



SEPT 22 - 24, 2026 | NAIROBI, KENYA

The Agentic Healthcare Enterprise: A New Operating System for Care



Nirmal Fernando

Director and Head of Engineering,
Solutions BU

What happens next?

A shared question for patients and the teams supporting them




Healthcare organizations - Present day




Patients & members


Care teams



Staff & partners



WSO2 Identity Platform



WSO2 API Platform



Patients & members

Care teams

Staff & partners



Healthcare Interoperability - Accelerates agent development

Why?

One shared foundation for all your agents.

01

Common data model

Healthcare domain is covered through the FHIR data model and LLMs understands FHIR already.

02

Prebuilt tools

For FHIR, there are prebuilt MCP servers, prebuilt knowledge bases that exists already. You do not have to write custom MCPs or build custom knowledge bases.

03

Reusable interfaces


FHIR integration work becomes a shared infrastructure instead of recurring work for each agent.




Patients & members


Care teams


Staff & partners



WSO2 OPEN HEALTHCARE SOLUTION

One platform for healthcare interoperability

FHIR APIS & SERVICES

FHIR APIs
FHIR server (Postgres based)
Terminology & validation
Bulk export, FHIRPath

INTEGRATION & MAPPING

HL7v2 v2.3 to v2.8
C-CDA to FHIR
X12 278/275 and FHIR
Epic, Cerner, athenahealth connectors

AI & AGENTS

FHIR MCP server
Healthcare aware Copilot for developers
Healthcare specific guardrails
Public knowledge base tools

IDENTITY & CONSENT

SMART on FHIR, OAuth 2.0
Granular consent control
Consent lifecycle & revoke
ATNA audit (IHE ITI-20)

REGISTRIES & EXCHANGE

Master patient index
OpenHIE layer
IHE PDQm, PIXm, ATNA
Provider registry

DEVELOPER TOOLING

Ballerina health tool
FHIR package generation
API & CDS Hooks templates
VS Code integration

WSO2 Agentic Enterprise Fabric

40+

healthcare libraries

15+

prebuilt services

Apache 2.0

open source, no lock-in



Fastest **open** **source** FHIR server

Highest throughput in CRUD,
import, search test suites ran by
an independent performance
benchmark.



<https://github.com/wso2/fhir-server>

Throughput ⓘ

5K req/s

4K req/s

3K req/s

2K req/s

1K req/s


0 req/s

Name	Mean
hapi	2.97K req/s
medplum	1.81K req/s
microsoft	359 req/s
wso2	4.64K req/s




Highest conformance to the FHIR standard

Highest conformance to the FHIR standard based on an independent conformance benchmark.



The Agentic Healthcare Enterprise - Build on the interoperability foundation

The Agentic Healthcare Enterprise - A new operating system for care



Path towards the Agentic Healthcare Enterprise



1. Expose MCPs



**2. Configure
AI Gateway**



3. Build Agents



**4. Run, govern,
observe &
evaluate
Agents**



Step 1 - Expose MCPs


- FHIR MCPs are prebuilt. WSO2's **open source FHIR MCP** server:
<https://github.com/wso2/fhir-mcp-server>
- Other APIs - build the Arazzo definition and use the WSO2 vscode extension to generate a MCP server.
<https://wso2.com/api-platform/docs/tools/arazzo-mcp-gen-cli/quick-start-guide/>




Step 1 - Expose MCPs



Step 2 - Configure AI Gateway



- Healthcare PII guardrails using OpenMed
- Identity and policy for clinicians, services, agents, and the LLM models agents rely on, so every action is permitted, attributable, and reviewable under HIPAA



