THE HEALTH PLAN OPERATOR'S BRIEF Issue No. 005

November 2026 | Q4

Institutional Intelligence

How a health plan learns from its own decisions – and gets better over time.
An executive briefing on the emerging operating model of health plans.

Letter from the Editor

Health plans make thousands of consequential decisions every day.

Which members need intervention. Which services require authorization. Which providers should be in the network. Which quality gaps demand attention. Which markets deserve investment. Which policies should change.

Increasingly, those decisions are supported by more information, better analytics and artificial intelligence. But having more information does not necessarily mean an organization is getting better at deciding what to do next.

That distinction has been on my mind.

We spend a great deal of time designing systems that help organizations make decisions. We spend less time considering what happens after the outcome is known.

Did the decision work? Why? Would we make the same decision again?

And perhaps most importantly: does the organization actually remember what it learned?

Health plans already preserve extraordinary amounts of information about what happened: claims, clinical records, authorizations, member interactions, transactions, utilization, costs and outcomes. Learning requires something more.

It requires preserving enough of the decision itself—the evidence available at the time, the context in which it was interpreted, the policy or model applied, the authority exercised, the action taken and the outcome that followed—to understand not simply what happened, but why.

That becomes increasingly important as artificial intelligence makes information, analysis and prediction easier to obtain.

The scarce capability may not be knowing more. It may be learning better.

This edition of *The Health Plan Operator's Brief* explores what happens when the outcome of one decision becomes evidence for the next.

It considers how organizational memory makes learning possible, why experience alone does not make an enterprise smarter, and what begins to happen when learning compounds across thousands of decisions.

Because there is a difference between an organization with experience and an organization that learns from experience.

The first accumulates history.

The second gets better.

© 2026 John G. Murtha. All rights reserved.

Section I — An Industry With No Shortage of Experience

The health plan industry is not suffering from a shortage of experience.

If anything, the past several years have provided an extraordinary amount of it.

Plans have navigated a pandemic, Medicaid redeterminations, rapidly changing Medicare Advantage economics, persistent medical cost pressure, regulatory changes, shifting utilization patterns and increasingly demanding expectations from members, providers and government.

The third quarter added another chapter.

Several national plans reported improving performance and greater confidence in their earnings trajectories. Medicare Advantage Star Ratings, however, remained uneven. Ratings slipped for several large plans while Humana reported a sharp recovery in the share of its Medicare Advantage membership expected to be enrolled in four-star-or-higher plans in 2027. At the same time, Medicaid margins remain under pressure, medical-cost performance continues to demand attention and plans continue to reconsider markets where they cannot see a path to sustainable performance.

Those are not simply financial developments.

They are the outcomes of thousands of decisions made over time—about pricing, benefits, networks, clinical programs, market participation, provider relationships, member engagement, capital allocation and operations.

And the responses are already producing another generation of decisions.

Plans are adjusting Medicaid markets and seeking better alignment between rates and underlying acuity. They are changing Medicare Advantage portfolios. They are tightening oversight of expensive therapies. They are introducing artificial intelligence into prior authorization and payment integrity. They are changing where and how care is delivered. They are becoming more selective about where they compete and where they invest.

The operating question is therefore larger than whether any individual response is correct.

What does the enterprise retain from all of this experience?

Consider Medicare Advantage.

A plan can know that a market underperformed, that a product lost money, that utilization exceeded expectations or that a Star Rating deteriorated. It can respond by changing benefits, networks, pricing, clinical programs or market participation.

But the outcome alone does not tell the organization which underlying decisions were sound, which assumptions proved wrong, which interventions worked, or which lessons should influence the next product, market or member population.

The same problem appears at a much smaller scale.

A member does not respond to an outreach strategy. A prior authorization is overturned. A care-management intervention fails to change utilization. A quality gap remains open. A provider relationship produces different outcomes than expected.

Each produces more data.

Each also creates an opportunity to learn.

