

Facing operational reality

A leader's guide to building connected intelligence across the health plan



Most health plans don't have an AI shortage—they have an AI coordination problem. Rather than adding isolated tools to existing workflows, leading organizations are redesigning how intelligence moves across the enterprise. This guide provides a practical framework for evaluating enterprise AI, sequencing investments, and building the connected intelligence needed to improve clinical, operational, and financial performance.

Five executive takeaways

Before evaluating enterprise AI platforms or planning your next investment, consider these five ideas every health plan leadership team should align on.

- 1

Enterprise AI is an operating model—not a collection of AI tools
- 2

Redesigning from within beats retrofitting—the highest-value path runs through your existing systems and teams, not around them
- 3

Connected intelligence creates compounding value
- 4

Evaluate platforms—not point solutions
- 5

Strategically sequence investments to build a scalable intelligence layer

What you'll learn

Section	Executive question answered
Enterprise AI has entered a new phase	Why do we need to rethink our AI strategy?
The operating model behind enterprise AI	What actually separates enterprise AI from point solutions?
Evaluating an enterprise AI platform	How should we evaluate vendors?
Sequencing enterprise AI investments	Where should we start—and what comes next?
How Cohere Health supports a connected enterprise AI	How does Cohere Health support this strategy?

Enterprise AI has entered a new phase

Today, AI is embedded across nearly every health plan, supporting functions from prior authorization and claims to care management and payment integrity. At the same time, health technology vendors have largely rebranded their solutions as "agentic AI," making enterprise AI appear more accessible than ever.

Yet widespread adoption has not translated into enterprise intelligence. Instead, many organizations have accumulated disconnected AI capabilities that optimize individual workflows while creating new silos of data, governance, and operational knowledge. For health plan leaders, the real challenge is how to build an enterprise strategy in which each investment sets the groundwork for the next.

Why most AI initiatives never become enterprise transformation

Typical AI Pilot	Enterprise AI Platform
Solves one workflow	Creates reusable capabilities
Fixed deployment model	Flexible deployment that fits existing environments
New integration for every use case	Common intelligence layer
ROI measured per project	ROI compounds across functions
Expansion requires another implementation	Expansion builds on existing infrastructure

One of the biggest misconceptions in enterprise AI is that assembling the best point solution for every workflow will naturally create the best enterprise outcome. In reality, every standalone implementation introduces another data model, governance process, and integration point, making AI harder—not easier—to scale across the organization.

Leading health plans are taking a different approach. Rather than treating every AI initiative as a separate project, they're building a **shared intelligence foundation** where clinical context, policies, and operational knowledge can be reused across functions. As a result, each new deployment builds on the value of the last, rather than starting from scratch.

What actually separates enterprise AI from point solutions?

The difference between point AI and enterprise AI is how intelligence moves across the organization. Every day, health plans generate valuable clinical knowledge through prior authorization, appeals, payment integrity, care management, and quality operations.

When every AI application operates, those insights stop at the boundaries of each workflow. Teams repeat work, decisions become inconsistent, and every new implementation requires its own foundation. Connected intelligence changes that model by creating a system where every decision strengthens the next.

How connected intelligence compounds



Connected intelligence doesn't replace individual workflows—it strengthens them by ensuring every interaction contributes to a shared enterprise knowledge base.

Workflow	Enterprise value created
 Prior authorization	AI extracts clinical evidence, policy interpretation, and utilization insights.
 Payment integrity	Clinical context supports more accurate claim validation and reduces unnecessary rework.
 Quality	Existing documentation helps identify quality opportunities without requiring separate chart retrieval efforts.
 Appeals	Existing evidence and policy logic are reused rather than recreated, accelerating reviews and improving consistency.
 Care management	Existing clinical context—member history, prior auth decisions, and care gaps—is automatically surfaced, so care managers can focus on outreach and coordination, rather than locating and synthesizing information.

Enterprise AI is defined not by the sheer number of capabilities deployed, but by how effectively intelligence is leveraged and scaled across the organization.

Enterprise AI success depends less on individual AI models than on the platform supporting them. As health plans evaluate enterprise AI investments, three characteristics separate platforms built for long-term scale from solutions designed for individual use cases.

Evaluation criteria #1: Trust the AI
Can you trust the AI in production?



Clinical foundation	Proven performance	Human-led automation
Models trained on real clinical decisions—not generic AI. Designed to understand unstructured clinical records and apply clinical policies, guidelines, and medical necessity criteria.	Benchmarked against expert reviewers in real workflows before deployment. Performance validated against clinical and operational outcomes—not just model benchmarks.	AI is designed to augment—not replace—clinical expertise. Clearly defined operating boundaries determine where AI acts, where it assists, and where human judgment is essential, enabling clinicians to work at the top of their license.

Things to consider before deployment


- ✓ Confirm the AI is trained on real clinical decisions—not adapted from generic AI
- ✓ Validate performance against expert reviewers
- ✓ Define where AI acts independently, where it assists, and where human expertise is required
- ✓ Plan for clinical and operational change management from the outset, not after launch

Evaluation criteria #2: Scale the platform


Will every deployment make the next one easier?

The most powerful platforms don't just automate workflows—they build a connected intelligence foundation that bridges clinical, operational, and decision support functions. As new use cases emerge, these platforms should reduce complexity, strengthen governance, and multiply enterprise value—not create new silos.

