

GOVERNANCE WHITEPAPER · OCTOBER 21, 2026

Governing the decision: evidence, authority and AI agents

What chief executives and boards in senior housing & care
should ask before an AI recommendation becomes an action

Coherence before Intelligence.

For chief executives, boards, chief operating, financial and clinical officers, and the
owners and lenders who rely on their reporting.

CONTENTS

	Executive summary	3
1	The question capability cannot answer	4
2	Why the evidence disagrees.....	5
3	A recommendation is not permission.....	7
4	A governance sequence from evidence to execution	9
5	A synthetic portfolio case study	10
6	For the chief executive	13
7	For the board	15
8	Evaluating claims with their limits attached	16
9	Objections executives raise	18
10	Where SeniorCRE fits, and where proof remains open	18
A	Board checklist	20
B	Chief executive questions.....	21
C	Glossary	22
D	Sources and boundaries	22

This paper is educational governance analysis. It is not legal, clinical, investment, staffing or regulatory advice.

SUMMARY

Executive summary

An AI agent can compare thirty communities in moments. That does not mean the thirty communities are describing the same operating condition, and it does not mean anyone has authorized what the agent suggests.

Senior housing & care organizations are adding AI to staffing, occupancy, expense and reporting workflows. The analysis is often good. The harder question is institutional: when a consequential recommendation is made, which evidence governed it, who had authority to approve the action, and what proves the action actually happened?

Model confidence never creates organizational authority.

This paper sets out a practical way to govern that question. It makes seven arguments.

1. **Capability and authority are different questions.** A more capable model improves analysis. It does not decide which of two conflicting records is authoritative, and it cannot hold accountability for a resident-care or workforce decision.
2. **Portfolios rarely share meaning by default.** The same word, such as occupancy or agency hours, can describe different conditions at different sites. Mathematically sound comparison can sit on inconsistent inputs.
3. **A recommendation is not permission.** The agent proposes. An accountable person determines, within defined limits and for a defined period. Execution is a separate fact that needs its own evidence.
4. **Governance is a sequence.** Preserve source evidence, surface disagreement, route it to the person with authority, seal the determination, let analytics read only sealed records, and verify that the action occurred.
5. **Chief executives should start with a decision, not a promise.** A first governance register of six questions, applied to one recurring decision, is more useful than an enterprise automation pledge.
6. **Boards should test the boundary.** Oversight needs to distinguish analytical output from authorized action, and should run exercises in which the agent is asked to exceed its permission, loses a permission, or meets corrected evidence.
7. **Every claim needs its limits attached,** including ours. Ask any provider whether a capability is designed, built, validation-exercised, accepted in a deployment, or measured.

Status of SeniorCRE's proposal. SeniorCRE is developing operator-controlled operating infrastructure for senior housing & care. As of October 21, 2026, its Operator-Controlled Operating Record is designed and not yet implemented in any community. This paper claims no operational, clinical, legal or financial benefit.

The paper is written for readers who may never use SeniorCRE. The questions and exercises apply whether an agent comes from an incumbent software vendor, an internal team or a new platform, and they are not settled by a provider's brand or a model's benchmark score.

PART 1

The question capability cannot answer

Better intelligence makes a recommendation more persuasive. It does not make it authorized.

What Microsoft reported, and what it does not tell us

On June 2, 2026, Microsoft published a feature describing its second-generation Majorana qubits. It reports that the new qubits can maintain their quantum state 1,000 times longer than the first generation, with a mean lifetime of 20 seconds and some runs up to a minute. The feature credits agentic AI, delivered through Microsoft Discovery, with helping manage workflows, automate measurements, optimize fabrication and find previously unnoticed flaws. In one example an agent reportedly caught an uncalibrated temperature sensor that was skewing fabrication data.