© 2026 John G. Murtha. All rights reserved.

The difference is whether the organization can connect the outcome back to the evidence, assumptions, decisions and actions that preceded it—and then carry what it learns into the next decision.

That is where experience begins to become an enterprise capability.

Section II – Closing the Loop

In earlier editions of The Health Plan Operator's Brief, I described the operating rhythm of a Decision Enterprise as a progression:

SIGNAL → EVIDENCE → DECISION → ACTION → OUTCOME → LEARNING ↺

A need, risk or opportunity is detected. Relevant evidence is assembled. A decision is made with the appropriate authority. The decision is translated into action. An outcome becomes observable. And what the enterprise learns becomes part of the evidence available the next time a similar decision must be made.

The cycle has always ended in learning. But that final step deserves more scrutiny. Because "learning" can easily become the least examined word in the model.

Measuring an outcome is not the same as learning from it.

A dashboard can show that utilization increased. A quality measure can show that performance declined. An authorization system can record an overturn. A financial report can show that a market missed its target.

Those observations tell the organization what happened.

Learning begins when the organization compares what it expected to happen with what actually happened, examines why they differed, and uses that understanding to change what happens next.

That review is the critical gate between experience and learning. Without it, the next cycle may begin with more data but essentially the same organizational knowledge.

With it, the outcome of the first decision becomes evidence available to the second.

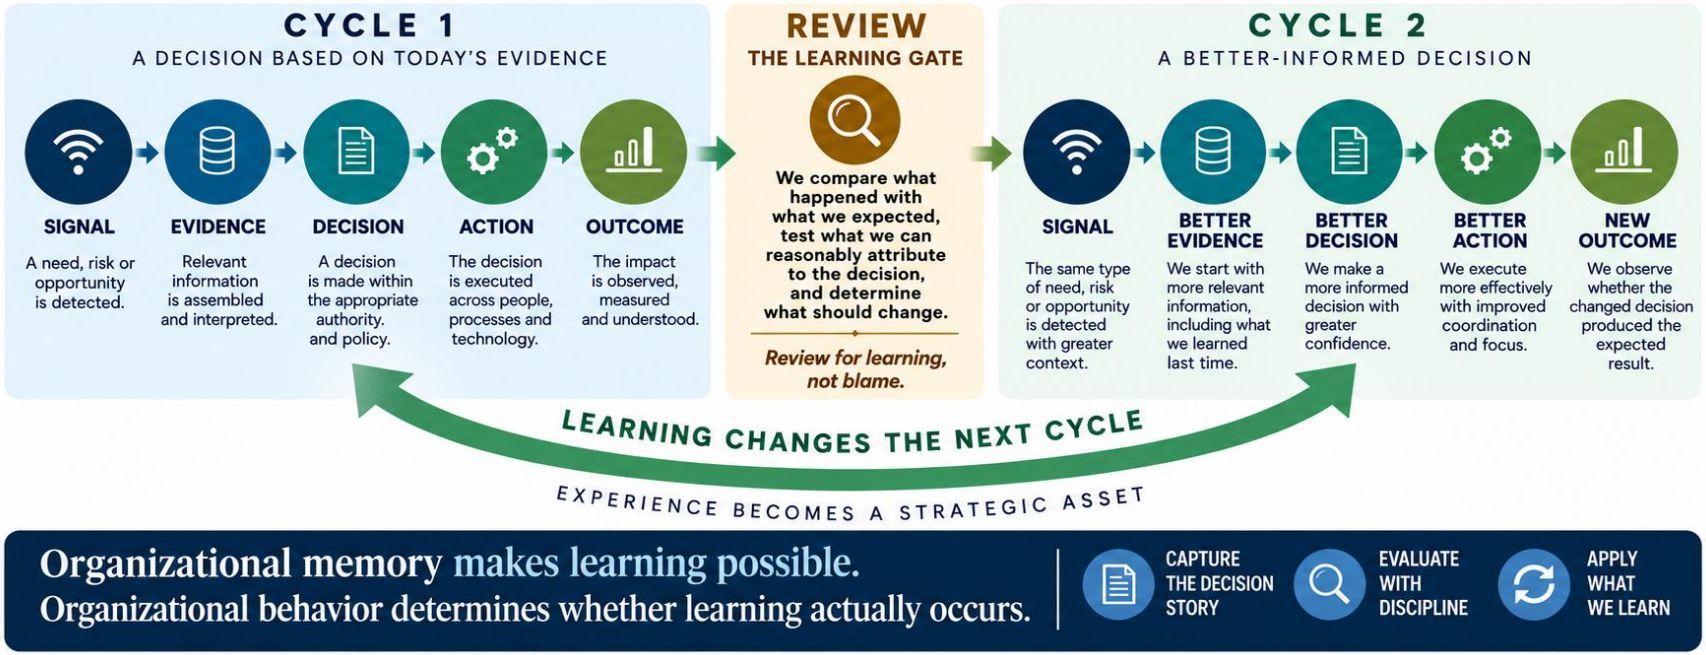
The structure of the decision does not necessarily change. The quality of the decision can.

