

The Domain-Native AI Organization

How to Leverage Domain Expertise to Build Better AI Products




Dr Christopher Lovejoy, MD
Founder, Notius Labs

2026-04-10

medical doctor



AI engineer



AI Engineer
World's Fair

Anterior

MAKE YOUR LLM APP A DOMAIN EXPERT:
How to Build an LLM-Native Expert System

19:18

Make your LLM app a Domain Expert: How to Build an Expert System —...

84K views • 7 months ago

AIE AI Engineer

Vertical AI is a multi-trillion-dollar opportunity. But y

Microsoft 15 chapters Introduction | The last mi

Chris Lovejoy, MD @ChrisLovejoy_

My hot take: The vertical AI products that will win are not the ones with the most sophisticated models or techniques - but those with **the best system for incorporating domain expertise.**

The LLMs in the product need to understand the industry-specific and customer-specific way that a workflow is performed. **This is solved by better domain context, not better reasoning models.**

This is the bet we're taking at Anterior - and why we built a system to translate domain insights into improvements.

We call it an **adaptive domain intelligence engine** and here's how it works:

OUR BET FOR VERTICAL AI APPLICATIONS:

system for incorporating domain insights	>	the sophistication of your models and pipelines
--	---	---

1:41 pm · 30 Jul 2025 · **20.4K** Views

9 17 161 228

“but how should I build my org to enable this?”

MY BET FOR VERTICAL AI APPLICATIONS:

winning in vertical AI is a
model problem.

MY BET FOR VERTICAL AI APPLICATIONS:

winning in vertical AI is a
~~model~~ problem.

organizational



Oracle

Directly add
domain expertise



Evaluator

Define and
measure quality



Architect

Build self-improving
systems

Tomorrow's Titans: Vertical AI

by Tiffany Luck and James Kaplan

AI

Harvey Hits \$11B as Sequoia Triples Down on Vertical AI Inflection

Legal AI startup's valuation surge signals enterprise AI crossing from experimentation to mission-critical infrastructure in professional services.

Trend 4: Vertical AI shows potential to dwarf legacy SaaS with new applications and business models

Vertical SaaS proved to be a sleeping giant that transformed industries during the first cloud revolution. Today, the top 20 US publicly traded vertical SaaS companies represent a combined market capitalization of ~\$300 billion, with more than half of these companies having IPO'd in the last ten years.

The US Bureau of Labor Statistics cites the Business and Professional Services industry at 13% of US GDP making this sector alone, dominated by repetitive language tasks, ~10x the size of the software industry.

~50% of generative AI
projects were
abandoned in 2025

Source: Gartner (<https://www.gartner.com/en/articles/genai-project-failure>)

The most common mistakes I see

- 1 . Not hiring domain experts (or hiring them too late)
- 2 . Hiring the wrong kind of domain expert
- 3 . Not fitting them into your org, not leveraging them appropriately

The most common mistakes I see

1. Not hiring domain experts (or hiring them too late)

1. Do I really need a domain expert(s)?

2. Hiring the wrong kind of domain expert

2. Who do I need?

3. Not fitting them into your org, not leveraging them appropriately

3. How should I leverage them?

The most common mistakes I see

1. Not hiring domain experts (or hiring them too late)

1. *Do I really need a domain expert(s)?*

2. Hiring the wrong kind of domain expert

2. *Who do I need?*

3. Not fitting them into your org, not leveraging them appropriately

3. *How should I leverage them?*

1. Do I really need a domain expert?

Appraising AI
quality requires
judgment.

Judgment
requires
domain
expertise.