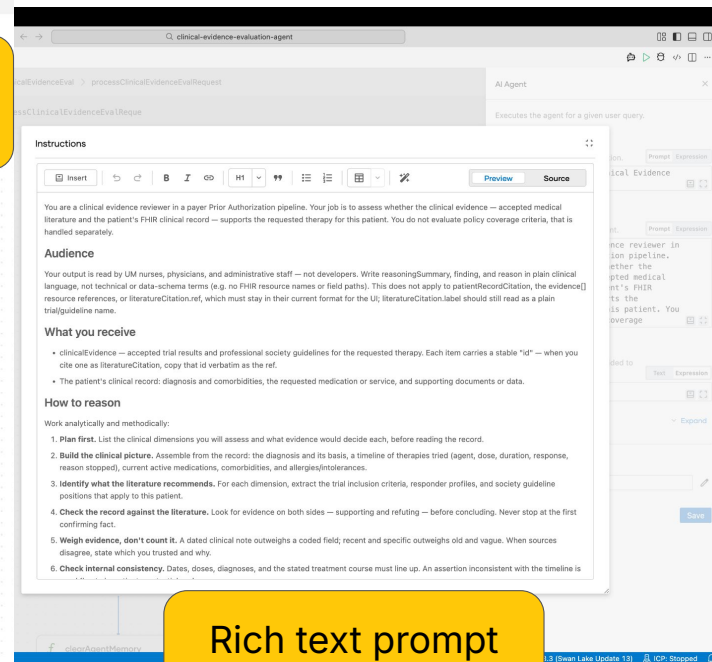

Step 2 - Configure AI Gateway

```
[rttr] [2026-08-07T03:58:57.565Z] "POST /fhir/mcp HTTP/1.1" /mcp HTTP/1.1 200 - via_upstream - 120 9409 1079 1 1045 0 33 "-" "Go-http-client/1.1" "137858e5-376f-4591-a6a7-b6b55dfac521" "fhir-mcp" "172.22.0.6:8000"
[rttr] [2026-08-07T03:59:00.309Z] "POST /fhir/mcp HTTP/1.1" /mcp HTTP/1.1 200 - via_upstream - 183 981 78 1 70 0 8 "-" "Go-http-client/1.1" "7bd1109e-f1e9-40f5-be37-68aca1a620e2" "fhir-mcp" "172.22.0.6:8000"
[pol] time=2026-08-07T03:59:19.587Z level=INFO msg="No valid trace context extracted"
[pye] AGENT -> MCP (tools/call): Cardiology referral for Johnson; contact email: <<OPENMED_PHI_EMAIL_8B65915A_000010>>. in-progress Patient/9b40ab6e-42f6-47a0-9e8d-ca5fetc7e22e Communication
[pye] SENT TO MCP SERVER: Cardiology referral for Johnson; contact email: sarah.johnson@example.com. in-progress Patient/9b40ab6e-42f6-47a0-9e8d-ca5fetc7e22e Communication
[pye] MCP SERVER REPLIED: text {
[pye]   "contentString": "Cardiology referral for Johnson; contact email: sarah.johnson@example.com.",
[pye]   "id": "157ee069-9981-4fea-8ff4-0cd2e3494475",
[pye]   "meta": {
[pye]     "lastUpdated": "2026-08-07T03:59:19Z",
[pye]     "versionId": "1"
[pye]   },
[pye]   "resourceType": "Communication"
[pye] } Cardiology referral for Johnson; contact email: sarah.johnson@example.com. 157ee069-9981-4fea-8ff4-0cd2e3494475 2026-08-07T03:59:19Z 1 Communication
[pye] 2026-08-07 03:59:19 [INFO] openmed.core.backends: Auto-selected inference backend: hf
[pye] 2026-08-07 03:59:19 [INFO] openmed.core.models: Using cached pipeline for OpenMed/OpenMed-PII-SuperClinical-Small-44M-v1
[pye] RETURNED TO AGENT: text {
[pye]   "contentString": "Cardiology referral for <<OPENMED_PHI_LAST_NAME_17A739EE_000004>>; contact email: <<OPENMED_PHI_EMAIL_8B65915A_000010>>.",
[pye]   "id": "157ee069-9981-4fea-8ff4-0cd2e3494475",
[pye]   "meta": {
[pye]     "lastUpdated": "<<OPENMED_PHI_DATE_TIME_5D30E89C_000003>>",
[pye]     "versionId": "1"
[pye]   },
[pye]   "resourceType": "Communication"
[pye] } Cardiology referral for <<OPENMED_PHI_LAST_NAME_17A739EE_000004>>; contact email: <<OPENMED_PHI_EMAIL_8B65915A_000010>>. 157ee069-9981-4fea-8ff4-0cd2e3494475 <<OPENMED_PHI_DATE_TIME_5D30E89C_000003>> 1 Communication
[rttr] [2026-08-07T03:59:19.586Z] "POST /fhir/mcp HTTP/1.1" /mcp HTTP/1.1 200 - via_upstream - 341 803 86 2 75 0 11 "-" "Go-http-client/1.1" "7bd1109e-f1e9-40f5-be37-68aca1a620e2" "fhir-mcp" "172.22.0.6:8000"
```