© 2026 John G. Murtha. All rights reserved.

The Learning Loop

FROM INFORMATION TO BETTER DECISIONS

An organization doesn't learn simply because it has more information.
It learns when it consistently reviews outcomes, evaluates what it learned, and applies those insights to the next decision.



Organizational memory makes learning possible.
Organizational behavior determines whether learning actually occurs.

CAPTURE THE DECISION STORY | EVALUATE WITH DISCIPLINE | APPLY WHAT WE LEARN

© 2026 John G. Murtha. All rights reserved.

The second cycle is not better simply because more time has passed or more information has accumulated.

It is better because the organization has deliberately carried something forward from the first.

Evidence can become more relevant. Decisions can be made with greater confidence. Actions can become better coordinated. Outcomes can improve.

Experience begins to compound.

But that raises a more difficult question.

What must an organization preserve in order to learn?

Most health plans already possess extensive records of what happened.

Far fewer preserve the complete story of how the decision happened.

That is where organizational memory becomes important.

Section III – From Data to Organizational Memory

Health plans are not short of data.

Claims record diagnoses, procedures, utilization and payments. Clinical records contain encounters, observations, laboratory results and notes. Authorization systems capture requests and dispositions. Member systems record interactions. Enrollment and eligibility systems preserve transactions. Quality, financial and operational systems measure outcomes.

Together, they create an extraordinarily rich record of activity.

