Modernizing the CMS for an AI-First Future

Preparing Healthcare Marketing for AI and Agent-Based Search

Nardeep Singh, MSITM
Director of Marketing, Communications & Technology
Renown Health

Jovana Trudell
Manager of Marketing Technology
Renown Health

SHSMD Society for Health Care Strategy & Market Development℠

Outline & Learning Objectives

1. **Is your CMS actually AI-ready?**
2. **Where AI belongs in the workflow**
3. **From keywords to prompts**
4. **What this sets up**

LEARNING OBJECTIVES

Assess whether their current CMS and content model are AI-ready by identifying gaps in structure, governance, and reuse that limit search intelligence, automation, and future personalization.

Apply AI safely within content workflows by defining where assistive AI adds value (ideation, editing, SEO guidance, data quality) while maintaining human review, accuracy, and healthcare governance controls.

Use AI and search intelligence reporting to interpret user prompts, query behavior, and AI Overview results, and translate those insights into data-informed content prioritization and updates aligned with real user intent.

Most health systems have adopted AI.

Very few have content AI can actually use.

Buy the AI tool › Generate drafts › Nobody can publish them quickly › No search lift › **“AI doesn’t work for healthcare”**

The tool wasn’t the problem.

AI tools fail because the content underneath them isn’t structured, governed, or reusable — not because the model is weak.

AI reads structure. A page is a blob.

PAGE-BASED CMS

One rich-text field. Headings, body copy, phone numbers, provider names — all fused into markup.

You cannot ask a blob what a service line's conditions are.

STRUCTURED CONTENT MODEL

```
conditionName · symptoms[] · treatments[]
providers[] →ref · locations[] →ref
faq[] · reviewedBy · reviewDate
```

Each fact in its own field. Queryable, reusable, retrievable.

Schema

The blueprint — what fields every content type must have.

Structured content

Each fact in its own field, not fused into a page.

Headless CMS

Content lives in one place and powers the site, apps, AI tools, and internal systems at once.

SHSMD Society for Health Care Strategy & Market Development℠

Key terms before we start

Seven words that come up all day in this conversation.

Schema

The blueprint for your content. It defines what fields every page needs and what type of data goes in each one.

Structured content

Content stored as organized data, not as pages. Each piece of information lives in its own field.

Headless CMS

A system that separates where you manage content from where it appears — one place powering site, apps, AI tools, and internal systems at once.

API

How software talks to other software. When one system needs data from another, it uses an API.

MCP (Model Context Protocol)

A standard that lets AI tools connect directly to your CMS — look at your content, understand its structure, and take action through conversation.

Next.js

A framework for building websites and web applications. It turns your content and design into the site users interact with.

Vercel

The platform that hosts and deploys the site. Push a change, the site updates automatically.

“When content is modeled intentionally, AI stops guessing and starts reasoning.”

SANITY · CONTENT OPERATING SYSTEM ANNOUNCEMENT, MARCH 2026

SHSMD Society for Health Care Strategy & Market Development℠

What “modeled intentionally” actually means

Every provider document in the Renown Provider Hub is keyed:

```
provider.{npi}
```

NPI is a federal standard — stable and identical across Epic, MdStaff, Loyal, renown.org, and the Hub.

Direct lookup

The document ID is the identifier. No query required to find a provider.

Deduplication is automatic

A deterministic ID cannot create a duplicate record.

Cross-system matching

Any future pipeline can join records across five systems on one key.

An identity-key decision is what “AI-ready” looks like in practice. It is not a content decision or a design decision — it is the one that determines whether anything downstream can reason over your data.

Is your CMS AI-ready? Five layers.

1 **Content model**
Is content typed and fielded, or freeform rich text?

2 **Structured data**
Is schema generated from the model, or hand-written per page?

3 **Governance**
Who owns accuracy — and is the review trail recorded?

4 **Reuse**
Does one fact live in one place and appear on many surfaces?

5 **Retrievability**
Can crawlers and agents reach the answer without executing JavaScript?

START WITH LAYER 5

Three questions about your own site:

01 Does the answer exist in the HTML, or only after JavaScript runs?

02 What does robots.txt say to AI crawlers — and did anyone decide that on purpose?

03 Can a machine reach your content through an API, or only by scraping?

What we did: legacy CMS → Sanity

2,000+
pages migrated

47
document types

1
content lake

This was not a CMS swap.

It was a shift from page-based content — HTML blobs locked to templates — to structured, queryable data that lives independently of any front-end.

New tech stack

NEXT.JS › Sanity › Vercel

WHAT WE'D DO DIFFERENTLY

We modeled the document types while we migrated instead of before. Given the time again we would have inventoried the content first — which types actually earn their place, what is genuinely reusable, and what was just a page somebody made once. We are still cleaning that up.

We picked this stack because it was modern. We did not know that treating content as data would be the thing that fast-tracked everything after it.

Why that structure is AI enablement

- Sanity stores content as structured JSON documents.
- AI can identify fields such as provider name, specialty, location, biography and metadata.
- References connect providers, services, locations and events.
- AI can search, analyze and update content through APIs.
- The presentation layer does not trap the content inside a specific webpage.
- Sanity's Content Agent already understands Renown's schemas and relationships.
- Better organized data means better output. An AI writing into named fields produces something you can publish; an AI writing into a blank page produces something you have to rebuild.

FOR EXAMPLE, YOU CAN ASK CONTENT AGENT

```
Find all provider profiles missing
biographies, identify their associated
specialties and locations, and return the
records by department.
```

It can interpret doctor, specialty, location and department as separate but connected document types. Sanity says Content Agent can find, analyze, create and update schema-aligned content through natural-language requests.

Source: Sanity Content Agent

PART 2

Where AI belongs in a healthcare content workflow

Learning Objective 2 — defining where assistive AI adds value while maintaining human review, accuracy, and governance controls.

SHSMD Society for Health Care Strategy & Market Development℠

THE OPERATING PRINCIPLE

AI assists.
Humans decide.

Nothing clinical publishes without clinical review.

SHSMD Society for Health Care Strategy & Market Development℠

Sanity Content Agent: what it is, how it works

An AI assistant inside the CMS. Plain-English requests — no code, no GROQ queries.

Schema-aware

Understands our content structure, field types, references, and relationships.

Permission-based

Operates with your existing user permissions. It can only access what you can.

In the Dashboard

Alongside the Studio, on the same content — no copy-paste between tools.

1 You ask

A plain-English request in the Dashboard.

2 Agent reads schema

It knows the field types and relationships.