Three details matter for governance. First, Microsoft describes the AI as offering guidance, with scientists in the loop. Second, the materials stack itself came from research that predates the agents. Third, the feature is Microsoft’s own account built on company statements. It does not report independent review, and its forward-looking claim, a scalable quantum computer by 2029, is a company expectation.

We cite this research as a prompt, not as support for our architecture. It does not mention senior housing & care, and Microsoft has not endorsed or validated SeniorCRE. Our interpretation is narrow: even where agents are capable enough to notice a flawed sensor, the organization still decides what to do with what they find. The more capable the agent, the more the surrounding question is about evidence, permission and accountability.

CAPABILITY AND AUTHORITY ARE DIFFERENT QUESTIONS



Source: Microsoft Source, “Majorana 2, made more reliable with Microsoft Discovery agentic AI,” June 2, 2026. <https://news.microsoft.com/source/features/innovation/majorana-2-microsoft-discovery-agentic-ai/>

Three separate questions

Question	What it asks	What is often missing
Evidence	Which original records support the conclusion, including material contradictions and missing information?	The source period, whether competing records were reconciled, and what was absent.
Authority	Who was permitted to make the determination or authorize the action, for what purpose and within what limits?	A written grant. Permission is inferred from convenience or from the agent's confidence.
Accountability	Can management reconstruct the recommendation, approval, response and later correction without silently filling gaps?	Proof of execution, and a clear record of what was superseded.

A confident recommendation is not an approval. An approval is not evidence that the action occurred. A retained record is not proof that the decision was appropriate. Keeping those four things apart is the discipline this paper describes.

Why senior housing & care feels this first

Decisions here touch residents and employees directly, run across many communities with different needs, and draw on several systems that were never designed to agree. Staffing, occupancy, agency spend and care-plan activity each live in a different platform with its own definitions and cutoffs. When an AI agent sits on top of that landscape, it inherits every disagreement beneath it.

That is why the question is practical rather than philosophical. An agent that reports a coherent portfolio picture may be doing exactly what it was asked. The risk is that the organization treats the result as settled before anyone has decided what the inputs mean.

PART 2

Why the evidence disagrees

The portfolio is described in one language by its dashboards and in thirty by its communities.

One site counts agency hours when an invoice arrives. Another counts them when attendance is verified. One measures occupancy at midnight and another reports the day's average. A portfolio recommendation can be mathematically coherent while its operating inputs remain institutionally inconsistent.

These differences are rarely errors. Each can be legitimate for its local purpose: payroll, billing, clinical staffing or reporting to an owner. The governance problem is not that figures differ. It is that nobody has said which definition and which source governs a given decision, so a disagreement can be resolved silently by whichever system a person happened to open.

SAME WORD, DIFFERENT MEANING ACROSS COMMUNITIES

	One community	Another community
Agency hours	Counted when the invoice arrives	Counted when attendance is verified
Occupancy	Measured at midnight	Reported as the day's average
Staffing exception	Judged against the posted plan	Judged against resident needs that shift



An agent compares all of them in moments.
The arithmetic can be sound while the meaning is not shared.


A single shift shows the pattern

Consider a synthetic shift in one community. The approved staffing plan calls for twelve caregivers. The scheduling platform shows twelve assigned. Timekeeping shows ten clocked in. An agency invoice shows one contracted caregiver. Eleven people were physically present against a plan of twelve, and the dashboard reports full coverage.

EXHIBIT 1

One synthetic shift, four sources

ONE SYNTHETIC SHIFT, FOUR SOURCES



Variance: 11 present against 12 required, one caregiver short. The dashboard says none.

None of the four systems is wrong by its own definition. Scheduling records who was assigned, timekeeping records who clocked in for payroll, the invoice records what was billed, and the plan records what was required. A governed answer needs a rule for which source decides what, a named person who resolves the conflict, and evidence of what happened next.

Common meaning without erasing local context

Portfolio oversight needs shared definitions. It also needs legitimate differences to remain visible. A staffing exception may depend on resident needs, staff qualifications, local requirements and the shift being examined. A common dashboard must not erase those conditions, and a common definition must not be used to override clinical judgment at the bedside.