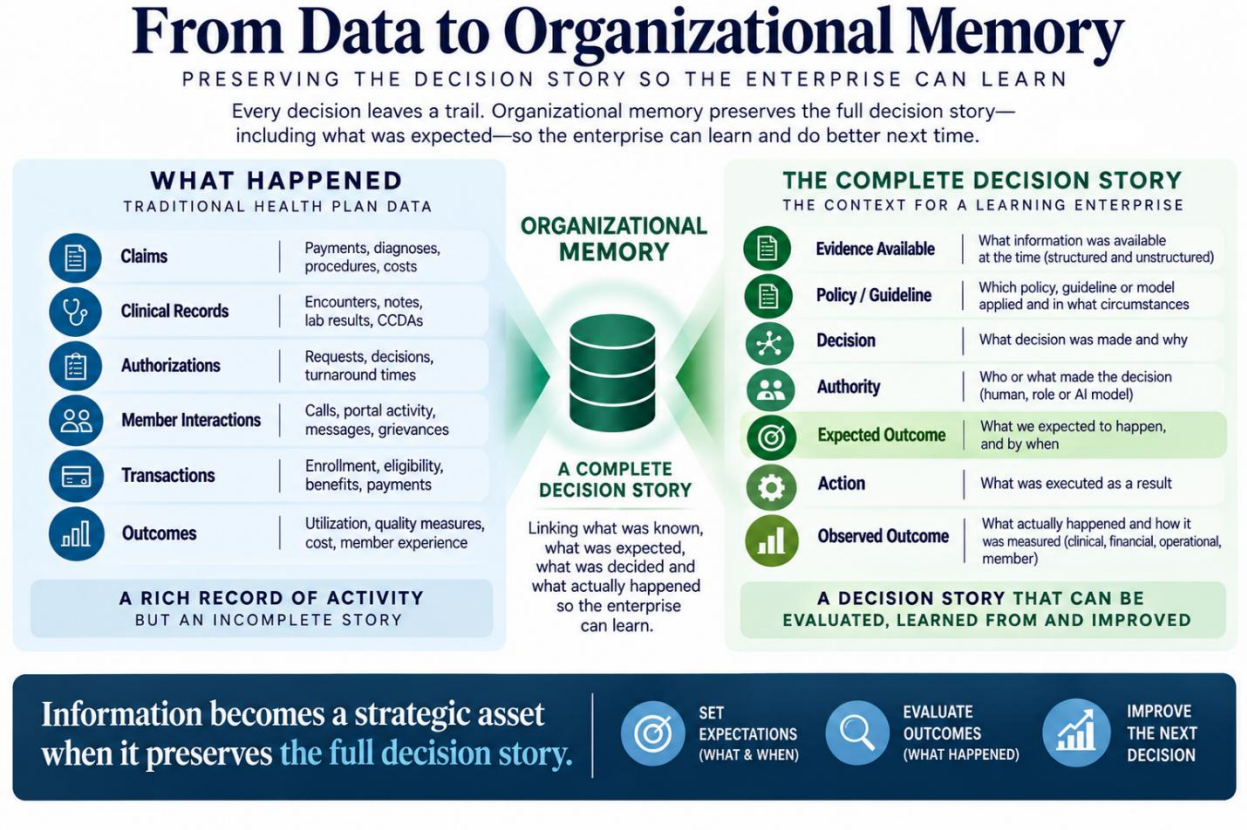
But activity is not the same as memory.

If an organization wants to learn from a decision, knowing what happened afterward is not enough.

It also needs to reconstruct the decision itself.

- What evidence was available when the decision was made?
- What policy, guideline or model applied?
- What decision was made, and why?
- Who (or what) had the authority to make it?
- What action followed?

© 2026 John G. Murtha. All rights reserved.



- What outcome was expected?
- And when was that outcome expected to become observable?

That last pair of questions matters.

Without an expectation recorded at the time of the decision, hindsight can easily substitute for learning. A favorable outcome can make a weak decision look sound. An unfavorable outcome can make a reasonable decision look wrong.

The organization needs enough of the original decision context to distinguish judgment from luck. That is the difference between a record of activity and organizational memory.

Organizational memory preserves the decision story: the evidence, context, expectation, authority, action and eventual outcome in a form that can be examined later.

It does not require preserving every piece of information forever. It requires preserving enough of the decision to make meaningful review possible.

© 2026 John G. Murtha. All rights reserved.

Section IV – AI Makes Memory More Important, Not Less

Artificial intelligence is rapidly increasing the ability of organizations to process information, identify patterns, generate predictions and support decisions.

Health plans are already applying AI across decision-intensive processes. Prior authorization is becoming faster and more automated. Payment-integrity models can identify potential errors and anomalies at scale. Clinical and operational data can be interpreted more quickly. AI-enabled systems can increasingly prioritize work, recommend actions and support decisions that once required substantially more manual effort.

These developments can create enormous value.

They also make organizational memory more important.

As more decisions become machine-assisted, the enterprise needs to preserve more than the recommendation or result.

It needs to know what evidence was available, which model or policy was applied, what the model recommended, who or what had authority to act, what action followed, what was expected to happen and what actually occurred.