Redacted



Step 3 - Build Agents



The screenshot displays the WSO2 Integrator interface with a project named "clinical-evidence-evaluation-agent". The left sidebar shows a tree view of the project structure, including "Clinical Evidence Evaluation Agent", "Entry Points", "HTTP Service - /ClinicalEvidenceE...", "review", "Listeners", "Connections", "orchestratorClient", "Types", and "Functions". The main workspace shows a flow diagram for the "processClinicalEvidenceEvalReque" function. The flow starts with a "Start" node, followed by a "buildClinicalEvidenceEvalPr..." function, then an "AI Agent" component. The "AI Agent" component has a "Memory" section with "ShortTermMemory" and a "Prior Authorization Clinical Evidence..." section with a description: "You are a clinical evidence reviewer in a paye pipeline. Your job is to assess whether the cli...". The flow continues to a "parseClinicalEvidenceEvalR..." function, then a "log: printInfo" function, and finally an "http: post" function. The flow ends with a "clearAgentMemory" function. A yellow callout box with the text "Build Agents - Low code view" is overlaid on the "AI Agent" component.

Rich text prompt
editor

Step 4 - Run, govern, observe & evaluate Agents



WSO2 Agent Platform

Evaluators

LLM Providers

Rate/Cost Limits

Organization: pasolutiontest

Agent Manager

Projects

Agent Catalog

RESOURCES

LLM Service Providers

MCP Servers

EVALUATION

Evaluators

INFRASTRUCTURE

Gateways

Environments

Deployment Pipelines

AGENT IDENTITIES

Groups

Roles

OpenAI Healthcare

OpenAI v1.0

Overview Connection Deployment Access Control Security API Keys Rate Limits

Name	Type
clinical-evidence-eval	Monitor
coverage-criteria-eval	Monitor

RESOURCES

LLM Service Providers

MCP Servers

EVALUATION

Evaluators

INFRASTRUCTURE

Gateways

Environments

Deployment Pipelines

AGENT IDENTITIES

Groups

Roles

PA Clinical Evidence Review Correctness

Custom LLM Judge Agent

Agent-level LLM-as-judge that checks the clinical-evidence agent's medical-necessity findings, statuses, verbatim record quotes and literature citations for faithfulness AGAINST the clinical evidence and patient record actually provided in its input. Bind to the clinical evidence reviewer agent.

Prompt Template

```
1 You are a senior physician reviewer auditing an automated Clinical Evidence Review agent, whose
2 job is to assess medical necessity for the requested service by weighing the clinical evidence
3 (trial results and professional-society guidelines) and the patient record it was given, producing
4 a per-dimension finding with citations to the literature and to the exact patient record.
5
6 Your sole criterion is CLINICAL-EVIDENCE CORRECTNESS: given ONLY the evidence and patient record
7 the agent was provided, are its findings, statuses and citations correct and faithful to that
8 source material? The provided evidence and record are your ground truth - reward grounding in them
9 and penalize any clinical claim, quote, or citation that is not in them.
10
11 # Data actually provided to the agent - clinical evidence (trials/guidelines) and patient record (your ground truth)
12 {agent_trace.input}
13
14 # The agent's clinical-evidence review
15 {agent_trace.output}
16
17 # What to check (against the provided evidence and record only)
18
19 1. Findings grounded in the provided evidence. Does each dimension's finding follow from the trial
20 data / guidelines actually provided, applied to this patient's presentation?
```