3 Finds & acts

Queries the dataset and proposes changes.

4 Stages for review

Every change arrives as a proposed draft.

Content Agent can publish. We don’t let it.

Every proposed change is staged as a draft and lands in a review queue with a per-field diff and a named approver. That is our policy, not a product limit — and it is what lets bulk operations and autonomous agents run without a new governance conversation each time.

SHSMD Society for Health Care Strategy & Market Development℠

What Content Agent can do

Six capabilities, all through conversation in the Dashboard — no code, no GROQ queries.

Search & Query

Find content across all projects using natural language. No need to write GROQ queries.

Create Content

Generate new documents that automatically match your schema — proper fields, references, and metadata.

Update & Edit

Make changes to existing documents. Update single fields or transform entire content sections.

Bulk Operations

Apply changes across hundreds of documents in one conversation. Track progress in real time.

Content Audits

Find missing metadata, stale pages, broken references, and quality issues across your entire library.

Web research

Search the web, extract info from URLs, and combine external data with your content.

> **Schema-aware means enforced: required fields are respected, references stay intact, data types match the model.**

What it can and can’t do

CONTENT AGENT CAN

- + Read your entire Sanity schema and understand content structure
- + Create, edit, and publish documents (blog posts, news, pages)
- + Audit content for missing fields, thin pages, outdated info
- + Generate meta titles, meta descriptions, and alt text
- + Search the web for competitor examples and best practices
- + Write blog post drafts, news articles, and content briefs
- + Add internal links between documents
- + Batch-update fields across multiple documents

CONTENT AGENT CAN’T

- — Publish to social media (Instagram, LinkedIn, Facebook)
- — Access Google Analytics, Search Console, or Piwik PRO
- — Edit frontend code, templates, or CSS
- — Add or modify schema definitions (that’s dev work)
- — Implement structured data / JSON-LD output
- — Fix build pipeline or rendering issues
- — Access password-protected systems or third-party tools
- — Upload images — it can only reference what’s in the library

When Content Agent tells you something is outside its scope, ask what it can do instead. It will usually offer a workaround — like drafting social captions inside a blog post that your social team can lift directly.

Start with operations, not writing

Missing Alt Text Detection

Scanned every image across 2,000+ pages and returned the list.

An ADA audit that took days — done in one conversation.

Find Active Providers Missing Specialty Degree

Cross-referenced provider documents to surface data gaps.

No database query. No report pull. No ticket.

Active Providers Image Count

Instant inventory of provider photos across the directory.

A question you'd never bother filing a request for.

Content operations is the lowest-risk, highest-return place to start.

Most teams reach for AI as a writing tool. An agent that can read every document in your dataset does in an afternoon what a content audit does in a quarter — and it touches no clinical claims.

Then creation, with the schema as the guardrail

Create Alert Banner

Content Agent generated production-ready alert banners that matched Renown's schema exactly.

No copy-paste from ChatGPT, no manual formatting.

Blog Draft Style Review

Content Agent reviewed drafts against brand guidelines and gave feedback on tone, style, and consistency.

Editorial review at the speed of a first pass.

FIVE THREADS, FIVE CATEGORIES

- Accessibility & compliance auditing
- Provider data validation
- SEO strategy at scale
- Content creation with schema awareness
- Editorial review & brand voice

> **The schema is the guardrail: output that doesn't fit the model doesn't save.**

SHSMD Society for Health Care Strategy & Market Development℠

Research, strategy, and a finished page — in one conversation

1

Ask for market research

“What are other health systems publishing about this service line?” It searches the web and brings the findings into the chat.

2

Talk through the plan

You argue with it. What we cover, what we skip, what our differentiator actually is.

3

Have it write

It drafts in our brand voice, against the outline you just agreed on.

4

It builds the page

Because it is connected to our content, it creates the document with the right fields filled in — not a Word file you have to re-key.

You never leave the conversation. That only works because the content underneath is structured.

1. Ask for market research

You don’t need to leave the CMS to search competitors, collect links, and bring everything back into another tool.