2. Who do I need?



Oracle

Directly add
domain expertise



Evaluator

Define and
measure quality



Architect

Build self-improving
systems

Three models for incorporating domain expertise



Three models for incorporating domain expertise



Oracle



Domain expert

assess

improve




Domain expert



Three models for incorporating domain expertise



Evaluator



Domain expert

assess

improve



Engineer



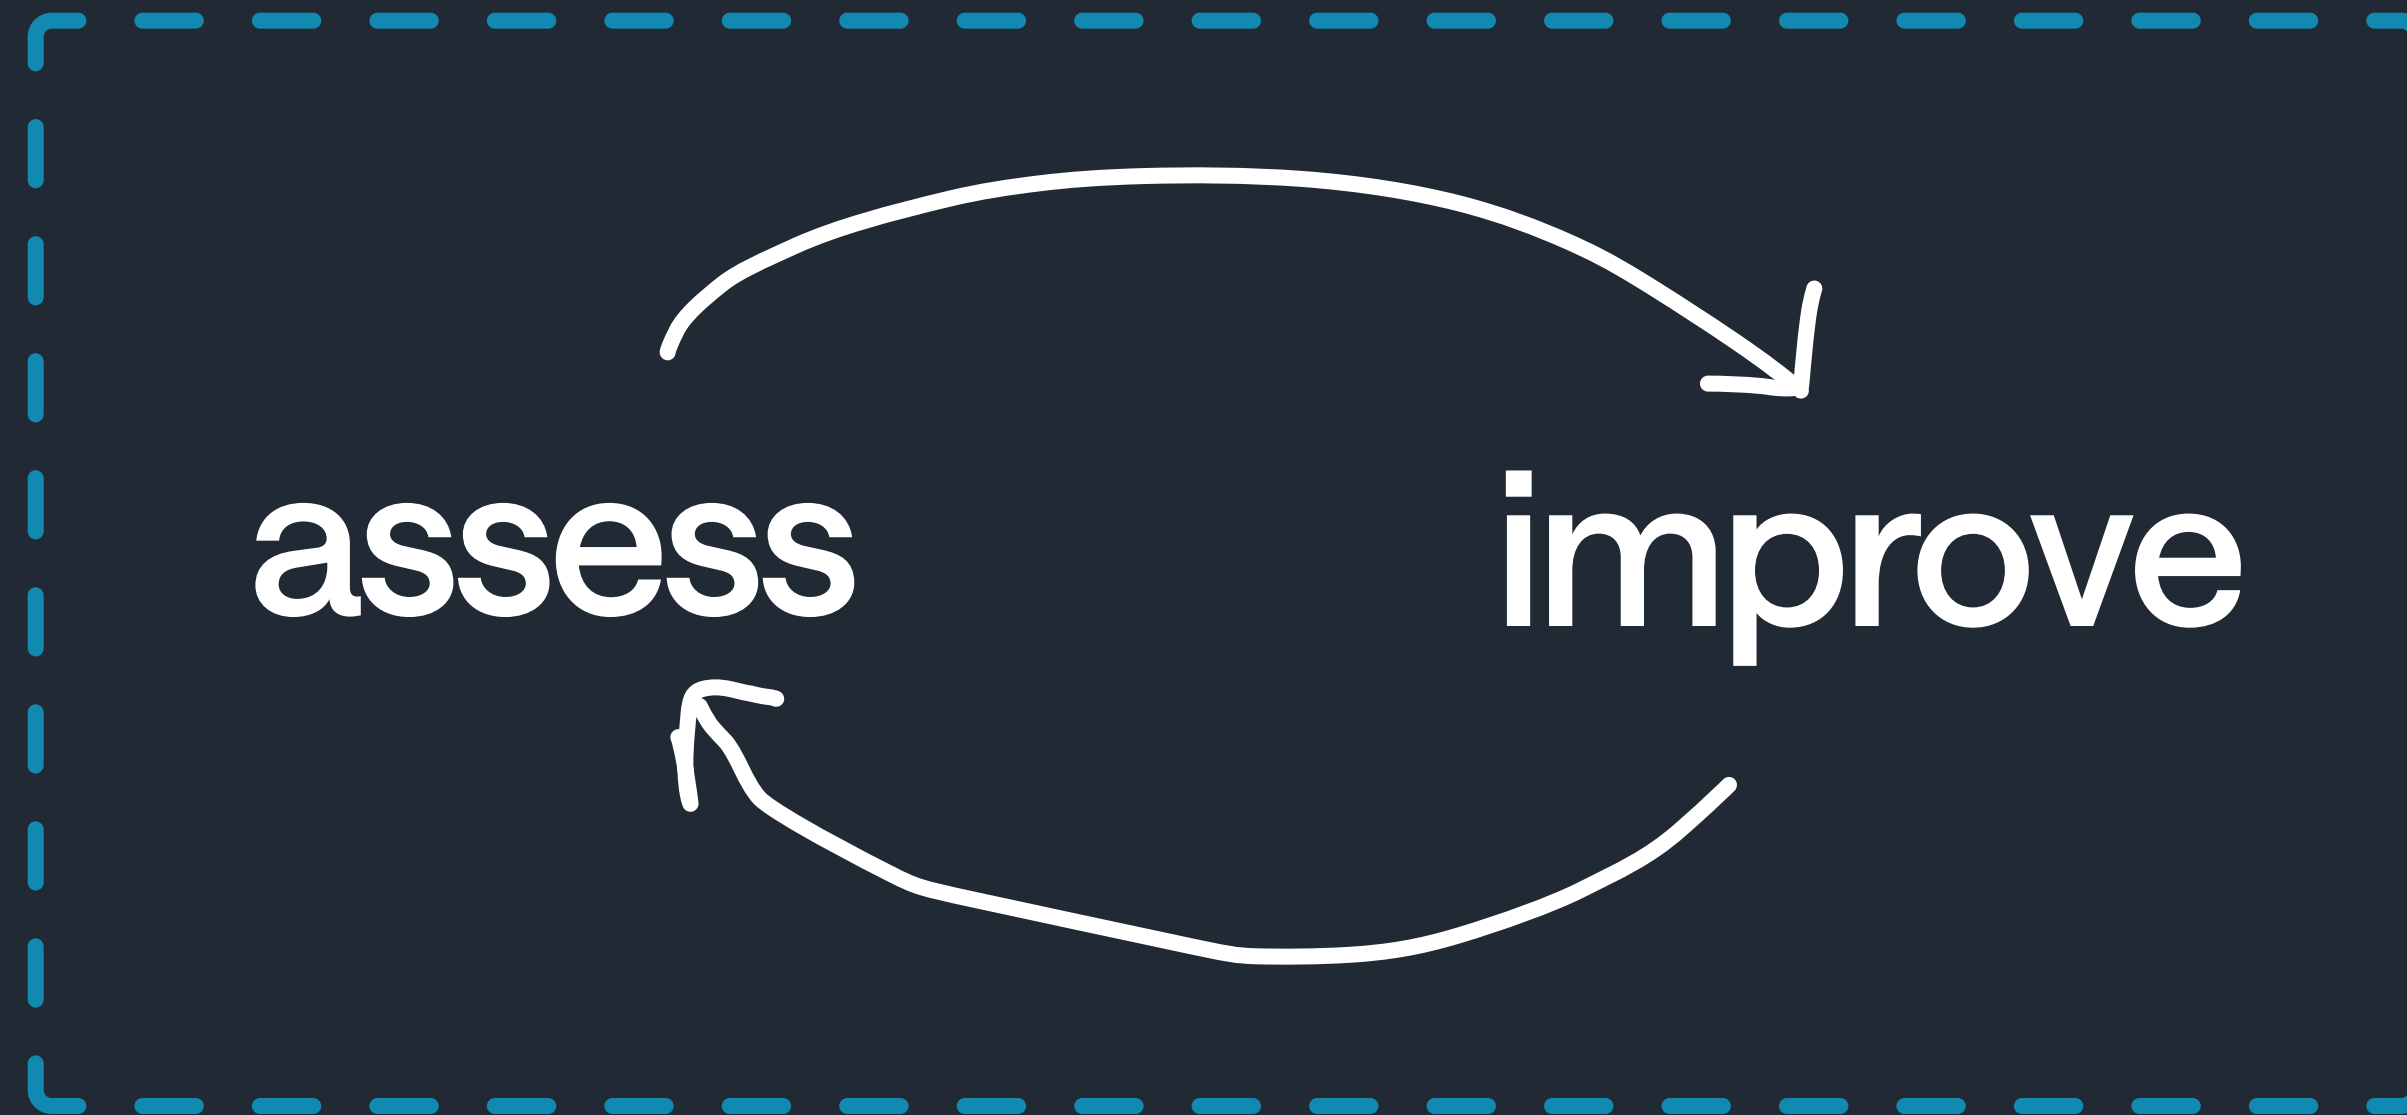
Three models for incorporating domain expertise