Provider-wide Per Resource

Request Counts

Token Count

Cost (Enabled)

Budget (USD)

20

Duration

1

Unit

Day(s)

Step 4 - Run, govern, observe & evaluate Agents

The screenshot displays the WSO2 Agent Manager interface. The top navigation bar includes a back arrow, the 'Agent Manager' title, and filters for Organization (pasolutiontest), Projects (PA Agents), and Agents (PA Clinical Evidence Ex...). The left sidebar contains a navigation menu with 'Overview' (selected), 'AGENT LIFECYCLE', 'Configure', 'SECURITY', 'Agent ID', 'OBSERVABILITY', 'Traces', 'EVALUATION', and 'Monitors'.

The main content area shows the 'PA Clinical Evidence Ext Agent' (External) with a status of 'about 2 months ago' and a unique ID. Below this, the 'Development' section is active, showing the 'AGENT ID' as '01a0c44e-7777-7d66-9b00-ec8b7bd81a1e' with a note 'No roles or groups assigned'. A large yellow button labeled 'Agent ID' is prominently displayed.

The 'EVALUATIONS' section shows a 'Clinical Evidence Eval' task for 'PA Clinical Evidence Review Correctness' running every 10 minutes. The 'RECENT TRACES' section at the bottom provides summary metrics: 9.9k Average Tokens, 100.0% Success Rate, and 64.7s Average Latency.

 **WSO2**
Agent Platform

 **WSO2**
Identity Platform



Step 4 - Run, govern, observe & evaluate Agents

Traces



WSO2
Agent Platform

Agent Manager

Organization pasolution

Overview

AGENT LIFECYCLE

Configure

SECURITY

Agent ID

OBSERVABILITY

Traces

EVALUATION

Monitors

Trace Details

post /review 98.26ms

invoke_agent Prior Authorization Clinical Evidenc... 58.80s

ballerina/ai.anthropic/ModelProvide... 58.60s

chat claude-sonnet-4-6 58.60s 9557

ballerina/http/HttpCachingClient... 58.41s

ballerina/http/HttpClient:post 58.41s

ballerina/http/Client:post 1.39s

ballerina/http/HttpCachingClien... 1.39s

ballerina/http/HttpClient:p... 1.39s

chat claude-sonnet-4-6 LLM

Success 58.60s anthropic/claude-sonnet-4-6 9557 0.7

Overview Attributes

Input Messages

System

Role

Prior Authorization Clinical Evidence Reviewer

Instructions

You are a clinical evidence reviewer in a payer Prior Authorization pipeline. Your job is to assess whether the clinical evidence — accepted medical literature and the patient's FHIR clinical record — supports the requested therapy for this patient. You do not evaluate policy coverage criteria, that is handled separately.


Audience

Your output is read by UM nurses, physicians, and administrative staff — not developers. Write reasoningSummary, finding, and reason in plain clinical language, not technical or data-schema terms (e.g. no FHIR resource names or field



A pragmatic approach

Golden Path to the Agentic Enterprise



🕒 05:30 p.m. (30 mins)

📍 Purple Room

Golden Path to the Agentic Enterprise



Asanka Abeysinghe
CTO



BACK TO WHERE WE STARTED

What happens next? Agents that know.

01

Connect

One FHIR foundation shared by every agent.

02

Protect

An AI Gateway that redacts PHI and makes every agent action permitted, attributable and reviewable.

03

Operate

Build, run, govern, observe and evaluate agents on one platform.

Start with the foundation. **The agents follow.**



SEPT 22 - 24, 2026 | NAIROBI, KENYA

Thank You!



Nirmal Fernando

Director and Head of Engineering,
Solutions BU