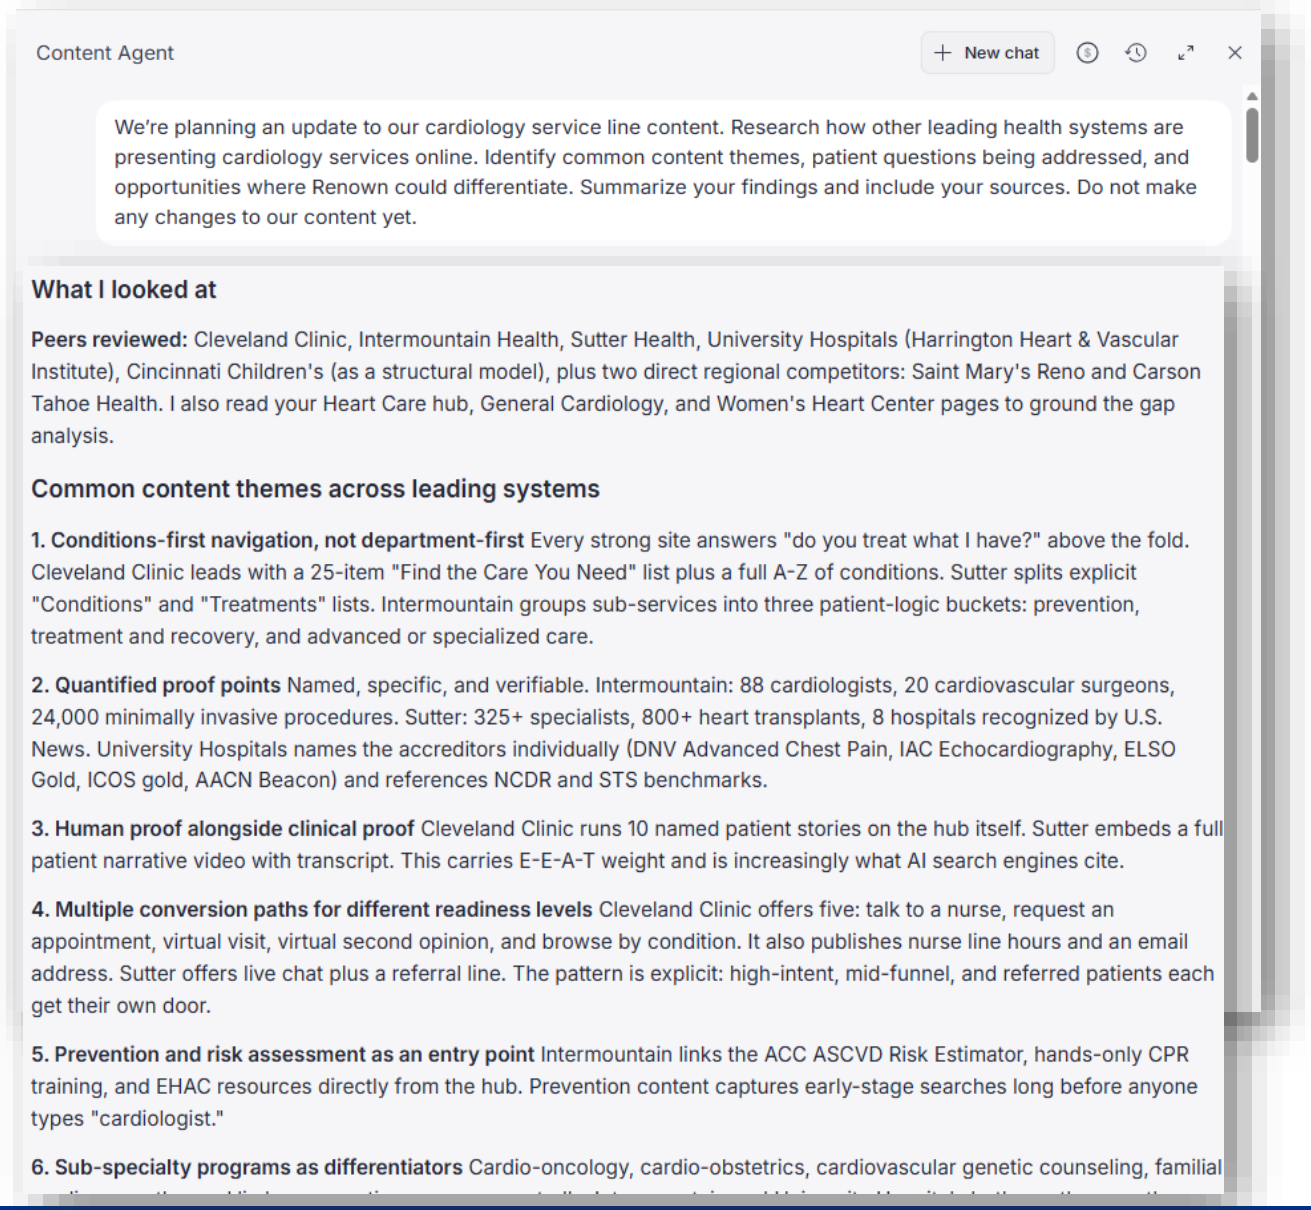

2. Talk through the plan

This is where you can challenge the output.

Ex: Let's move “Why Renown” lower. I want the patient's immediate needs and next steps to come first. Put access to care near the top and simplify the number of sections. Show me the revised outline.

