AI MedAgent

Medication screening over CDS Hooks · coverage reported on every call

> **A demonstration of a patent-pending method. Advisory only. Not a medical device. The endpoint is unauthenticated and must not be connected to a production EHR.**

What it is

A medication screening service that a hospital's electronic health record calls at the moment a clinician selects an order, and that reports, on every single call, how many of the eight standard screening classes it was actually able to perform. It is reachable today at `https://aimedagent.net/cds-services` and anyone can verify it from a browser without asking permission.

It speaks CDS Hooks, a published HL7 standard. Epic supports it and lists medication ordering decision support among its own integration playbooks. Oracle Health and MEDITECH are both named contributors to Da Vinci Coverage Requirements Discovery, which is built on CDS Hooks. A service is an HTTPS endpoint, which means there is no agent to install, no interface engine change, no database, and no vendor in the middle.

The method, which is the protected asset

Every clinical decision support product ever built goes silent when it finds nothing, and every clinician has learned to read that silence as *checked, and fine*. It is not. It is *checked whatever I happen to check, which you have never been told*.

This service inverts that. It returns a **coverage card on every invocation, findings or not**, stating how many screening classes ran, how many ran against a maintained drug knowledge base, and naming the ones that did not with the reason each was refused. The count is measured against a taxonomy the system did not invent: the eight published module names of Wolters Kluwer's Medi-Span clinical screening APIs. Measured against a vendor's own published list, the figure can be checked by anyone.

> **The design rule that makes it useful rather than harmful.** The coverage card is returned with indicator `info`, never `warning`. A warning that fires on every order is noise, and within a week it is invisible noise. An informational card that appears every time is a status line. Promoting it to a warning would turn this into the alert fatigue it exists to prevent.

How it connects

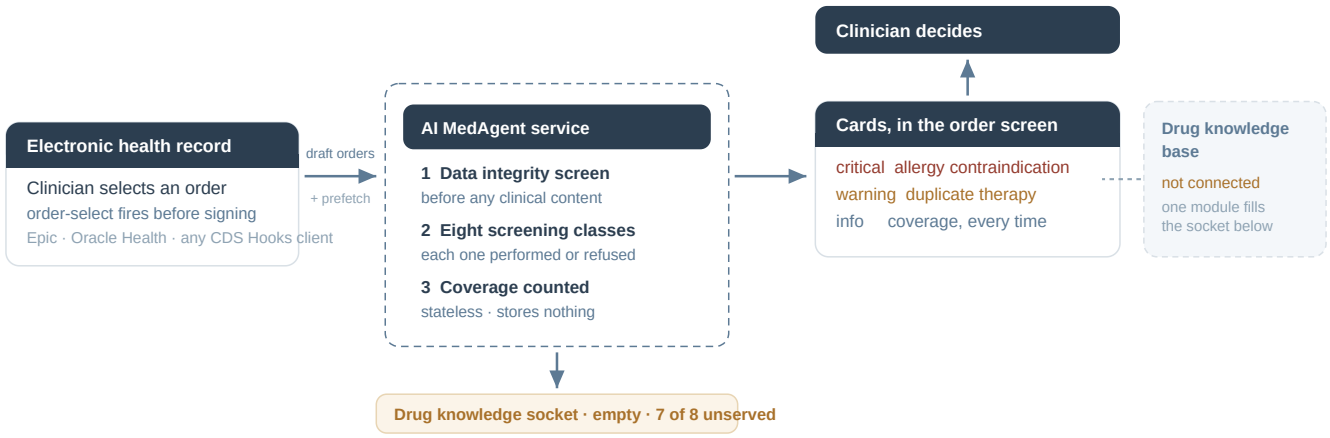

The socket is drawn empty on purpose. That is the honest state of the system, and it is the shape of what a drug data partner would fill. Nothing else in the diagram changes when it is filled.

The eight screening classes

Class names are transcribed from Wolters Kluwer's published Medi-Span clinical screening API module list. Status is what the service performs today, without any commercial drug data source connected.

SCREENING CLASS	WHAT IT ANSWERS	TODAY
Drug Allergy Screening	An ordered medication against a recorded allergen	**Partial. Substance-name text match, drafts included. No ingredient or cross-reactivity check**
Drug Interaction Screening	Drug against drug, food and alcohol	**No source. NLM discontinued the free RxNav interaction API on 2 January 2024 and named no replacement**
Duplicate Therapy Screening	Two orders that do the same job under different names	**Partial. Ingredient-name collision after stripping strength and form. Blind to two drugs of one class**
Dose Screening + Drug Orders	Dose, frequency and duration against the accepted range	**No source**
Drug Disease Contraindications	A medication against a disease state. Renal and hepatic adjustment live here	**Context only. Surfaces renal function and counts the orders. Maps no drug to any disease**
Pregnancy, Lactation, Age, and Gender Screening	Patient attributes that change whether a drug is appropriate at all	**No source**
Route Contraindications	A route contraindicated for this product	**No source**
US Controlled Substances	DEA and state schedule status	**No source**

Every output the service produces carries this table's live state as a count. On a representative record that count reads: one of eight performed, zero against a maintained drug knowledge base.