Architect



Domain
expert



How to build an LLM-Native Expert System
AI Engineer World's Fair (SF, June 2025)
<https://www.youtube.com/watch?v=MRM7oA3JsFs>

2. Who do I need for my org.?

Can measure performance
in metrics?

No

Yes

1 person enough?

Manual iteration fast enough?

Yes

No

Yes

No

Oracle



Decentralized
Oracles



Evaluator



Architect





Decentralized Oracles





Oracle



Evaluator



Architect



Progression
necessary if:




Oracle

Progression
possible if:



Decentralized Oracles

Breaking at
current scale



You can reliably
measure quality



Evaluator

Manual iteration
is too slow



You can
automate
improvement



Architect

Case Studies

Case Study 1: (meeting notes)



Background: Writer / journalist

Jo Barrow

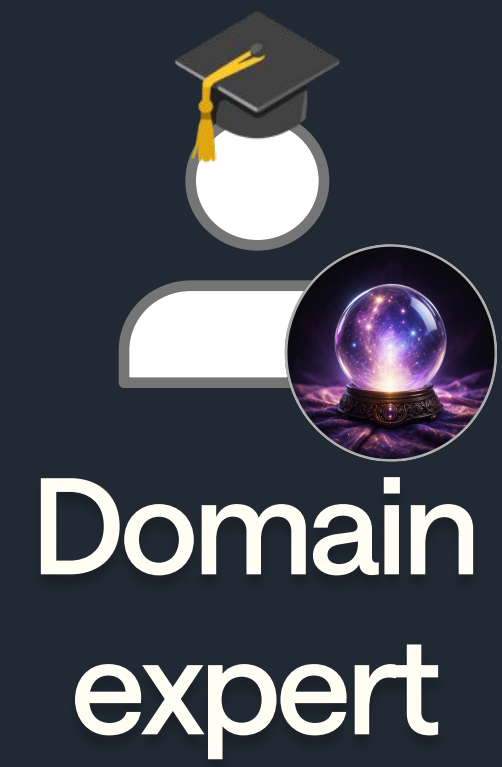
Case Study 1: (meeting notes)



Background: Writer / journalist


Jo Barrow

- 1st employee, wrote all the prompts
- Extensive research into what makes good meeting notes
- Primary gatekeeper of AI quality



assess

improve



Case Study 1: (meeting notes)



Background: Writer / journalist

Jo Barrow

- 1st employee, wrote all the prompts
- Extensive research into what makes good meeting notes
- Primary gatekeeper of AI quality

Oracle has worked well because:

- 1 . There is no objectively “perfect” meeting note - human taste is really important.
- 2 . There is a single core output (the meeting note) which is amenable to direct human review-and-improvement

Case Study 2:



(medical AI scribe)



Background: Medical doctor
then McKinsey

Roy Pekny

Case Study 2: Tandem (medical AI scribe)



Background: Medical doctor
then McKinsey

Roy Pekny

- Early employee, reviewed medical notes, updated prompt templates (**oracle**)
- As scaled, became impossible for one person

Case Study 2:



(medical AI scribe)



Roy

Lauren

Yan

Artin

- Updated platform to support long tail of prompt customizations
- Hired doctors from different countries and specialties