Evaluate these capabilities



Questions every enterprise AI platform should answer



- How is clinical context shared across different use cases?
- What capabilities are reused versus rebuilt for each deployment?
- How does intelligence move between workflows?
- How does your platform prevent new AI deployments from creating additional silos?
- What becomes easier after the first implementation?
- How does enterprise value increase as additional use cases are added?

Evaluation criteria #3: Govern with confidence

Can you scale AI safely?

While enterprise AI governance can be perceived as a constraint on innovation, the healthcare sector presents a contrasting reality. In this context, robust governance frameworks are not an obstacle, but a critical enabler. The right model—one based on auditability, human oversight, and gradual expansion—allows AI to earn trust and take on more responsibility over time.

1 Auditability

Every AI-supported recommendation is traceable to supporting evidence and decision rationale.

2 Calibrated human oversight

AI is designed to augment—not replace—clinical and operational expertise. Decision thresholds define when AI acts independently, when it assists, and when human judgment is essential.

3 Validation before expansion

AI expands only as accuracy thresholds are met and benchmarks confirmed. Routine administrative work may be automated first, while more complex clinical decisions continue to incorporate human expertise.

4 Continuous production monitoring

Performance should be continuously monitored for accuracy, drift, bias, and operational impact.

Things to put in place before scaling

-  Define decision lineage requirements—every AI output should be traceable to its source before deployment.
-  Set accuracy thresholds and volume requirements that trigger expansion of AI responsibility per use case.
-  Establish production monitoring as a standing operational activity, not a post-launch response.

Where should we start—and what comes next?

Enterprise AI isn't implemented all at once. Leading organizations sequence investments, so each deployment expands the value of the next. Rather than solving isolated operational problems, they build an intelligence foundation that compounds over time.



PHASE 1: Establish the foundation

Start with a known pain point and a well-defined use case where AI can have a significant impact and accuracy can be benchmarked against human performance.

Establish governance from day one—auditability, oversight calibration, and validation thresholds before launch.

Build internal confidence so that clinical leaders, auditors, and operational teams can see the AI perform in production before they advocate for expansion.



PHASE 2: Expand connected intelligence

Expand into adjacent functions that share clinical context, such as appeals, payments, care management, and quality.

Each deployment should reuse the same intelligence foundation, rather than introduce a new one.

Track compounding value: measure ROI not just within each use case, but across functions as intelligence is shared.



PHASE 3: Operate at enterprise scale

Establish a consistent clinical and operational foundation across functions, including UM, appeals, care management, quality, claims, and payment integrity.

AI can support appropriate administrative tasks with appropriate human oversight, while complex clinical decisions continue to incorporate human expertise.

The objective isn't simply enterprise-wide AI adoption—it's operational coherence, where every AI-supported decision strengthens the clinical intelligence available to the next workflow.


How Cohere delivers connected enterprise AI

Cohere Health was purpose-built around the same evaluation criteria outlined in this guide. Rather than offering disconnected AI tools, Cohere Unify™ provides a connected intelligence platform that helps health plans trust AI, scale it across functions, and govern it responsibly.

Where other AI vendors offer general-purpose automation or point solutions that solve individual tasks, Cohere Health delivers connected clinical intelligence—trained on real clinical decisions, deployed to meet the health plan where it is, and governed through a progressive model where AI responsibilities expand as performance is validated.

Evaluation criteria	How Cohere Unify™ delivers
Trust the AI	Precision AI trained on 15M+ real clinical decisions. 27M unique patient profiles. 7,500+ codified policies. Designed to work alongside existing systems and clinical workflows.
Scale the platform	A connected intelligence platform that supports clinical context, policies, and operational knowledge to be leveraged across utilization management, appeals, care management, quality, claims, and payment integrity. Composable AI capabilities allow each new use case to build on the last, rather than start from scratch.
Govern with confidence	Full decision lineage on every output. Progressive autonomy governance model—AI earns autonomy as accuracy thresholds are met. Continuous performance monitoring across 34,000+ daily UM transactions. HITRUST r2 certified.

Proven Enterprise Impact*



*Based on 2026 Cohere Health data

One platform. Enterprise-wide impact.

Enterprise AI delivers maximum value when intelligence seamlessly integrates across core health plan workflows. Rather than deploying siloed AI solutions, Cohere Unify™ establishes a connected intelligence foundation—bridging utilization management, appeals, care management, quality, payment integrity, and claims so that every new capability builds on the last.

Operational area	AI capabilities	Business value
Utilization management	• Policy digitization • Intelligent intake • Clinical review support	Faster, more consistent utilization management while maintaining clinical oversight.
Payment integrity & claims	• Claims intelligence • Anomaly detection • Payment validation	Improve payment accuracy, increase processing efficiency, and reduce manual investigation.
Quality	• Clinical data ingestion • Gap closure identification • Exclusion review prioritization	Support digital quality measurement while reducing manual chart abstraction.
Appeals	• Intake digitization • Case summarization • Review note drafting	Accelerate appeal reviews and reduce manual administrative effort by reusing existing clinical evidence.
Care management	• Case summarization • Gap assessment • Care plan drafting	Supports care managers to spend more time engaging members and less time documenting care.

Redesign—not retrofit

Enterprise AI isn't achieved by adding another tool to the stack. It requires a connected intelligence foundation in which clinical context, governance, and AI capabilities are shared across workflows, so that every new deployment strengthens the next.

Schedule a working session

Whether you're evaluating your first enterprise AI investment or looking to expand an existing one, our team can help you identify the highest-value starting point and build a roadmap for connected intelligence.