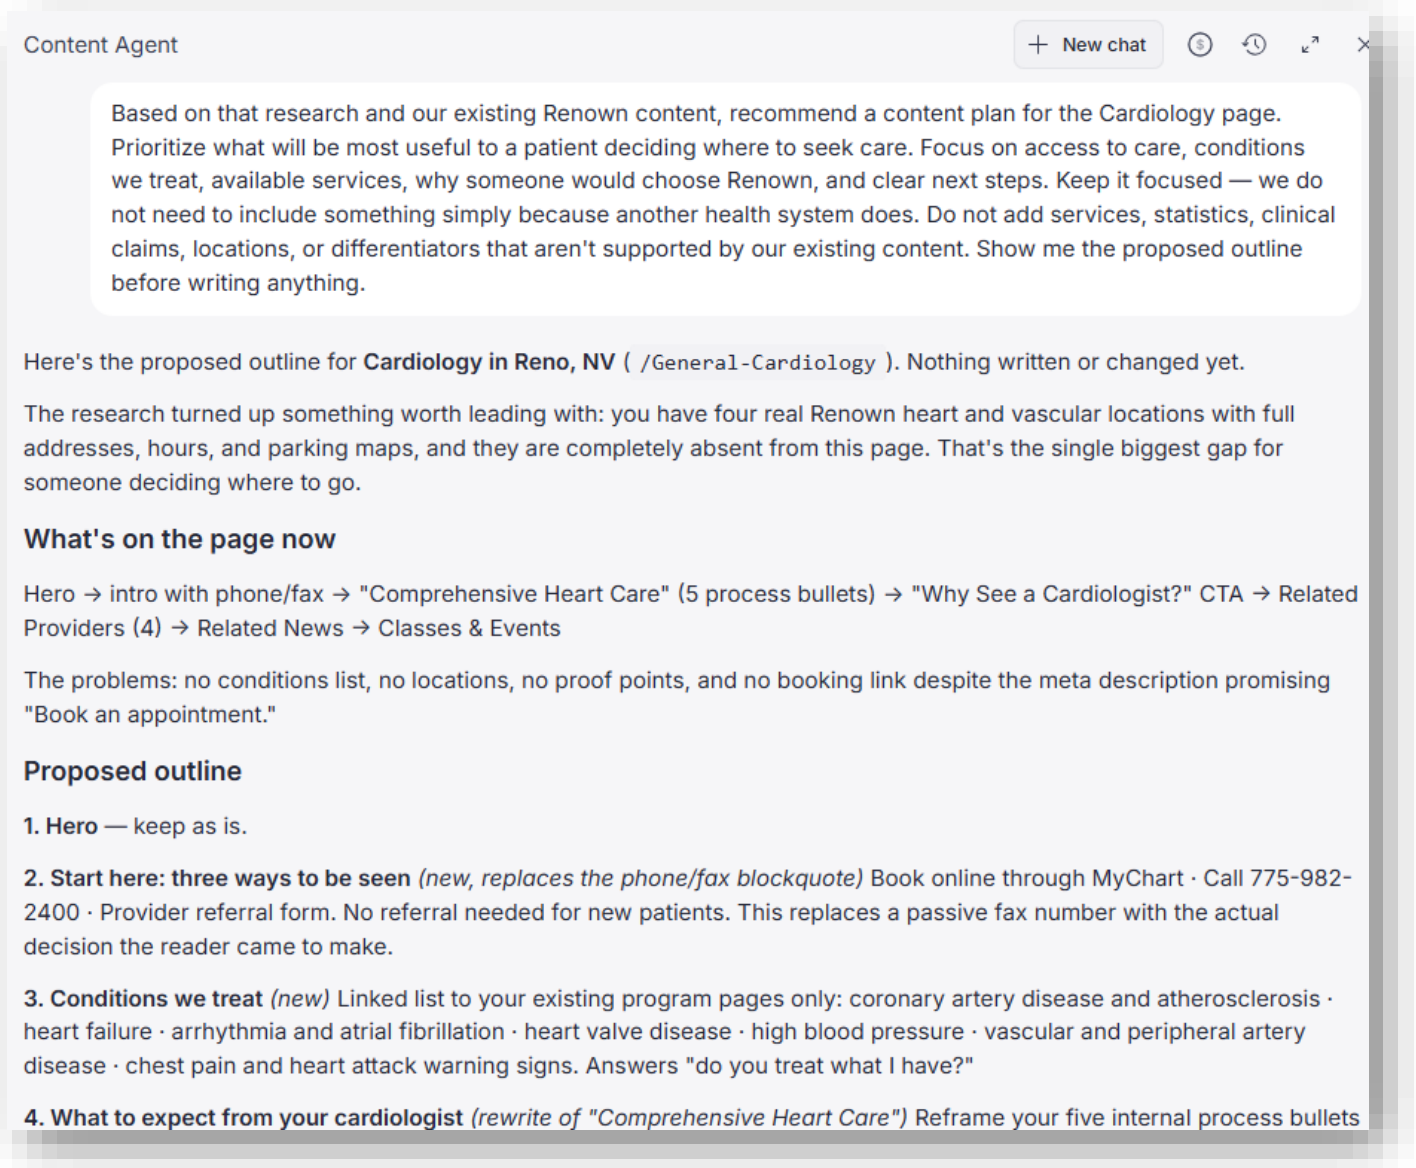

3. Have it write

Once you're satisfied with the outline, ask it to write it for you.