Case Study 2: Tandem (medical AI scribe)



Roy

Lauren

Yan

Artin

- Updated platform to support long tail of prompt customizations
- Hired doctors from different countries and specialties

Oracle to Decentralized Oracle has worked well because:

- 1 . Medical expertise is required to assess the quality of medical notes
- 2 . Medical note quality has some subjectivity
- 3 . Variations require geographical and specialty-based domain expertise
- 4 . There are so many customizations that many domain experts are needed

Case Study 3: Anterior (prior auth)



Background: Medical doctor
then AI startups

Chris Lovejoy

- First technical employee, built initial product (prompts and code)
- Reviewed AI outputs, updated prompts and code (**oracle**)

Case Study 3: Anterior (prior auth)



Background: Medical doctor
then AI startups

Chris Lovejoy

- First technical employee, built initial product (prompts and code)
- Reviewed AI outputs, updated prompts and code (**oracle**)
- Defined metrics and failure modes, built review dashboard, hired clinicians (**evaluator**)

Case Study 3:

 **Anterior**

(prior auth)



Background: Medical doctor
then AI startups

Chris Lovejoy

- First technical employee, built initial product (prompts and code)
- Reviewed AI outputs, updated prompts and code (**oracle**)
- Defined metrics and failure modes, built review dashboard, hired clinicians (**evaluator**)
- Designed methods for automatic improvement (architect)

Case Study 3: Anterior (prior auth)



Background: Medical doctor then AI startups

Chris Lovejoy

- First technical employee, built initial product (prompts and code)
- Reviewed AI outputs, updated prompts and code (**oracle**)
- Defined metrics and failure modes, built review dashboard, hired clinicians (**evaluator**)
- Designed methods for automatic improvement (architect)

Oracle to **Evaluator** to **Architect** worked well because:

- 1 . The AI quality is measurable: either correct or incorrect Prior Auth decision
- 2 . Clinical reasoning required to determine
- 3 . Large variation in how prior auth rules are interpreted between customers
- 4 . Understanding failure modes (as Oracle) is key for defining how to assess and improve (as Evaluator, Architect)

2. Who do I need for my org.
and what skills should they have?



Oracle

Directly add
domain expertise



Evaluator

Define and
measure quality



Architect

Build self-improving
systems

What skills should they have?



Oracle

core skills

- relevant domain expertise - direct experience of the use case

other relevant

- experience with prompting, context engineering
- depth of insight, attention to detail
- customer communication




What skills should they have?



Evaluator

core skills

- relevant domain expertise
- data science intuition



Domain expert

assess

improve



Engineer

other relevant

- everything from Oracle
- statistics
- leadership experience
- industry connections
- product management

What skills should they have?



Architect

core skills

- relevant domain expertise
- experience working on LLM-powered products

other relevant

- everything from Evaluator
- engineering skills

assess

improve



Domain expert

3. How should I leverage them?

Define a Principal Domain Expert

Move faster by avoiding consensus by committee

Give them ownership

Truly leverage them to build a differentiated product

Hire for breadth - and evolve their role over time

Your best Architect or Evaluator might be your former Oracle

Takeaways

- Assessing the AI quality in your product requires **domain expertise** – whether it's formal (e.g., healthcare, law) or informal (e.g., productivity tools)
- Three roles that **domain experts** can play:
 - bake domain expertise directly into your application (the **Oracle**)
 - define quality so you can measure and improve it (the **Evaluator**)
 - design systems that improve and learn over time (the **Architect**)
- The appropriate approach depends on your use case and scale.
- Each approach requires different adjacent skills in the **domain expert**.
- My advice is to hire a **Principal Domain Expert** early (with relevant **domain expertise** and a **breadth of adjacent skills**), give them ownership, start as an Oracle and evolve their role as your product and org evolves.



Thank you

Dr Christopher Lovejoy, MD

hi@chrislovejoy.me

www.chrislovejoy.me