

**AI IN LONG-TERM CARE:
PROMISE, RISK & READINESS**

WHITEPAPER

What leaders need to know now, before adoption
outpaces oversight.

A strategic briefing for executives,
compliance leaders, and care
innovators navigating AI's rise in elder
care environments.

This guide helps leaders balance
innovation with resident safety,
workforce impact, and regulatory
realities.

Executive Summary

Artificial intelligence (AI) is no longer theoretical in long-term care, it's arriving in real workflows, right now. From predictive staffing tools to AI powered resident engagement platforms and automated clinical documentation, technology vendors are actively marketing solutions that promise to streamline operations, reduce burnout, and improve quality of care.

Some of these tools offer real promise. Others introduce new risks that are poorly understood, and in environments as complex and sensitive as long-term care, unknown risk is unacceptable.

This briefing is designed for executives, clinical leaders, and compliance officers who recognize the urgency to innovate but refuse to do so blindly.

You'll find:

- 3 use case snapshots of how AI is being used today in LTC environments
- Key risk domains to evaluate before adoption, from data governance to clinical oversight
- Guidance for pilot testing, monitoring outcomes, and preparing for future regulation
- A readiness framework to help your leadership team determine where you are now and what steps to take next



AI can amplify quality. It can also amplify liability

LTC leaders face a dual responsibility: staying ahead of innovation while safeguarding clinical integrity, regulatory compliance, and resident safety. With minimal regulatory guidance available today and rising pressure to adopt, the decisions made in 2025 could shape both outcomes and exposures for years to come.

This resource equips leaders with a practical lens for evaluating AI through a long-term care specific lens, not just from a tech standpoint, but through the operational, clinical, and reputational risks it creates or mitigates.

Whether your organization is testing new tools or simply monitoring the space, this guide is designed to help you make smarter, safer, and more strategic decisions about how AI fits into your care environment.

AI Use Case Snapshots in Long-Term Care

Innovation is happening fast. Here's how AI is already being deployed in LTC environments and what leaders should consider before adoption.

Staffing Optimization Tools

AI powered platforms are analyzing shift patterns, resident acuity, historical call outs, and staffing gaps to recommend real time scheduling decisions.

Promise

- Reduces burnout by identifying proactive coverage solutions
- Improves cost control through smarter resource allocation
- Adjusts staffing levels in response to dynamic resident needs

Risk

- Over reliance on predictive models that may not reflect real time nuance (e.g., behavioral needs)
- Lack of transparency in algorithms, what data is being used, and is it biased?
- May conflict with collective bargaining agreements or create staffing inequities

Strategic Questions

- Does the AI system allow clinical overrides based on human judgment?
- How does it account for non-quantifiable factors like resident rapport or emotional labor?
- Who is responsible if an AI driven staffing decision results in a lapse in care?

Leadership Reflection Points

"If a staffing tool consistently overrides our instinct or experience, are we leading, or following a black box?"

"If this tool replaced a core task, what leadership visibility would we lose in the process?"

Use these prompts to open a discussion with clinical and operations leaders:

- Where do we draw the line between augmentation and abdication?
- Who remains accountable when algorithmic decisions go wrong?
- Are we losing valuable data in the name of efficiency?

Resident Engagement Platforms

AI chatbots and cognitive stimulation tools support social interaction, mental stimulation, and loneliness mitigation.

Promise

- Extends engagement capabilities without increasing staff workload
- Provides data on engagement patterns to inform care planning
- Can support residents with dementia using tailored prompts or reminders

Risk

- May cause resident confusion or distress if not properly introduced or monitored
- Raises ethical concerns if residents believe they are interacting with a person
- Privacy concerns around data collection during interactions

Strategic Questions

- How will families and staff be educated on what the tool is and isn't?
- Is usage opt-in and tracked for consent?
- What steps are in place to intervene if a resident reacts negatively to the interaction?

What Could Go Wrong?

An AI chatbot is introduced in the memory care unit to reduce resident isolation. Staff notice early signs of success, one resident, Margaret, seems calmer and talks to it frequently. But over time, she begins referring to the chatbot as her son, asking staff to “call him” during meals and becoming agitated when the device updates or disconnects. Staff are unsure how to respond. Eventually, the family expresses concern that the tool was used without their understanding, and a regulatory complaint is filed.

► This isn't just a tech hiccup, it's a breakdown in ethical oversight, staff training, and family communication.

Yet just because something could go wrong doesn't mean it should be avoided entirely. Engagement tools hold real promise when used with intention, guardrails, and proactive support.

The key question isn't “Is this risky?”, it's “Have we planned for and mitigated the risk responsibly?”

Clinical Documentation & Note-Taking Tools

One of the most promising and widely adopted tools for saving time across the healthcare spectrum, speech to text and ambient AI are used to automate care notes, shift reports, and incident logs.

Promise

- Reduces staff time spent on documentation
- Minimizes delayed or missed entries
- Creates structured, standardized records for clinical review and billing

Risk

- Inaccurate transcription of nuanced or sensitive content
- Risk of misclassification or misinterpretation of clinical terms
- Legal exposure if documentation errors are tied to adverse outcomes

Strategic Questions

- Are clinicians reviewing and approving all AI generated notes?
- Is there a clear audit trail to distinguish human versus machine inputs?
- How would you handle a legal challenge if an AI generated record is called into question?