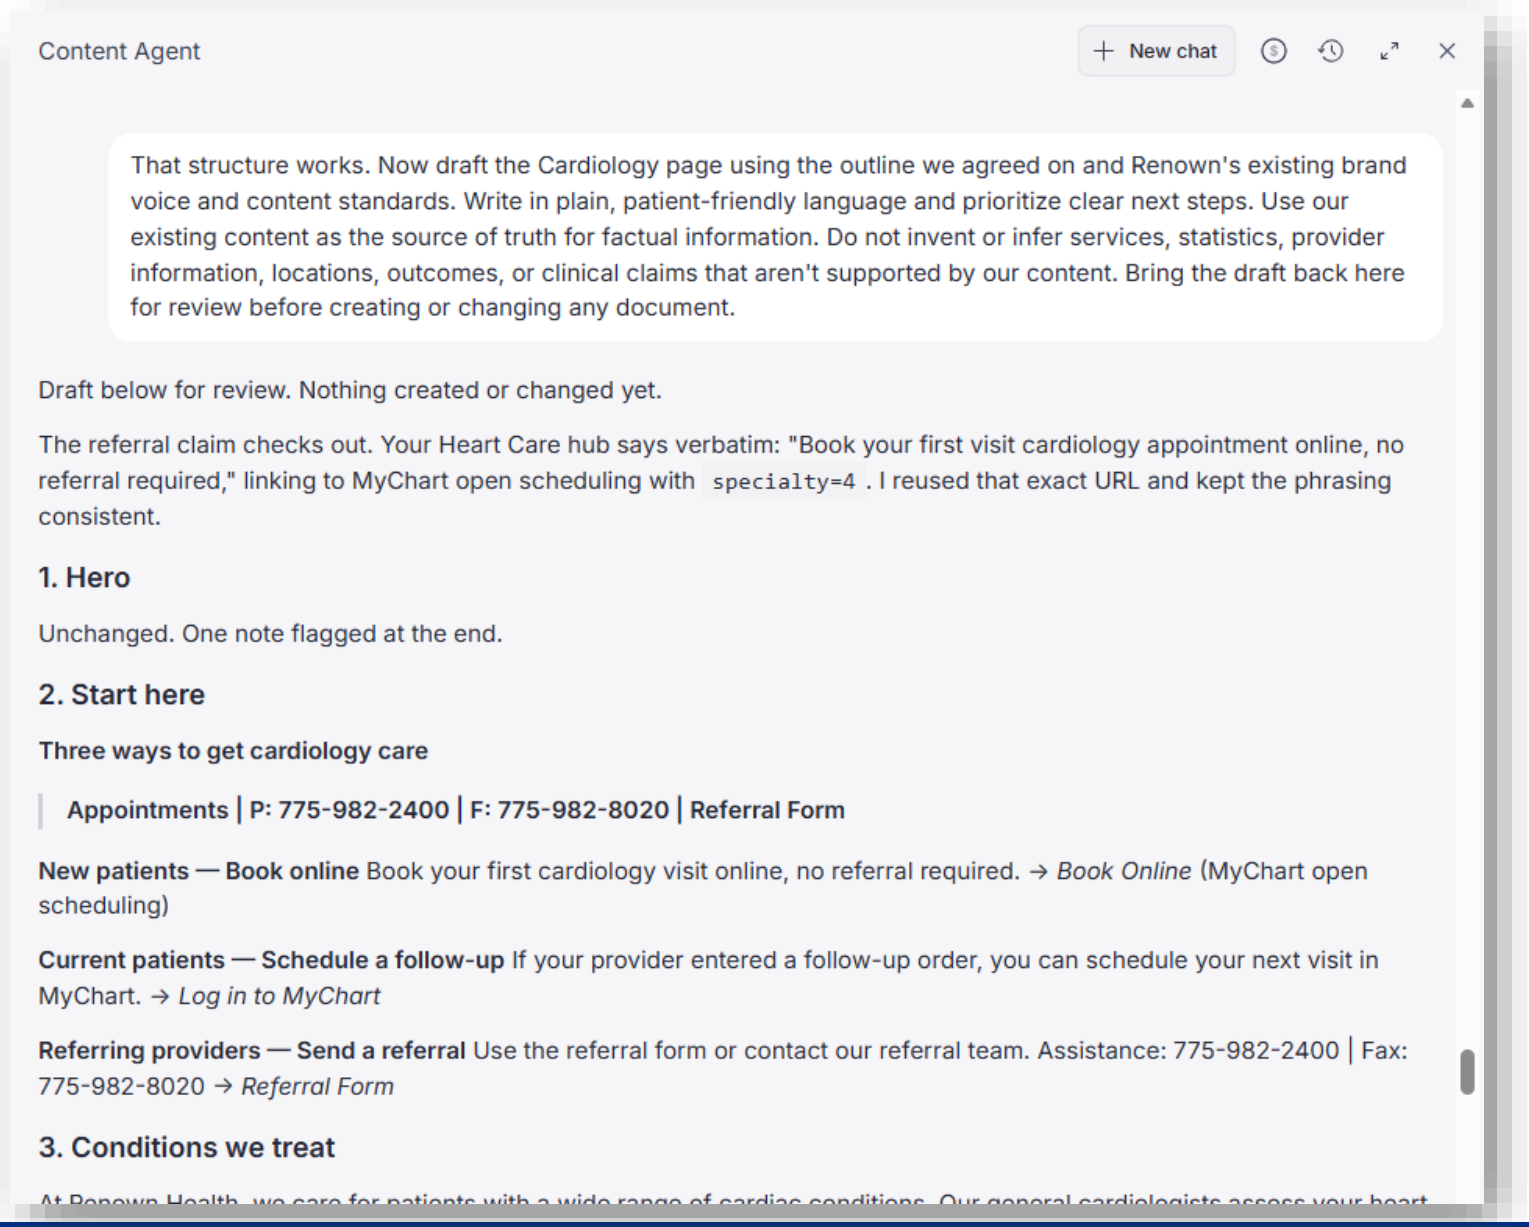

4. It builds the page

Open changes withing the Dashboard.

1 **A real service line document**

The typed fields — the structure from Part 1, in production.

2 **A data quality check**

Across a document set. What it flags and what it doesn't.

3 **An SEO structure suggestion**

On one page. The diff, not the output.

4 **The staged draft**

And the approval step. Nothing published itself.