The practical balance is to standardize the meaning of the figure that crosses the boundary between community and corporate, such as which hours count as direct care, while leaving community leaders responsible for care delivery and immediate adjustments. Corporate leaders gain visibility into recurring exceptions without taking over local decisions.

Quietly selecting the most convenient value is the failure to avoid. If records disagree, identify the dispute, preserve both assertions and route it to someone with authority.

PART 3

A recommendation is not permission

The agent proposes. A person with authority decides. A separate record shows what happened.

A synthetic example

Suppose a portfolio agent recommends moving personnel among three communities to reduce overtime. Its recommendation may depend on outdated acuity information or an unverified assumption about available coverage. The recommendation should not itself authorize the transfer. An accountable manager must review the relevant care and workforce conditions. Any authorization needs a defined scope, an effective period and a retained rationale. Evidence of implementation must remain separate from both the recommendation and the approval.

This is a synthetic example. It is not a staffing instruction or a measured SeniorCRE result, and real decisions require professional judgment and applicable clinical, employment and regulatory review.

What the agent may read, and what it may do

Permission has two halves. The first is read access: what data an agent may use, and for what purpose. The second is action authority: what it may initiate, recommend or merely flag. Both should be explicit grants. If a required grant is absent, permission should not be inferred, and the agent should refuse or escalate rather than proceed on a reasonable-sounding assumption.

- **No silent merging.** No platform, integration or spreadsheet should average, merge or pick a favorable number to clear a conflict between records.
- **No algorithmic judgment on authority.** No model decides which conflicting record is authoritative. That is a determination for a designated person under a written rule.
- **Human accountability stays human.** Software cannot bear clinical or legal accountability. A named role in a documented chain of command does.

Five states a decision can truthfully be in

Most workflow tools collapse the life of a decision into a single status. A more honest record distinguishes five states, and each one has to correspond to something that exists.

EXHIBIT 2

Unresolved, approved, executed, verified and superseded

FIVE STATES A DECISION CAN TRUTHFULLY BE IN



Each label must match a real record.

A completed workflow indicator must not disguise a missing approval, and an approval is not evidence the action occurred. Where evidence is absent, show the gap, its owner and the next review.

State	Means	Evidence it needs
Unresolved	Sources disagree or evidence is missing.	The disagreement itself, preserved, with its owner.
Approved	An authorized person decided, within limits and for a period.	Name, role, scope, effective period and rationale.
Executed	The action is recorded as taken.	A record from the system or person that took it.
Verified	Independent evidence confirms the action occurred.	For example a sign-in log, a timekeeping adjustment or a care plan reassessment.
Superseded	A later determination replaced this one.	A link from the old record to the new one. The old one stays reconstructable.

These labels matter because a completed-workflow indicator can disguise a missing approval, and a retained approval can disguise an action that never happened. A board packet or an executive dashboard that uses these states honestly can show the gaps instead of fabricating certainty.

PART 4

A governance sequence from evidence to execution


The sequence is deliberately human-authorized, and it preserves disagreement instead of overwriting it.

The sequence has six stages. Each has one job, and a figure that skips a stage has no recorded authority behind it. We describe the sequence as a governance model that any organization could adopt. It is also the sequence on which SeniorCRE's proposed Operating Record is designed.

EXHIBIT 3

Six stages from raw evidence to verified execution

SIX STAGES FROM RAW EVIDENCE TO VERIFIED EXECUTION



Stage 6 feeds back to Stage 3: closed-loop proof, not a log entry.

