

October 2026 | Q4

The Operating Rhythm of the Modern Health Plan

How a Decision Enterprise coordinates decisions across different clocks without centralizing the enterprise

Letter from the Editor

When should a Health Plan make a decision?

In August, I wrote that a health plan does not operate on one calendar.

Actuarial teams price populations that may not enroll for months. Product teams design benefits for a future plan year. Clinical teams make decisions whose value can change in hours. Risk adjustment and quality teams evaluate evidence from care already delivered. Technology leaders make investments whose consequences may unfold over years.

I described these as the **many clocks of the health plan**.

The clocks do not need to run at the same speed. They shouldn't.

But that leaves a more difficult operating question:

How should a modern health plan operate when different decisions need to happen at different speeds?

That question has become more important as artificial intelligence moves deeper into healthcare workflows. Technology can increasingly acquire evidence, interpret information, generate recommendations and initiate action faster than the organizations using it have historically been able to operate.

That sounds like an argument for speed.

I don't think it is.

The objective should not be to make every decision as quickly as technology allows. Some decisions lose value when they are delayed. Others lose legitimacy when they are made before the evidence is sufficient, before appropriate judgment has been applied, or by someone—or something—without the authority to decide.

The operating objective is therefore more precise:

Make the earliest appropriate decision.

That requires a health plan to understand three things about every consequential class of decision:

- Is the evidence sufficient?
- Is the authority established?
- Is the decision window still open?

The answers will differ across functions. They should.

What should increasingly be common is the enterprise capability to support those decisions—and to learn from what happens next.

Last month, I called the emerging operating model capable of doing this the **Decision Enterprise**.

October asks the next question:

How does that enterprise actually operate?

—John G. Murtha

Section I – One Enterprise. Many Clocks.

In August, I argued that health plans operate across multiple decision horizons.



Care, enrollment, pricing, revenue and enterprise capability all operate on different clocks.

That is not evidence of dysfunction. Different decisions legitimately require different operating rhythms.

An expedited prior-authorization decision should not operate on the same clock as a network strategy. A care-management intervention may have value today that it will not have next month. A benefit decision may appropriately require months of analysis before affecting members a year later.

The mistake is assuming that coordination requires synchronization.

It does not.

The operating problem begins when the organization cannot move appropriately from evidence to decision to action while the opportunity to create value remains.

Evidence may exist but remain inaccessible to the function that needs it. A recommendation may be ready while authority remains unclear. A decision may be correct but arrive after the opportunity to act has passed.

The important question, then, is not simply how quickly the enterprise can move.

It is whether the enterprise can make an appropriate decision **while that decision still has value.**

That brings us from the many clocks of the health plan to something more precise:
the decision window.

Section II – Inside the Decision Window

Every consequential decision has a window.

That window does not begin merely because information exists. It begins when sufficient evidence exists to responsibly make the decision and the appropriate authority can act.


And it does not remain open indefinitely. At some point, the value of the decision begins to decay.

That gives us a more useful definition:

A decision window is the interval in which sufficient evidence exists, the appropriate authority can act, and the value of the decision has not materially decayed.

Inside the Decision Window

A decision is successful when sufficient evidence exists, the appropriate authority can act, and the value of the decision has not materially decayed.



This creates two different operating failures.

A decision made **too early** may lack sufficient evidence, required judgment or appropriate authority.

A decision made **too late** may be correct and still have lost much of its value. The opportunity to intervene may have diminished, downstream work may have accumulated, or another workflow may already have moved forward.

The operating objective lies between the two:

Not the earliest possible decision. The earliest appropriate decision.

Prior authorization provides an unusually clear example because both boundaries of the decision window are becoming explicit.

Since January 2026, CMS has required impacted payers, other than QHP issuers on the federally facilitated exchanges, to make prior-authorization decisions for covered medical items and services as expeditiously as the patient's condition requires, subject to maximum timeframes of **72 hours for expedited requests and seven calendar days for standard requests**. CMS also requires public reporting of prior-authorization metrics, including average time from submission to decision.