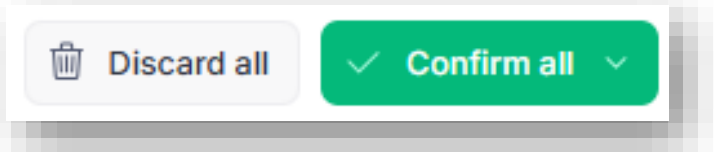

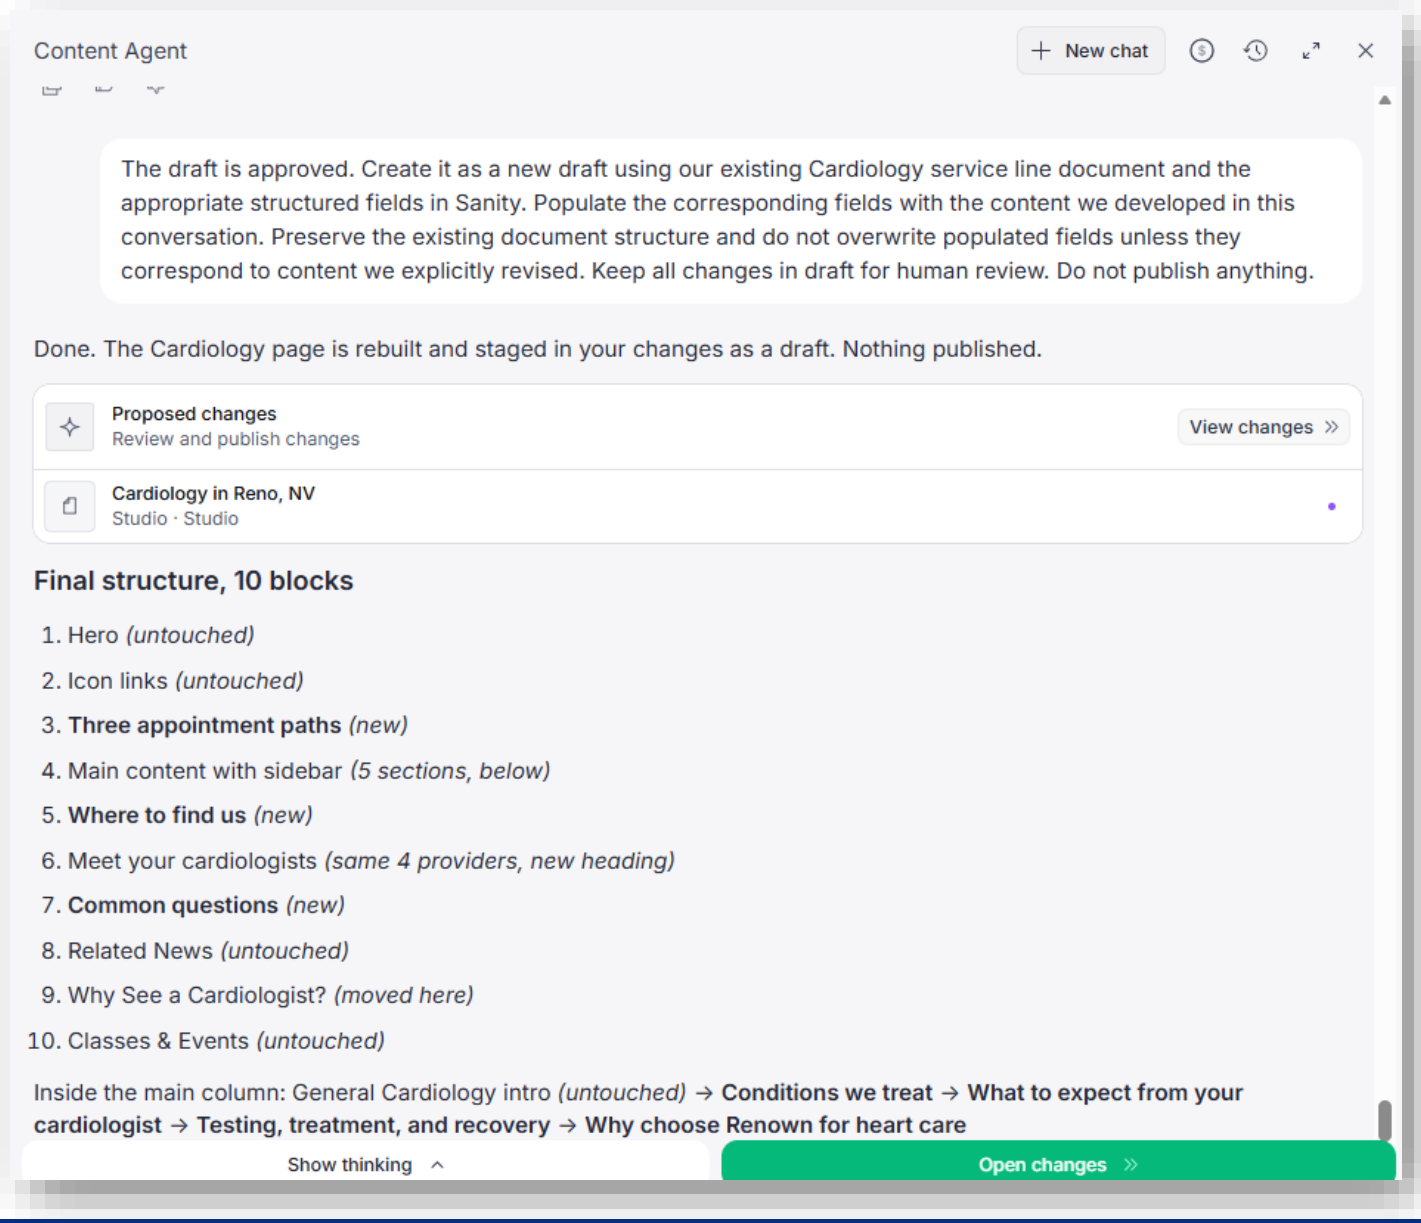

What a real prompt actually looks like

PROMPT — VERBATIM

```
Find all documents of type blog post, news article, and content
page where the main image field exists but has no alt text —
meaning both the document-level mainImage.alt field is empty or
undefined AND the referenced image asset's altText in the Media
Library is also empty or undefined. If either field has a
value, exclude that document. Return the results as a browsable
list showing the document title, type, and a link to edit it.
```

WHAT CAME BACK

The prompt is long, specific, and unglamorous.

That is the point. Showing the real one is the difference between a method and a magic trick.

SHSMD Society for Health Care Strategy & Market Development℠

Governance: three tiers of content risk

Tier	Content type	What AI may do	Who approves
1. Clinical	Conditions, treatments, symptoms	Structure and readability only. No clinical claims.	Clinical reviewer + marketing
2. Marketing	Service lines, provider bios, campaigns	Draft and edit	Marketing + service line owner
3. Operational	Alt text, tags, metadata, internal links	Act, with sampled QA	Content ops spot-check

THE GUARDRAILS THAT GOT US TO YES

- No PHI in prompts
- No AI-generated clinical claims
- Permission inheritance
- Staged drafts, named approver
- Full audit trail
- Legal reviewed vendor terms first

PART 3

From keywords to prompts

Learning Objective 3 — interpreting user prompts, query behavior, and AI Overview results, and translating them into content prioritization.

SHSMD Society for Health Care Strategy & Market Development℠

Search changed shape — so did what we track

Keywords were three words.	→	**Prompts are sentences.**
Queries came one at a time.	→	**Conversations come in threads.**
Results were links.	→	**Answers are summaries — with or without you in them.**

WHAT WE TRACK NOW

AI Overview presence
Which of our priority queries trigger an AI answer at all.

Citation
Whether we are cited in it.

Who gets cited instead of us
The competitive question that rankings never answered.

Branded prompt visibility
How we appear when someone asks about care in our market.

Search intelligence in practice

Intent clusters

Grouped by service line, not by keyword volume.

Content gaps

Questions with real demand and no page of ours to answer them.

AI Overview tracking

Presence and citation on our priority terms.

Three intent types, one condition: symptom-checking (“is this serious”) · comparing (“which treatment, which hospital”) · ready to act (“who takes my insurance”). One page can’t serve all three well. Structured content serves all three from the same facts.