Stage	What happens	Control
1 Evidence	Source assertions are held unaltered, with their original identities, definitions, timestamps and lineage.	Chain of custody. Raw data isolated from analytics.
2 Disagreement	Rules flag material conflicts between sources, and between sources and the approved plan.	Diagnosis only. No correction and no smoothing.
3 Authority	Each flag routes to the person with decision rights for that kind of conflict.	Software cannot resolve its own flag.
4 Governing record	The human determination is sealed with the conflicting sources, definitions, decision-maker, rationale and effective period.	Later corrections are added, never overwritten.
5 Intelligence	Dashboards, analytics and agents read sealed determinations rather than raw feeds.	No unverified streams in executive reporting.
6 Execution	Evidence shows that the decision was carried out.	Closes the loop and returns an exception to stage 3.

Coherence before intelligence

The order is the point. Operational facts should be governed, reconciled and authorized by accountable people before they are consumed by analytics, AI engines or predictive tools. Deploying advanced analysis on conflicting source data produces flawed intelligence and unstable dashboards, and it makes compliance assertions harder to defend. We call this principle coherence before intelligence.

A record that shows a gap and no response

A sealed record of a known problem is only half a control. If a record shows leadership knew of a shortfall and nothing shows a response, the record documents awareness without action. The remedy is not to keep fewer records. It is to pair the determination with evidence of execution, and to record a missing piece as a gap with an accountable owner rather than to invent a conclusion.

Where evidence is missing, record the gap and its accountable owner. Do not estimate, backfill or let software average the missing number.

This is a risk-control observation, not legal advice. How any record would be treated in a proceeding depends on the facts and the jurisdiction, and counsel should review any policy built from it.

PART 5

A synthetic portfolio case study

The sequence is easier to judge on a concrete week. This case is invented. It shows method, not product capability, and it is not a staffing instruction.


The setting

A regional operator runs thirty senior housing & care communities. A portfolio agent reads scheduling, timekeeping, agency invoices and the clinical record every morning. It is asked a recurring question: where is overtime rising, and can people be moved between communities to reduce it? Three communities matter this week.

EXHIBIT 7

What the agent saw and what people decided

SYNTHETIC CASE: WHAT THE AGENT SAW, AND WHAT PEOPLE DECIDED



Illustrative and synthetic. Not a staffing instruction or a measured result.

Step 1: the agent's recommendation

Community A has two caregivers short on nights and rising overtime. Community C's schedule shows two more people assigned than its plan requires. The agent recommends moving four caregivers from C to A and B for two weeks. The arithmetic is sound and the recommendation is persuasive. Under the scheduling system's definition, C does have a surplus.

Step 2: the disagreement surfaces

The governed process does not act on the recommendation. It lays the sources for each community side by side and flags where they diverge.

Community	Scheduling	Timekeeping and agency	Clinical record	Flag
A	Two short against plan	Overtime rising; no agency cover	Stable acuity	Shortfall confirmed by two sources
B	Two above plan	Clock-ins match the schedule	Stable acuity	No conflict; spare capacity verified
C	Two above plan	Three call-outs on nights; two agency shifts	Two recent admissions raise acuity	Conflict: schedule overstates present coverage

In C, the scheduling surplus is an artifact. Three people did not work, two agency workers filled in, and acuity went up. Moving four caregivers out would have cut coverage where residents need it. The recommendation had read the one source that supported it.

Step 3: routing to authority

The conflict in C goes to the community's Director of Nursing, who holds authority over clinical coverage, and the cross-community transfer goes to the chief operating officer. Neither the agent nor a workflow rule resolves it. Each person sees the original records, the definitions in force and the proposed action.

Step 4: the determination

The people decide three things. They decline the move from C. They approve moving two caregivers from B to A for one week, because B's spare capacity is verified by two sources. They refuse the remaining part of the recommendation. Each decision is recorded with the evidence, the approver's name and role, the scope, the end date and the reason.

Step 5: execution and verification

Two things must be shown. A's sign-in log and timekeeping record show the two caregivers arrived and worked the shifts. A's clinical record shows the care plans were reassessed against the new coverage. If either is missing at the end of the week, the record shows a gap and its owner rather than a completed transfer.