That is decision provenance.

Without it, automation can increase the number and speed of decisions without increasing the organization's ability to understand them.

A model can perform well overall and poorly for a particular population. A recommendation that worked under one set of circumstances may perform differently as utilization, clinical practice or member behavior changes. An automated action can produce an unintended consequence that becomes visible only later.

If the enterprise preserves only the outcome, it may know that performance changed.

It may not know why.

And if it cannot connect that outcome back to the evidence, model, decision and action that produced it, it has little basis for determining what should change next.

This creates an important distinction.

AI can help an organization make a decision.

A learning system helps the organization determine whether that decision worked.

The two capabilities should reinforce one another.

© 2026 John G. Murtha. All rights reserved.

Better models can improve decisions. Better organizational memory can make those decisions more observable. Disciplined review can determine where the model performed as expected, where it did not and what should change.

Over time, the outcome of one AI-assisted decision can become evidence for improving the next.

Without that loop, an organization may become very good at making more decisions, faster.

It may also create organizational amnesia at greater speed.

The opportunity is not simply to automate judgment.

It is to create an enterprise capable of learning from the consequences of both human and machine-assisted decisions.

And once that learning occurs not around a single decision, but across thousands of decisions throughout the enterprise, something more consequential becomes possible.

Section V — When Learning Compounds

One decision can produce a lesson.

Thousands of decisions can reveal a pattern.

That distinction matters because health plans do not make decisions in isolation. They make them continuously, across populations, markets, functions and time.

Care managers decide when and how to intervene with members. Utilization-management teams make coverage and medical-necessity decisions. Quality teams determine where to focus interventions. Network leaders make decisions about access, configuration and provider relationships. Risk-adjustment teams evaluate clinical evidence. Benefit leaders make choices intended to influence access, behavior and cost. Executives decide where to invest, compete and build new capabilities.

Each decision occurs within its own context.

But collectively, they create an extraordinary opportunity to learn.

Consider a care-management program.

One member's outcome may tell us relatively little. Across hundreds or thousands of members, patterns can begin to emerge.

Which members responded?

Which interventions were associated with better outcomes?

When did those outcomes become observable?

© 2026 John G. Murtha. All rights reserved.

Under what circumstances?

Which evidence was most useful in identifying the opportunity?

Which actions produced no measurable change?

Did the same intervention perform differently across populations, providers or markets?

Those patterns can inform the next decision.

But they introduce an important discipline.

A pattern is not necessarily a cause.

Members are not randomly distributed across most health-plan programs. Interventions may be directed toward people who are sicker, more engaged or otherwise different from those who do not receive them. Outcomes may appear months or years after the original decision. Market conditions, provider behavior, benefit changes and other interventions may change at the same time.

A learning enterprise therefore needs more than organizational memory.

It needs evaluation discipline.

For some decisions, a clear expectation recorded in advance and a disciplined retrospective review may be sufficient. For others, credible learning may require comparison groups, holdouts, matched populations, longitudinal analysis or other methods appropriate to the question.

The objective is not to turn every operating decision into a research study.

It is to match the strength of the conclusion to the strength of the evidence.

That distinction matters.

There is a difference between observing that an outcome followed a decision and concluding that the decision caused the outcome.

A sophisticated learning system should know the difference.

The same logic applies elsewhere.

A health plan can examine which prior-authorization decisions are most frequently overturned and why. It can evaluate which approaches to closing quality gaps are associated with durable changes in care. It can identify where network configurations appear to improve access and where they create friction. It can test whether benefit changes alter member behavior or simply redistribute cost.

The opportunity becomes more consequential when learning crosses functional boundaries.

A pattern observed in utilization management may have implications for network strategy. A care-management outcome may change how a population is identified. Clinical evidence

© 2026 John G. Murtha. All rights reserved.