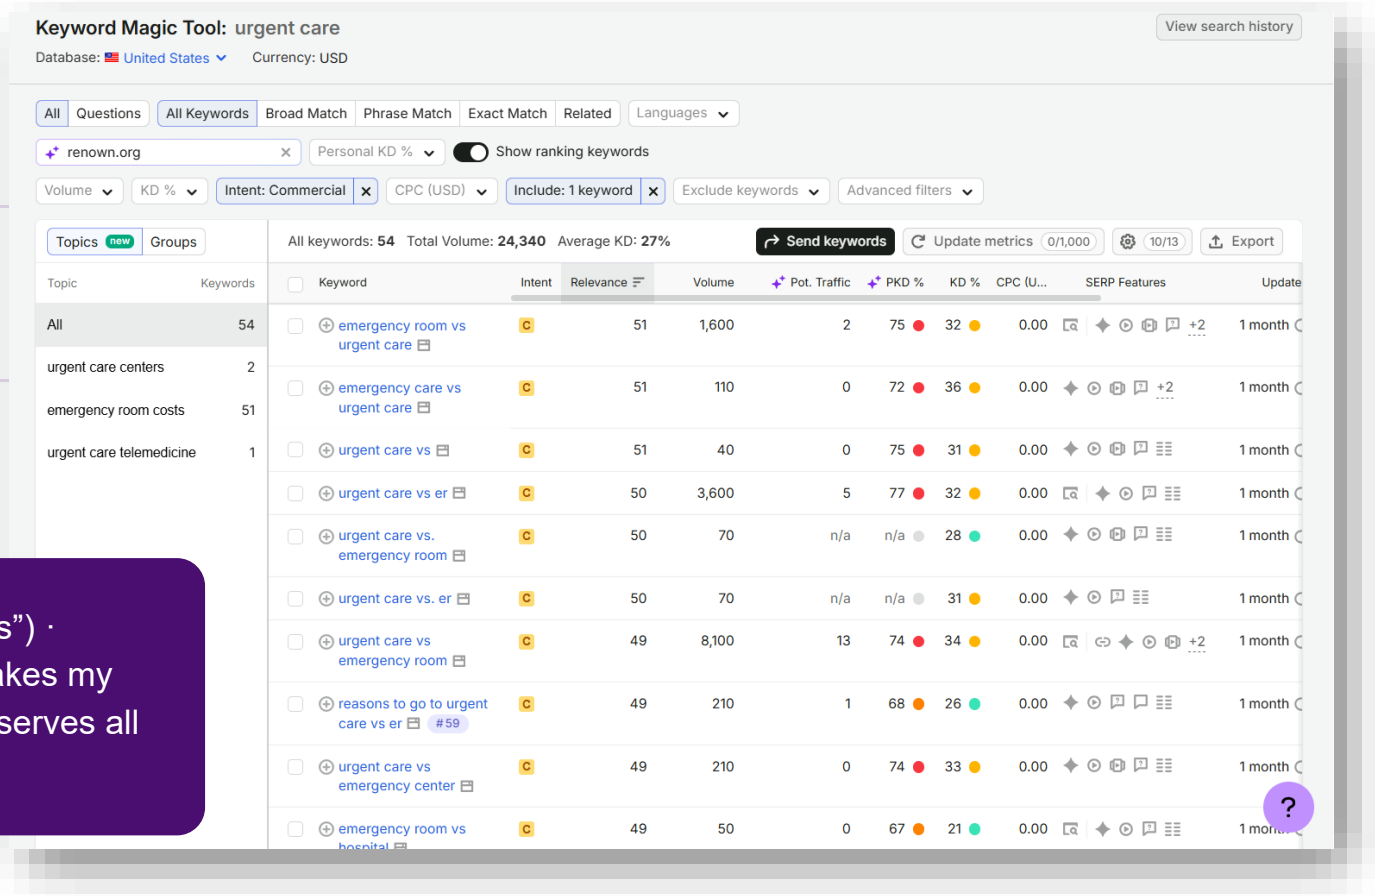

SHSMD Society for Health Care Strategy & Market Development™

From insight to a content queue

```
Score = (demand × AI answer presence × our absence × service line value) ÷ effort
```

Run it monthly. It produces a queue, not a wish list.

WHAT WE CHANGE ON THE PAGE

- H2s rewritten as the questions people actually ask
- The direct answer in the first 40–60 words
- Entities named plainly: condition, treatment, location, provider
- FAQ and MedicalWebPage schema generated from the model
- Duplicate and near-duplicate pages consolidated
- Review dates surfaced, not buried

> **The method matters more than the exact weights — build the version your team will actually run.**

Three Content Agent threads in one month on urgent care alone — the strategy developed in conversation with the content itself.

SHSMD Society for Health Care Strategy & Market Development℠

PART 4

What this sets up

Structured content is the prerequisite for personalization and for agent-based search. Both need facts they can retrieve. What follows is our latest experiment in building on that — Provider Hub.

Renown’s web technology stack

One content layer underneath, seven things built on top of it.

renown.org	The public surface. Next.js, deployed on Vercel.
Structured content	Sanity Content Lake. 47 document types, every field queryable.
Content Agent	Content operations by conversation, inside the Dashboard.
MCP + Sanity Context	Read-only access for external AI tools and our own agents. They can query; they cannot write.
Custom builds	Provider Hub — our own application on the same content.
Event-driven automation	Functions and the Agent API, triggered by changes in the dataset.
Human oversight	Every change staged for a named approver. Nothing reaches the live site unreviewed.

> Every layer above the content is replaceable. That is the whole point of putting the content underneath.

What the content model made possible

Provider Hub

A self-service PWA built in-house on the same stack — 15 provider screens plus an admin console inside the Studio. It replaced a spreadsheet audit and a third-party middleware vendor.

```
714 providers · 637 submissions
362 (57%) from providers
8 reviewers · 4.4/5 satisfaction
```

Specified, built, and shipped by the marketing team.

An autonomous drafting agent

Built on the Claude API, working directly against Sanity content. It researches topics, writes in our brand voice, and pushes structured drafts in.

Every draft enters the same review queue as human work.

A door we can now open

MCP lets an outside AI tool read our content directly — our schema, our published pages, and nothing else. Read-only: it can ask, it cannot write.

That is the door a patient FAQ bot or an internal search tool would go through, answering from content a human approved instead of guessing.

We haven’t shipped one yet. The connection exists — what we build on it is a decision, not a rebuild.

The build method is the transferable part: a written specification document served as the authoritative brief, and the AI built from it. Spec first, not vibes.

SHSMD Society for Health Care Strategy & Market Development℠

A CLOSER LOOK

Provider Hub

Not every answer is an AI answer. Sometimes it is being able to build the right tool, quickly.

One problem, one application, and the data it generated — which is where AI came back into it.

SHSMD Society for Health Care Strategy & Market Development℠

A directory problem, solved with an application

Provider profiles were out of date and nobody owned fixing them. Not an AI problem.

The problem

714 providers, profiles maintained by a spreadsheet audit and a third-party middleware vendor. Slow, expensive, and always behind.

What we built

A web app providers open on their phone to verify and update their own Find a Doctor profile — 15 screens, plus a reviewer console inside the Studio.

Why it could exist

The content was already structured and keyed by NPI. No new CMS, no vendor, no procurement cycle — just an application on content we already had.

4 Weeks
from written spec to providers using it

Six figures
estimated financial impact
(tool savings + avoided outsourced development costs)

0
new vendors, systems, or logins

Specified, built, and shipped by the marketing team, spec first. The agility comes from the content layer being right — not from anyone on the team becoming an engineer.

> And then it started producing something we didn't plan for: a record of every correction our reviewers make.

SHSMD Society for Health Care Strategy & Market Development℠

Your review log is a training set

Every time a reviewer edits a submission before approving it, that is a recorded judgment: this is what good looks like.

691
submissions

209
corrections

0.30
corrections per submission

WHERE THEY LAND

Education & training — 138 (66%)

Board certifications — 26 (12%)

Everything else, 12 fields — 45 (22%)

FOUR PATTERNS, OVER AND OVER

Abbreviation → full name · ABIM → American Board of Internal Medicine

Missing year supplied · Larner College of Medicine → … (2022)

Punctuation and casing · psychiatry → Psychiatry

Institution named properly · Legacy Emanuel → Legacy Emanuel Medical Center

We didn’t write the agent’s instructions. We derived them from what the team already does.

SHSMD Society for Health Care Strategy & Market Development℠

What a reviewer can’t fix

A reviewer cannot correct an answer nobody gave.

59%

of providers with a medical school have no residency at all

275 of 467

32%

of training entries have no year recorded

259 of 820

94

providers are bookable online — the intake form never mentioned scheduling

0 corrections

These generate zero corrections, so they are invisible in the review data — and they are the biggest quality hole in the directory. Absence is invisible in the data you collect by default. You have to go looking for it deliberately.

Sort the problems by what can actually solve them

Problem	Right tool	Why not an LLM
Punctuation, casing	**Plain code**	Deterministic, free, instant
Language misspellings	**A closed list**	It can’t invent a language
Institution names	**List built from your own data**	Your reviewers already decided
Missing residency, missing year	**The agent**	Needs a human to answer
“What is this clinic called now?”	**Nothing yet**	Local knowledge nobody has

37%

of past reviewer corrections would be caught inside Provider Hub — before a submission ever reaches renown.org

ONE HONEST NOTE

This reviewer went live weeks ago. The 37% is measured against past corrections, not results — whether it reduces the team’s workload is unproven. The test is re-running the 0.30 baseline in a few weeks.