Step 6: expiry and supersession

At the end of the week the approval expires. New evidence, such as a further admission in C or a recovery in A's call-outs, supersedes it. The old determination stays reconstructable and remains linked to the new one. Dashboards that report overtime read the sealed determinations, so they show which exceptions were approved, which were refused and which were verified.

What the board sees

Decision category	State	Evidence	Owner and next review
Cross-community transfer from C	Declined	Records show C's coverage was overstated	Chief operating officer; weekly review
Temporary transfer, B to A	Verified	Sign-in log; timekeeping; care plan reassessment	Chief operating officer; expires at end of week
Overtime trend in A	Unresolved	Persistent shortfall remains	Regional director; next board exception review

What the case does and does not show

- It shows how a persuasive recommendation can rest on one source, and how a governed process catches that without blocking the agent's useful analysis.
- It shows that refusal and escalation are results, not failures. Two of the three outcomes were refusals or partial approvals.
- It does not show that any system prevents this error today, that SeniorCRE has produced such a result, or that any cost, safety or liability benefit follows.
- Real decisions require professional judgment and applicable clinical, employment and regulatory review.

PART 6

For the chief executive

The responsibility is not to approve every transaction. It is to make the delegation structure clear and the decision reconstructable.

Begin with a decision, not an enterprise promise

A promise to automate across the portfolio invites broad claims no one can check. A decision is concrete. Choose one recurring question, such as whether a staffing exception needs escalation, whether an occupancy figure may support a management report, or whether an expense variance warrants investigation.

Then define the field, the period, the community, the intended use and the accountable owner. Preserve the original source assertions. If the records disagree, identify the dispute instead of quietly selecting the most convenient value.

The first governance register

For each selected decision, document six things.

EXHIBIT 4

Six questions for each decision

THE CEO'S FIRST GOVERNANCE REGISTER: SIX QUESTIONS PER DECISION

1 Read What may the agent read, and for what purpose?	2 Govern Which source and definition govern each input?	3 Escalate What disagreement requires refusal or escalation?
4 Approve Who may approve the action, within what limits?	5 Confirm What confirms execution, and who checks it?	6 Correct How are corrections and superseding decisions retained?

Question	Why it matters	A good answer
What may the agent read, and why?	Read access is a grant, not a default.	A written purpose and scope for each data source.
Which source and definition govern each input?	Without a rule, conflicts are settled by whichever system is opened.	A named source authority for each input in each decision context.
What disagreement requires refusal or escalation?	An agent should stop rather than guess.	Thresholds and triggers, and the person to whom it escalates.
Who may approve the action, within what limits?	Approval has scope and duration.	A role, a limit, an effective period and a retained rationale.
What confirms execution, and who checks it?	Approval is not proof the action occurred.	An artifact separate from the recommendation and the approval.
How are corrections and superseding decisions retained?	The old decision must remain reconstructable.	A link from the old record to the new one.

If a required grant is absent, permission should not be inferred. The chief executive need not personally approve every transaction. The responsibility is to establish a clear delegation structure, require consequential exceptions to reach the right owner, and ensure that oversight can reconstruct a decision without inventing missing evidence.

A delegation structure for consequential exceptions

Delegation can be simple. What matters is that every decision category has a named owner and a limit, and that exceptions above the limit rise to someone with authority.

Decision type	Community leader	Operations or clinical executive	Chief executive or board
Immediate shift adjustment	Decides and records	Reviews recurring patterns	Sees exceptions in the packet
Clinical staffing change	Director of Nursing decides within policy	Chief clinical officer sets policy and reviews	Informed of material exceptions
Cross-community staff transfer	Provides evidence	Chief operating officer decides within limits	Informed when limits are exceeded
Occupancy figure used in owner reporting	Supplies source figure	Chief financial officer approves the governing definition	Reviews material changes to definitions
Agency spend against worked hours	Confirms attendance	Chief financial officer validates	Reviews variances above limit

A bounded evaluation, not a portfolio promise

An initial evaluation can compare the same decision under present controls and under a proposed governed method. Keep the source assertions, disputed values, approvals, refusals and completion evidence available for review.