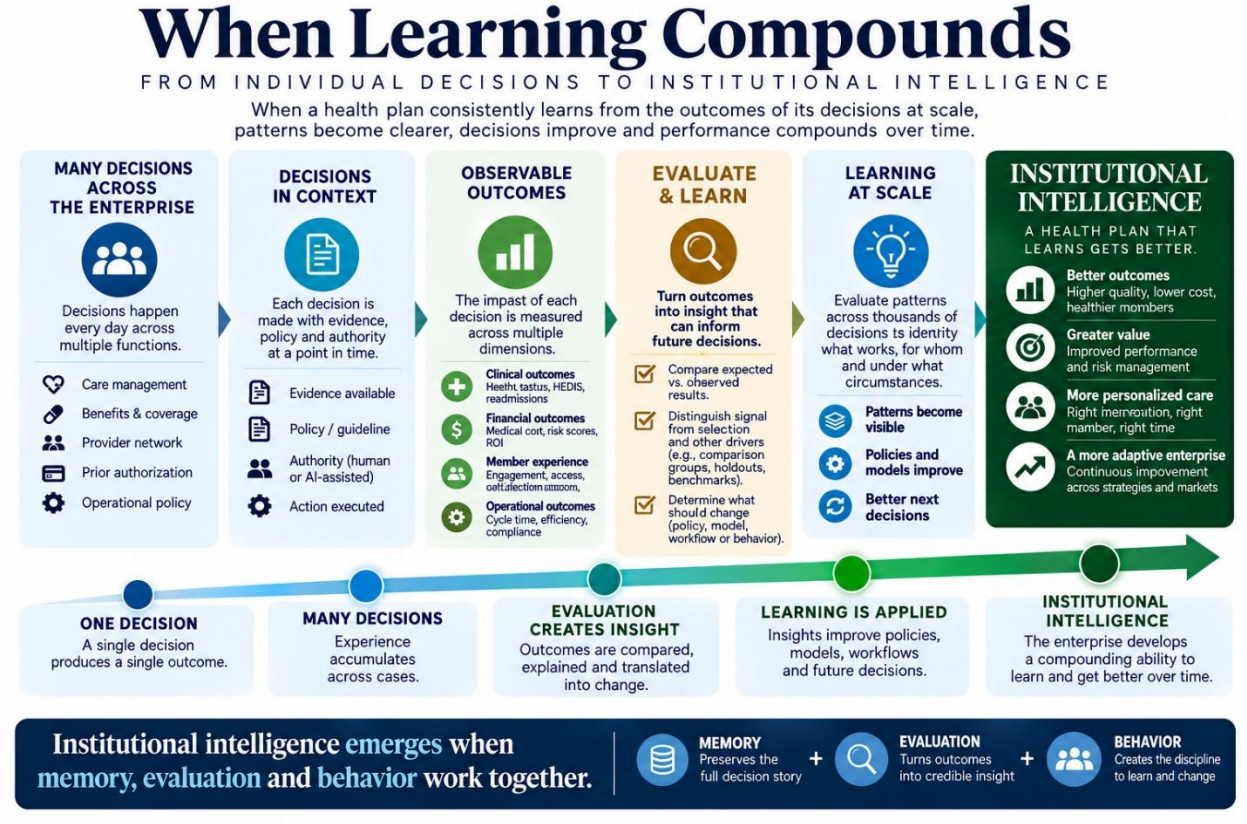
used for risk adjustment may reveal opportunities for quality or care management. Member experience may expose an access problem that claims alone would not explain.

The enterprise begins to learn not simply from individual workflows, but from the relationships among them.

The process becomes cumulative:

- Decisions produce outcomes.
- Outcomes create evidence.
- Evidence reveals patterns.
- Evaluation tests what those patterns mean.
- What survives that evaluation changes future decisions.

That is how experience begins to compound.



The progression illustrated above changes the nature of the capability.

At first, the organization is learning whether an individual decision produced the expected result.