Decision latency, in other words, is no longer merely an internal service-level concern. In prior authorization, it has become an externally measurable operating requirement.

But the other boundary matters too.

CMS guidance for Medicare Advantage permits algorithms to assist coverage determinations while limiting their ability to substitute population-level predictions for the circumstances of the individual patient.

Together, those requirements illustrate the operating challenge: **a decision can arrive too late, but technical ability to produce an answer does not by itself establish authority to make the decision.**

That distinction extends far beyond prior authorization.

As technology becomes capable of producing answers earlier, health plans will increasingly have to determine when those answers are sufficiently supported, appropriately authorized—and still valuable enough—to act upon.

Section III – Not All Latency Is the Same

The word *latency* usually sounds like something to eliminate. Operators should be more careful. Time between evidence and decision can represent two very different things.

Some latency improves or legitimizes the decision. Some reflects the cost of organizational fragmentation.

A qualified clinician taking time to interpret a complex case is not equivalent to a team reconstructing information that already exists. Deliberation over a consequential network or pricing decision is not equivalent to work sitting in a queue because decision authority is unclear.


Governance is not the same as bureaucracy. Judgment is not the same as delay.

The distinction suggests a useful question:

What happened during the wait—and did it improve the decision?

Not All Latency Is the Same

Time between a signal and a decision can reflect appropriate judgment and control — or unnecessary friction and waste. Great operators know the difference.



Some latency is the cost of appropriate authority and judgment. Some is the cost of fragmentation.
The operator's job is to tell them apart.

If additional time produced better evidence, necessary judgment, appropriate governance or deliberate consideration of consequences, that latency may be part of good decision-making.

If the time was consumed reconstructing existing evidence, repeating reviews, reconciling definitions, waiting for another transaction or determining who owns the decision, the latency is doing something else.

It is revealing the operating model.

The distinction matters because speed alone tells us very little about decision quality.

In a 2022 review of Medicare Advantage prior-authorization denials, the HHS Office of Inspector General found that 13% of the denied requests it examined met Medicare coverage rules and likely would have been approved under original Medicare. In some cases, MA organizations cited insufficient documentation even though OIG reviewers concluded that the existing medical records supported medical necessity.

The lesson is not that those decisions simply needed to happen faster. It is that reducing decision time is valuable only if the organization preserves the evidence, judgment and authority necessary to make the decision well.

A fast decision based on inadequate evidence or inappropriate criteria is not operating excellence.

Neither is a well-supported decision delivered after its value has disappeared.

Some latency is the cost of appropriate authority and judgment. Some is the cost of fragmentation.

The operator's job is to tell them apart.

Section IV – AI Changes the Speed. It Does Not Eliminate the Decision.

Artificial intelligence makes the distinction between speed and good decision-making more urgent.

AI can interpret unstructured information, identify patterns across large bodies of evidence, generate recommendations, organize tasks and increasingly participate directly in workflows.

That compresses portions of the decision cycle.

Signal → Evidence → Analysis → Recommendation

can increasingly happen at machine speed.

But faster technology does not repair fragmented operating design.

A fragmented enterprise with faster technology is still fragmented. It simply fragments faster.

And technical capability does not confer organizational authority.

The question therefore is no longer simply:

Can AI make this decision faster?

It becomes:

What role should AI play in this decision, given the evidence required, the authority involved, the consequences of error, and the window in which action has value?

For some decisions, automation may be entirely appropriate.

For others, AI may assemble evidence, identify an exception or make a recommendation while accountable human judgment remains essential.

And for high-consequence or ambiguous decisions, deliberation may remain an intentional feature of the operating model.

This is why I increasingly think AI is less interesting as a standalone technology strategy than as a **stress test of organizational design**.

It exposes assumptions that were easier to avoid when information moved more slowly: **who owns a decision, what evidence is sufficient, which policies govern it, what explanation is required, who is accountable for the outcome, and when the decision should occur.**

Those questions existed before AI.