- Does the process preserve disagreement rather than resolving it silently?
- Does it enforce the actual authority boundary, including refusal when a grant is absent?
- Does it distinguish an authorized action from one that occurred?

- Are failed tests recorded as findings, not removed from the report?

Do not infer financial improvement, resident-safety benefit or reduced liability merely because the record is easier to inspect. Those are separate outcome hypotheses and need appropriate evidence of their own.

PART 7

For the board

A board can receive an elegant AI-assisted report and still lack what it needs to oversee the decision behind it.

A report may contain accurate calculations while leaving questions open: which source period was used, whether competing records were reconciled, whether an agent had permission to initiate an action, and whether management verified completion. Board oversight should distinguish analytical output from authorized institutional action.

What a useful board packet contains

Boards generally need material exceptions and control evidence, not unrestricted access to resident or employee records. Management should provide appropriately limited information with the context needed for oversight.

EXHIBIT 5

What a board packet can show

WHAT A BOARD PACKET CAN SHOW WITHOUT EXPOSING RESIDENT RECORDS

Decision category	Authority boundary applied
Responsible executive	Response status
Source period	Evidence supporting closure
Unresolved disagreement	Gap, owner and next review

Limited, relevant information for oversight.

Detailed records follow privacy, professional and legal requirements.

The packet should distinguish unresolved, approved, executed, verified and superseded states, and those labels must reflect actual records. Where evidence is absent, the packet should show the gap, the owner and the next review rather than fabricate certainty. Access to detailed supporting records should follow privacy, professional and legal requirements. Board visibility is not a justification for broad data sharing.

Test the boundary, not just the successful answer

A demonstration that an agent can produce a correct result proves little about governance. Three synthetic exercises test the boundary instead.

Exercise	Setup	What to review
Exceed the grant	Present conflicting staffing records and an instruction outside the agent's authority.	Does the system refuse or escalate, preserve the disagreement, and avoid claiming the requested action occurred?
Revoke a permission	Remove a permission that was previously available.	Does the agent stop, and is the earlier authorization still reconstructable?
Introduce corrected evidence	Add a corrected source record after an approved determination.	Can users distinguish the old determination from the superseding one, and is both retained?

These are proposed evaluation methods, not claims of tested SeniorCRE production controls. Passing a demonstration does not establish deployment readiness, compliance or safety across an operating portfolio.

Human review must be meaningful

A human approval step can become ceremonial if reviewers cannot see the evidence, lack relevant expertise or face pressure to accept an automated conclusion. Ask management how reviewers encounter conflicting records, when they must escalate, whether they can refuse a recommendation, and how overrides are retained. Require the organization to distinguish what the agent proposed from what an authorized person determined.

Equally, human approval does not cure unsupported evidence or relieve professionals of their responsibilities. The board's role is not to select every operating value. It is to ask whether management has established accountable authority, proportionate controls and evidence that those controls behave as represented.

PART 8


Evaluating claims with their limits attached

A proposed control is not a demonstrated capability merely because a page, a diagram or an interface exists.

Ask every provider, including SeniorCRE, to distinguish five levels of evidence. For each material claim, request the evidence, the scope, the date, the exclusions and the acceptance gate.

EXHIBIT 6
The evidence scale

ASK EVERY PROVIDER WHERE A CLAIM SITS ON THIS SCALE



Operating Record, October 21, 2026: level 1, designed.

Not yet implemented in any community. Higher levels need evidence, scope, date, exclusions and an acceptance gate.