At scale, it begins identifying patterns across decisions—and, with appropriate evaluation, learning what works, for whom, under what circumstances and why.

© 2026 John G. Murtha. All rights reserved.

Policies can improve. Models can be refined. Workflows can change. Resources can be directed differently. Decisions can become more consistent where consistency creates value and more contextual where judgment is required.

The enterprise is no longer simply accumulating information about its operations.

It is developing the ability to become better because it has operated.

That is a different kind of capability.

It is institutional intelligence.

Section VI – Institutional Intelligence

Institutional intelligence is the capability of an enterprise to learn from the outcomes of its decisions and apply that learning to make better decisions over time.

It is not synonymous with artificial intelligence.

It is not a data platform, an analytics capability or a repository of institutional knowledge.

Each can contribute to it.

The idea has important antecedents in the learning organization and the learning health system. The distinction here is operational: institutional intelligence depends on preserving the decision story–the evidence available, the context and authority under which a decision was made, what was expected to happen, the action that followed and the outcome that resulted–so experience can change future decisions.

Institutional intelligence exists when experience changes future behavior.

Organizational memory preserves the decision story. Evaluation determines what can reasonably be learned from it. Organizational behavior determines whether that learning changes what the enterprise does next. Institutional intelligence emerges when all three work together.

For much of the health plan's history, information and learning have accumulated primarily within functions. Utilization management learns about utilization management. Quality learns about quality. Network management learns about networks. Care management learns about interventions.

Those functional capabilities remain essential.

But many of the most consequential health-plan problems do not respect functional boundaries.

Access can be influenced by benefits, network configuration, authorization, provider capacity, clinical circumstances and member behavior simultaneously. Medical-cost performance can reflect pricing, utilization, site of care, benefit design and clinical practice patterns. A quality

© 2026 John G. Murtha. All rights reserved.

outcome can depend on whether the right evidence reached the right person early enough to act.

Institutional intelligence does not require centralizing those decisions.

It requires the ability to learn across decisions that remain distributed throughout the enterprise.

The objective is not an omniscient central system deciding everything. Nor is it replacing professional judgment with an algorithm.

It is creating an enterprise in which the evidence generated by experience can travel far enough to improve the next consequential decision—wherever that decision appropriately belongs.

Repeated across populations, markets and operating cycles, improvements in decision quality can accumulate into meaningful differences in clinical outcomes, member experience, operating performance and financial results.

That is why learning should be understood as an operating capability rather than an organizational aspiration.

The value is not that the enterprise knows more about its past.

The value is that its past makes it better at navigating its future.

Which raises a harder question:

What prevents organizations with enormous amounts of experience from learning this way already?

Section VII — The Hardest Part of Learning May be Human

The infrastructure for organizational memory can be designed.

The behavior required for institutional learning cannot simply be installed.

Organizational memory makes learning possible. Organizational behavior determines whether learning actually occurs.

For an organization to learn from its decisions, someone eventually has to be willing to ask whether the decision worked.

Sometimes the answer will be yes.

Sometimes the outcome will be disappointing even though the decision was reasonable given the evidence available at the time.

And sometimes the organization simply made the wrong call.

That last possibility makes learning difficult.

© 2026 John G. Murtha. All rights reserved.

Accountability matters. Health plans operate in highly regulated environments where decisions can affect member health, provider relationships, financial performance and compliance.

But accountability and learning are not the same thing.

If every retrospective examination is experienced primarily as an attempt to identify who was wrong, people learn something very quickly:

Do not make uncertainty visible.

Assumptions become less explicit. Expected outcomes are documented less clearly. Decisions are reconstructed after the fact. Unexpected results are explained rather than examined.

The organization may preserve the data while losing the opportunity to learn from it.

A learning enterprise requires a different discipline.

At the time of a consequential decision, make the reasoning observable:

What evidence did we have?
What did we believe it meant?
What did we expect to happen?
When did we expect to know?