SHSMD Society for Health Care Strategy & Market Development℠

The same governance, one channel further

An AI reviewer inside the submission form, with voice powered by ElevenLabs. It reads the provider's own draft and asks for what's missing.

1 **The provider fills in their profile**
Same signed link, same form, same fields.

2 **The agent reads the draft**
No last name, no NPI, a board certification that names the specialty but not the board — plus eight more places detail would help, highlighted in place.

3 **The provider can ask it anything**
Tap the orb and ask about any note, out loud.

4 **Send it as it is**
The agent advises. It never blocks a submission, and the record is untouched until a human approves.

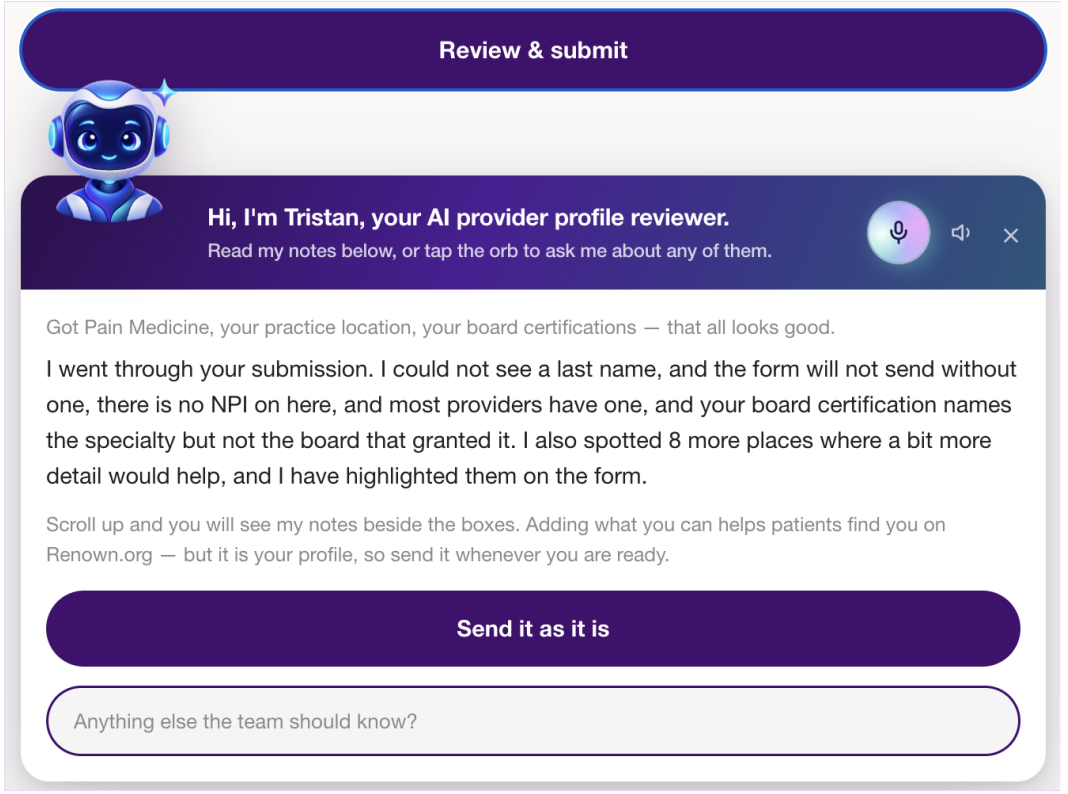

The agent gets no special privileges.

Same endpoint, same review queue, same approver. What it changes is where absence gets caught — at the point of entry, instead of never.

SHSMD Society for Health Care Strategy & Market Development℠

Three Key Take-Aways

1 Not everything is an AI problem.

Pick the right tools first, then introduce AI where it makes sense. Provider Hub started as a data problem, not an AI project — AI only came back in once the data existed. **Before the next AI purchase, ask what problem it solves and whether the data it needs exists yet.**

2 Treat your content as data.

Know your content types and get them into a structure — inside your current system if you can, and as the first question you ask if you are shopping for a new one. Structure is what makes you agile, and it is the stepping stone to personalization, MCP, and custom applications. **If you cannot list your content types, that is the first project.**

3 Marketing is not a silo anymore.

Martech is not a separate department. Consumer expectations move faster than our roadmaps, which means the people doing marketing have to be technical enough to understand the stack they are buying. **Be in the room for the next platform decision — and ask what it will let you build in three years.**

“When content is modeled intentionally, AI stops guessing and starts reasoning.”

SHSMD Society for Health Care Strategy & Market Development℠

Renown® HEALTH

Questions?

Please be sure to complete the session evaluation on the mobile app.

Nardeep Singh, MSITM
Renown Health

Jovana Trudell
Renown Health

SHSMD Society for Health Care Strategy & Market Development℠

Speaker Biographies

Nardeep Singh, MSITM

Director of Marketing, Communications & Technology · Renown Health

Brings more than 12 years of experience in enterprise marketing technology, with a focus on digital transformation, platform strategy and emerging technology. She led Renown's migration from a legacy CMS to Sanity across more than 2,000 pages and the development of its clinical trials directory and provider directory platform. A nationally recognized healthcare marketing technology leader and speaker, Nardeep has presented at HMPS, SHSMD and HCIC and is a SHSMD Rising Star Award recipient. She holds an M.S. in Information Technology Management and is a certified medical interpreter.

Jovana Trudell

Manager of Marketing Technology · Renown Health

Leads digital strategy and marketing technology initiatives that connect content, technology and the patient experience. Jovana's work spans CMS strategy, digital experience, marketing automation and emerging AI applications. She drives cross-functional initiatives including Renown's clinical trials directory and provider directory transformation, bringing teams together to move complex digital projects from strategy through execution. An AI-forward marketer and ClickUp Power User, Jovana is passionate about building practical, sustainable ways for healthcare organizations to innovate while keeping people at the center.

Renown® HEALTH

Organization logo may be placed on this slide only, per SHSMD template requirements.

SHSMD Society for Health Care Strategy & Market Development℠

Bibliography / References

Sanity. Content Agent and Agent Actions documentation.

Sanity.io. “Sanity Launches The AI Content Operating System for the AI Era.” Press release, March 4, 2026.

Semrush. AI Overviews and AI visibility tracking documentation.

Google Search Central. Structured data and AI features guidance.

Schema.org. MedicalWebPage, Physician, and FAQ Page vocabularies.

Centers for Medicare & Medicaid Services. National Provider Identifier standard.