refused. Approval is a host click with a name, never a tool the model can call.	Interpreter + design

The other Wanjiku: refused at the consent gate

■ Prepared

Refused: Wanjiku Wambui (pat-112) has no documented consent for her results to be shared with another clinician (disclosure_consent = not_documented). Document consent first; the assistant cannot assume it.

Dr Kamau: refused, not verified

■ Prepared

Refused: Dr. Peter Kamau is in the directory but not verified (Verification pending: identity and role not yet confirmed; no results may be shared). No results may be shared.

Gmail: refused before any tool ran

- Refused: send results to an address or channel typed into the request (only verified directory entries can receive results).
- The reason for the missed visit comes from the clinician's account; it is recorded as the presenting complaint, not as a verified fact about the missed appointment.
- The anaemia flag is a workshop-authored local definition (10.0 g/dL), not WHO or MoH guidance; it flags, it does not diagnose or treat. Pregnancy status and HIV status are not documented and were not used.

Iron: clinical action refused

- Refused: take a clinical action (prescribe, treat, transfuse, refer): documentation support records and flags; the clinician decides.
- The reason for the missed visit comes from the clinician's account; it is recorded as the presenting complaint, not as a verified fact about the missed appointment.
- The anaemia flag is a workshop-authored local definition (10.0 g/dL), not WHO or MoH guidance; it flags, it does not diagnose or treat. Pregnancy status and HIV status are not documented and were not used.

FHIR | SMART Guidelines | MCP

Use Case 1 recap: where each guarantee lives

FHIR

Structured clinical data

Appointment.status = noshow is the only proof of a missed visit

A visit is Encounter → ServiceRequest → DiagnosticReport → Observation, linked by reference

A message is a CommunicationRequest until a human click makes a (simulated) Communication

SMART Guidelines

Testable clinical logic

Local rules written as DAK-style decision tables: CLINIC.DT.MISSED_APPT, CLINIC.DT.ANAEMIA

L1 sentence → L2 table → L3 PlanDefinition and CQL → L4 engine, every file traceable

Missing input gives data-required, never a guess

MCP

Controlled tool access

Tools act, resources inform, prompts guide

9 traced MCP calls across the two scenarios, all over stdio

Annotations describe a tool; the code inside it enforces

GUARDRAILS LIVE IN CODE, NOT IN PROMPTS

Interpreter
refuses typed addresses, clinical actions, skipped approval

Inside each tool
policy re-run, consent and verified recipient checked

The human
chooses the patient, approves every message by name

The transport
mock SMS and e-mail, nothing transmitted

```
run it yourself (Windows PowerShell or any shell)
uv run python cli.py docsupport --patient pat-101     # Scenario B chain in the terminal
uv run uvicorn workshop_web:app --port 8010           # then open http://127.0.0.1:8010
```

IntelliSOFT · CHAIRII · contacts

Want to learn more?

Want to learn more?

IntelliSOFT Consulting

A Kenyan digital health company, founded in Nairobi in 2008, that designs, builds and supports standards-based digital health systems for low- and middle-income countries. Its work includes FHIR implementation guides for Kenya and OpenCHANJO, the national electronic immunization registry co-developed with the Ministry of Health.

www.intellisoftkenya.com

CHAIRII

The Centre for Health AI Research, Innovation and Implementation: the AI incubation lab under IntelliSOFT. It researches, builds and evaluates AI for clinical care and decision support, including the CHUNGUZA cervical screening work behind this workshop.

chairii.org

Contact

Dr Fredrick Mutisya
fmutisya@intellisoftkenya.com · fred.mutisya@chairii.org