Level	What it can support	What it cannot
Designed	A stated architecture and a way to evaluate it.	Any claim about operation, safety or results.
Built	That interfaces or components exist.	That they work end to end under real conditions.
Validation-exercised	Behavior on synthetic or controlled cases.	Deployment readiness, compliance or safety across a portfolio.
Operator acceptance	That an operator accepted a scoped deployment against agreed criteria.	Outcomes, savings or generalization to other operators.
Measured outcomes	Results against a stated baseline, with scope and dates.	Results outside that scope.

The same scale applies to AI features offered by incumbent vendors. A demonstration of an agent’s ability to draft a staffing recommendation sits at a different level from evidence that it respected an authority boundary in production for a quarter.

Ask for the date, scope and exclusions behind every capability claim, and ask where it sits on the scale. A provider that welcomes the question is easier to trust than one that treats it as an obstacle.

PART 9

Objections executives raise

These are fair questions. The answers are short because the position is narrow.

Objection	Response
We already have dashboards.	Dashboards show a number. Governance shows which source and rule produced it, who approved exceptions and what happened next. They complement each other when the dashboard reads governed determinations.
Human approval will slow us down.	Only exceptions need a person. Routine agreement between sources can pass on a written rule. The cost of approval falls on the disagreements, which are the cases that already take hours to settle.
Our vendors say their agents are safe.	Ask where each capability sits on the evidence scale and what the vendor will show for the authority boundary. Brand and benchmark scores do not answer that.
This sounds like more paperwork.	The record is a by-product of the decision, not a second task. If capturing it is burdensome, that is a design problem to fix, not a reason to leave the decision unrecorded.
Approvals could become rubber stamps.	Agreed. That is why reviewers need to see the evidence, be able to refuse, and have overrides retained. Test it with the boundary exercises.
Won't records create liability?	A record that shows a gap with no response can. That is why the paper pairs the determination with proof of execution and records gaps honestly. Counsel should shape the policy.
We cannot standardize across communities.	Standardize only the figure that crosses the boundary between community and corporate. Local clinical judgment stays local.
Why not wait for the technology to mature?	The agents are already in use. The questions of authority, evidence and verification apply now, and the exercises cost little to run.

PART 10

Where SeniorCRE fits, and where proof remains open

SeniorCRE is developing operator-controlled operating infrastructure for senior housing & care. Its thesis is operator-held authority over evidence, determinations and downstream use, not vendor ownership of an organization's decisions. The operator-held control plane governs what agents may read, what they may act on and what they must escalate.

Two intended configurations

The operator has two intended deployment choices. It can retain its existing systems and govern what they produce, which we call Road 1. Or it can choose SeniorCRE itself as the system running a function, which we call Road 2. The operator chooses the road. Replacement is not required to obtain governance, and it is not prohibited. Neither road is a claim of current production deployment.

Current status

Operating Record	Designed. Not yet implemented in any community as of October 21, 2026.
Native clinical configuration	Future and unavailable for production.
Medication administration and eMAR	Roadmap.
Required before reliance	Scoped implementation, function-specific safety review, end-to-end validation and operator acceptance.
Not claimed	Operational, clinical, legal or financial benefit; compliance; any owner, lender or operator endorsement; Microsoft endorsement or validation.

These limits are not caveats added to a finished story. They are part of the position. A company that asks owners and boards to demand evidence should apply the same standard to itself.

How to begin, with or without SeniorCRE

- Choose one recurring decision and write the six-question register for it.
- Reconstruct one recent consequential recommendation: the source records, the unresolved disagreement, the decision owner, the approval and the evidence that the action occurred. Record what you could not produce.
- Run the three boundary exercises on whatever agents you already use.
- Ask every provider to place its claims on the evidence scale.

Executives who would like to discuss the approach can request an Executive Briefing at <https://seniorcre.com/executive-briefing>. The conversation starts from your existing controls, not from an assumed product fit.

APPENDIX A

Board checklist

One page for a board meeting. Tick what management has shown, note what it has not, and name an owner for each gap. Where evidence is missing, record the gap. Do not infer.