AI simply makes it harder to avoid answering them.

Section V – Centralize the Capability, Not the Decision

If different decisions appropriately operate at different speeds and under different authorities, the answer cannot be to centralize every decision.

Nor should every function continue rebuilding the capabilities underneath its decisions independently.

There is a middle ground:

Centralize the capability, not necessarily the decision.

Care Management, Utilization Management, Risk Adjustment, Quality, Network, Finance and Member Operations will continue to require specialized workflows, judgment, decision rights and accountability.

They should.

But specialization does not require each function to maintain a separate version of the enterprise beneath it.

Centralize the Capability, Not the Decision

Different functions retain their workflows, judgment and decision rights.
The enterprise shares the foundation that makes trusted decisions possible.



If the same clinical evidence influences Care Management today, Quality next month and Risk Adjustment later in the year, the enterprise should not have to rediscover, reinterpret or reconstruct that evidence each time.

And increasingly, it should preserve more than the evidence itself.

What decision was made, under what policy, by whom—or what—and when? What action followed? What happened afterward?

Those questions turn a decision from a moment in a workflow into something the enterprise can observe.

And once the enterprise can connect the evidence behind a decision to the action that followed and the outcome that resulted, something more important becomes possible:

Learning.

Different functions can retain the authority and specialized workflows appropriate to their purpose while the enterprise develops a common memory of how decisions were made and what happened next.

The enterprise does not need to own every decision.

It needs to preserve what the decision was made from, how it was made, and what happened next.

Section VI – The Operating Rhythm of the Decision Enterprise

Put the pieces together and the operating model begins to look different.

The modern health plan does not need every function moving at the same speed. It needs the ability to support each consequential decision on the clock appropriate to it.

That means being able to answer, repeatedly:

- What evidence must exist?
- Who—or what—has authority to decide?
- When does the decision begin losing value?
- Where must the decision enter workflow for action to occur?
- What outcome should return so the enterprise can learn?

The answers will differ by decision.

The underlying discipline does not.

This is the operating rhythm of the Decision Enterprise:

Different clocks. Shared capabilities. Explicit authority. Observable decisions. Continuous learning.

The objective is not to make the organization move faster everywhere.

Speed without sufficient evidence can create error. Speed without authority can create risk. And speed inside fragmented workflows can simply accelerate fragmentation.

A Decision Enterprise operates differently.

It becomes faster **where speed creates value**, deliberate **where judgment creates value**, and increasingly automated **where evidence and authority allow it**.

Different decisions can move at different speeds without becoming disconnected decisions.

That is a more useful definition of operational agility.

The Operator's Test

Most leadership teams can tell you how quickly important systems process information. Fewer can answer how long it takes the organization to act once sufficient evidence exists. For a consequential decision, I would ask five questions:

1. What evidence must exist before this decision should be made?
2. Who—or what—has the authority to make it?
3. When does the decision begin losing value?
4. What portion of today's latency improves the decision—and what portion reflects fragmentation?
5. Can we observe what happened and use the outcome to improve the next decision?

If those questions cannot be answered clearly, the problem may not be technology.

It may be the operating model.

The Operator's View

Health plans will continue to operate across different clocks.

They should.

The advantage will not belong to the organization that forces every function onto a common cadence or automates every decision technology can reach.

It will belong to the organization that understands **which capabilities should be shared, which decisions should remain specialized, and when each decision should appropriately move from evidence to action.**

The objective is not maximum speed.

It is something harder:

The right decision.

With the right evidence.

By the right authority.

At the right time.

And then learning from what happened.

That is the operating rhythm of the modern health plan.

About The Health Plan Operator's Brief

A monthly executive briefing examining the strategic, operational and technological forces reshaping health plans. Written to inform discussion—not to promote products or predict the future. The views expressed are my own.

—John G. Murtha

Continue reading The Health Plan Operator's Brief

New issues and occasional ideas and frameworks on the future operating model of health plans.

[Subscribe at johngmurtha.com](http://johngmurtha.com)