Then return when the outcome becomes observable:

What happened?
What did we get right?
What did we misunderstand?
What should we do differently next time?

In a regulated and litigious industry, making assumptions and uncertainty explicit also creates legitimate questions about documentation, privilege and discoverability. Those questions require appropriate governance. But avoiding the record entirely carries its own cost: an enterprise that cannot reconstruct what it believed at the time will struggle to distinguish judgment from hindsight.

Someone must also own the return. The appropriate cadence will differ by decision—days for some operational decisions, months or even years for others—but the review cannot depend on someone remembering to revisit it. For consequential, repeatable decisions, the operating model should establish who returns to the decision, when, and with what evidence.

The purpose is not to eliminate accountability.

It is to make accountability more useful.

Poor judgment, insufficient evidence, flawed execution, changing circumstances and an inherently uncertain outcome are different problems. They should produce different responses.

Technology can help preserve that history.

© 2026 John G. Murtha. All rights reserved.

Leadership determines whether the enterprise is willing to learn from it.

Without honest review, organizational memory becomes an archive.

With it, memory can become intelligence.

Section VIII – The Operator's Test

Choose one consequential decision your organization makes repeatedly.

Not the easiest decision to measure.

Choose one that materially affects members, providers, clinical outcomes, operating performance or financial results.

Then ask five questions.

1. Can we reconstruct the evidence available when the decision was made?

   Not what we know now.

   What did the decision-maker (human or machine) actually know at the time?

2. Can we identify what shaped the decision?

   What policy, guideline, model, assumption or judgment was applied?

   Who (or what) had the authority to decide?

3. Did we record what we expected to happen—and when?

   Without an expected outcome, retrospective analysis can easily become hindsight.

   A result is much easier to interpret when the organization knows what it originally believed would occur.

4. Did the action produce the intended clinical, financial, operational or member outcome?

   And what would have happened otherwise?

   Observing an outcome is not the same as demonstrating that the decision caused it. The more consequential the claim, the stronger the evaluation required to distinguish the effect of the decision from selection, changing conditions and chance.

5. Does what we learn change the next decision?

   This is the test that matters most.

   Does the evidence change a policy?

   Refine a model?

   Alter a workflow?

   Change an intervention?

© 2026 John G. Murtha. All rights reserved.

Improve the information available to the next decision-maker?

If the organization can answer the first four questions but not the fifth, it may have excellent information and substantial experience.

It may not yet have a learning system.

The objective is not to document every decision made across the enterprise.

It is to identify the decisions consequential enough (and repeated often enough) that learning from them can materially improve performance.

Start there.

Then close the loop.

Section IX – The Operator's View

Over the past several months, The Health Plan Operator's Brief has explored a simple question from several directions:

What operating model does the modern health plan require?

Capabilities create options.

Shared evidence creates a more consistent basis for action.

Explicit decision rights allow authority to remain where judgment belongs.

Operating rhythm allows different decisions to move at different speeds without fragmenting the enterprise.

But something is still missing if the outcome simply becomes another historical record.

The loop has to close.

SIGNAL → EVIDENCE → DECISION → ACTION → OUTCOME → LEARNING ↺

Learning makes the next cycle different from the last.

Not simply because the organization has more data.

Because it has more experience that it knows how to use.

That distinction will become increasingly important as artificial intelligence makes information easier to process, patterns easier to detect and recommendations easier to generate.

The advantage may therefore belong not simply to the health plan with the most data, the most technology or even the most experience.

It may belong to the organization that can most consistently convert experience into better decisions.

Capabilities create options.

© 2026 John G. Murtha. All rights reserved.

Evidence informs decisions.

Operating rhythm turns decisions into action.

Outcomes create experience.

Learning makes the enterprise better.

That is the promise of institutional intelligence.

There is a difference between an organization with experience and an organization that learns from experience.

The first accumulates history.

The second gets better.

***

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

© 2026 John G. Murtha. All rights reserved.