Ask management	Shown	Gap and owner
Which consequential decisions now involve an AI agent or automated recommendation?	Yes / No	
For each, who is the accountable executive?	Yes / No	
What may each agent read, and where is that grant written?	Yes / No	
Which source and definition govern each input, and who wrote the rule?	Yes / No	
What makes an agent refuse or escalate, and to whom?	Yes / No	
Who may approve each action, within what limits and for how long?	Yes / No	
Does the packet separate unresolved, approved, executed, verified and superseded?	Yes / No	
Where a status says complete, is there evidence of approval and execution?	Yes / No	
Can management reconstruct one recent decision end to end today?	Yes / No	
Have boundary exercises been run: exceed the grant, revoke a permission, correct the evidence?	Yes / No	
Can reviewers see conflicting records, refuse a recommendation and have overrides retained?	Yes / No	
Where does each provider claim sit on the evidence scale, with date, scope and exclusions?	Yes / No	
What limits on supporting records do privacy, professional and legal duties set?	Yes / No	
What remains an assumption?	Yes / No	

Date reviewed: _____ Reviewed by: _____ Next review: _____

Chief executive questions

For the chief executive

1. Which recurring portfolio decision could we reconstruct today, including the evidence that the authorized response actually happened?
2. For that decision, which source and definition govern each input, and who wrote that rule?
3. What may our agents read, and where is that grant written?
4. What would make an agent refuse or escalate, and to whom?
5. Who may approve the action, within what limits and for how long?
6. What separates an approved action from an executed one in our records?
7. How would a correction or superseding decision be retained?
8. Which conclusions in our dashboards depend on raw feeds rather than sealed determinations?

For the board

1. If a consequential AI-assisted decision were challenged later, what could management produce today, and what would remain an assumption?
2. Does the packet distinguish unresolved, approved, executed, verified and superseded states?
3. Where a status says complete, is there evidence of an approval and of execution?
4. Who in management is accountable for each decision category?
5. How do reviewers see conflicting records, and can they refuse a recommendation?
6. Has management run boundary exercises, and what failed?
7. What privacy and legal limits govern the supporting records the board may see?
8. For each provider claim, where does it sit on the evidence scale?

APPENDIX C

Glossary

Agent	Software that can read, recommend or initiate actions with some autonomy. It has no organizational authority of its own.
Authority	A written grant to make a determination or authorize an action, held by a named role, with limits and a duration.
Coherence before intelligence	The principle that operating facts are governed and authorized before analytics or AI consume them.
Determination	A human decision that resolves a disagreement or authorizes an action under a stated rule.
Governing record	A sealed record of a determination with its sources, definitions, decision-maker, rationale and effective period.
Lineage	The traceable path from a reported figure back to the source records and approvals behind it.
Road 1	The operator retains existing systems and governs what they produce.
Road 2	The operator chooses SeniorCRE as the system running a selected function.
Source authority	The rule naming which source governs a given input in a given decision context.
Superseded	Replaced by a later determination that remains linked to the original.

APPENDIX D

Sources and boundaries

Microsoft Source feature, June 2, 2026	https://news.microsoft.com/source/features/innovation/majorana-2-microsoft-discovery-agentic-ai/
Coherence Before Intelligence (SeniorCRE)	https://seniorcre.com/book
The Institutional Operator (SeniorCRE)	https://seniorcre.com/the-institutional-operator
The Institutional Owner (SeniorCRE)	https://seniorcre.com/the-institutional-owner
Request an Executive Briefing	https://seniorcre.com/executive-briefing

Public discussion paper. Educational analysis, not legal, clinical, investment, staffing or regulatory advice. Synthetic scenarios do not demonstrate product capability or outcomes. Microsoft does not endorse or validate SeniorCRE. As of October 21, 2026, the Operator-Controlled Operating Record is designed and not yet implemented in any community.