What it refuses to answer, and why that is the feature

A screen that had no data to screen against returns NOT PERFORMED. It never returns zero findings. Zero findings and no data look identical on a screen and mean opposite things at a bedside. The refusal is enforced above the screening logic, so a more capable drug data source cannot argue its way past a missing input.

Four states of "no allergies", which are not one state

STATE	WHAT IT MEANS	ALLERGY SCREEN
`blocked`	The application is not permitted to read allergy data for this patient	**Refused**
`notstarted`	Epic's `List.emptyReason`. A list exists and has never been opened. Nobody has asked	**Refused**
`empty`	A search ran and returned nothing. This cannot distinguish a patient with no allergies from a patient nobody asked, because FHIR returns an identical empty Bundle for both	**Refused**
`recorded`	At least one allergen is on file	**Performed**

All four refuse the screen. They refuse it for different reasons, and the card says which. Collapsing them into "no known allergies" is the single cheapest way for a reasoning layer to harm a patient.

Data integrity, checked before any clinical content

You cannot reason over a record you have not validated, so the integrity screen runs first and its findings are printed above the clinical section rather than in a footnote. Seven checks run today. All are arithmetic, every finding cites the FHIR resource it came from, and each is reproducible from the record.

- **Physiologic impossibility, vital signs.** A blood pressure of 1/1 mmHg and a body temperature of 45 °C were retrieved from a live EHR sandbox, stored and served without complaint.
- **Impossible for this patient.** Height and weight checked as stability and rate. Temperature, pulse and blood pressure are deliberately *not* checked this way: variation is what a vital sign is for, and an engine that reports a fever as a data error trains clinicians to ignore it.
- **Medication list staleness.** A stop date in the past carried alongside a status of active.
- **Mutually exclusive staged diagnoses.** One patient carrying chronic kidney disease stages 1, 2, 3 and 4 simultaneously, all active. Renal dosing is decided by stage.
- **Same test, same day, results that disagree.** Reported once with a count, not once per occurrence.
- **One code, more than one unit of measure.** Any consumer that reads the value and ignores the unit is averaging two different quantities.
- **Laboratory plausibility.** 53 LOINC codes with impossibility bounds, published as a readable table rather than buried in code. A bound is applied only when the unit matches the unit it was written for; otherwise

the mismatch is reported and no bound is applied.

> **Impossibility bounds are not reference ranges.** Nothing in that table says a value is abnormal. It says a value is not a measurement of the thing its code claims. A potassium of 7.8 mmol/L is a medical emergency and belongs in the clinical section. A potassium of 78 is not potassium. Two of the published bounds were wrong on first release and were corrected; the table ships as data so that it can be argued with.

Security and governance

What the service does

- **Stores nothing.** No database, no cache, no state. It reasons over what the hook sends and forgets it when the response is written.
- **Logs no patient data.** Card counts and coverage numbers only. Never a request body, a patient identifier or a drug name.
- **TLS in transit**, served from managed edge hosting.

What it does not do yet

- **No authentication.** The CDS Hooks security model requires a signed JWT from the client, verified against its published JWK Set. Not implemented.
- **No write-back.** Read and reason only. Nothing is filed into any chart.
- Both service descriptions in the discovery document carry the words DEMONSTRATION ENDPOINT so it cannot be wired to a production EHR by accident.

What it is not

It is not a medical device, and it is not for clinical use. It is not a replacement for a pharmacist, a drug knowledge base, or an EHR's native decision support. It does not diagnose, and it does not prescribe. It makes no claim to improve any outcome, and no outcome has been measured. It is a demonstration that the method is doable and credible, built deliberately so that what it cannot do is visible in its own output rather than in a disclaimer.

Verification

CLAIM	HOW TO CHECK IT
The endpoint is real and conformant	`https://aimedagent.net/cds-services` returns a CDS Hooks discovery document
It runs inside somebody else's tool	Point the public sandbox at `sandbox.cds-hooks.org` to `https://aimedagent.net`
The eight classes are not invented	Wolters Kluwer publishes the Medi-Span clinical screening module list on their own site
Problem list duplication is common in real records	Wang EC, Wright A. Characterizing outpatient problem list completeness and duplications in the electronic health record. *JAMIA* 2020;27(8):1190–1197.

CLAIM	HOW TO CHECK IT
	doi:10.1093/jamia/ocaa125
Extremely low HbA1c has real, non-artifactual causes	Wang P. What clinical laboratorians should do in response to extremely low hemoglobin A1c results. *Lab Med* 2017;48(1):89–92. doi:10.1093/labmed/lmw050
The free federal interaction source is gone	NLM RxNav notice: the Drug Interaction API was discontinued on or about 2 January 2024, with no replacement named
The synthetic patients are reproducible	MITRE Synthea public sample set, by named patient identifier

Patent and provenance

AI MedAgent is a patent-pending method. This datasheet describes the medication screening and CDS Hooks interface only; the continuous encounter method is covered in the companion advisory datasheet. Live records shown in the demonstration are drawn from Epic's public FHIR sandbox, which contains no real patient. Synthetic records are generated by Synthea, an open MITRE project, from published incidence and prevalence statistics; they describe no real person.

Daniel Pettus · aimedagent.net · dan@aimedagent.net