Mini Checklist: AI Generated Documentation Review

As long-term care organizations explore AI tools to streamline clinical documentation, it's tempting to focus on efficiency gains and cost savings. But before adopting these technologies, leadership must ask: *Are we truly ready to integrate this into our care environment responsibly?*

Automating documentation isn't just a workflow shift, it alters how information is captured, interpreted, and relied upon during audits, reviews, and critical events. A well designed pilot program should go beyond functionality to assess trust, usability, and legal defensibility. That means ensuring strong human oversight, establishing clear definitions of success, and preparing staff not only to use the tool, but to understand its limitations.

The checklist below outlines key considerations for planning and evaluating a safe, scalable pilot.

Use this quick check during pilot planning or vendor demos:

- ☐ Are AI generated notes clearly flagged in the record?
- ☐ Is there a mandatory human review step before submission?
- ☐ Does the system store voice/audio securely and in compliance with HIPAA?
- ☐ Can staff easily revise, reject, or comment on AI entries?
- ☐ Is there version history for medico-legal defense or audits?

Vendor Vetting Questions: What to Ask Before You Adopt

When AI vendors present tools to long-term care leaders, the focus is often on features and potential. But effective adoption begins with the right questions, asked early and with strategic clarity. Your goal isn't to become a tech expert, it's to protect your residents, staff, operations and reputation by ensuring each tool aligns with your mission, risk tolerance, and regulatory obligations.

Below are two sets of questions, one for assessing external AI vendors, and one for evaluating your own internal readiness to implement the tools they offer.

Questions to Ask AI Vendors

- How was the algorithm trained and does it include data from long-term care settings?
- What level of human oversight is expected or required?
- Can the system's recommendations or outputs be audited or explained to regulators?
- What happens to resident or staff data once it's processed? Who owns it?
- How is HIPAA compliance maintained across all data collection, processing, and storage functions?
- What is your process for addressing errors, outages, or harm caused by AI generated content?
- Can this tool be configured or customized to match our policies or values?
- How do you support post implementation monitoring or quality control?

Questions to Ask Internally Before Implementing AI

- What specific problem are we trying to solve and is AI the best solution?
- Who is responsible for reviewing outputs, monitoring risks, and intervening when needed?
- Have we discussed how this tool might change workflows, roles, or expectations?
- Are we prepared to document our due diligence and oversight if surveyed or audited?
- Do our staff feel comfortable flagging concerns or pushing back on AI-driven decisions?
- Have we set clear success criteria, not just for performance, but also for safety, ethics, and communication?

Evaluation Isn't a One Time Event

AI tools don't sit still. Algorithms update. Vendor business models shift. Staff roles evolve. What seemed like a good fit during pilot testing may drift out of alignment over time, silently introducing risk into your operations. That's why AI adoption must be treated as a lifecycle, not a project.

Just as you regularly review care protocols, budget assumptions, or vendor contracts, your organization should build in recurring checkpoints for any AI system in use. This doesn't need to be overly technical or burdensome, but it does need to be intentional.

Bottom line: Responsible AI adoption isn't about locking in the right tool, it's about creating a culture and process that can adapt as the technology and your organization evolve.

Pilot Planning Framework: Test Small, Learn Fast, Scale Wisely

AI in long-term care isn't an all or nothing decision. The safest and most effective approach is a structured pilot, designed to surface questions early, validate value, and build confidence before full scale adoption.

Many pilot programs fail, not because the technology doesn't work, but because the planning wasn't deep enough. Success doesn't just mean "it ran." It means your staff trusted it, your data was protected, and your leadership gained clarity, not complexity.

Below is a simple framework to shape your next AI pilot:

1. Define the Problem, Not Just the Tool

Start by identifying the pain point. Is the goal to reduce documentation time? Improve staff scheduling accuracy? Increase resident engagement consistency?

Be specific, and be willing to consider non-AI solutions if they're simpler and safer.

2. Set Measurable Success Criteria

Before you deploy, define what success actually looks like. Go beyond technical performance.

Examples:

- Time saved per nurse per shift
- Accuracy compared to human generated documentation
- Staff confidence or satisfaction with the tool
- No unresolved privacy or data ownership concerns

3. Create a Feedback Loop for Users

Staff should never feel like passengers in an AI rollout. Build in weekly or bi-weekly check-ins where users can voice concerns, flag confusion, or propose adjustments. Early stage buy in is earned, not assumed.

4. Assign Oversight Roles

Who will review outputs, spot errors, escalate concerns, and track unintended consequences? Assign specific individuals, clinical, operational, and technical, who are accountable for monitoring the pilot and making go/no-go decisions.

5. Plan for the End (Not Just the Beginning)

Many pilots drag on with no conclusion, or scale prematurely. Before you launch, determine when and how you'll evaluate results. Will success be decided by leadership, staff feedback, or both? How will findings be shared across departments?

A good pilot doesn't just generate data, it builds decision-making clarity.

Strategic Reflections for Leadership

AI may be driven by technology, but its outcomes are shaped by leadership.

Before adoption, thoughtful leaders pause to ask not just “can we?”, but “should we?” and “how will we know?” The following prompts are meant to guide internal discussion and strategic planning, especially when the pressure to innovate feels urgent.

Organizational Alignment

- Is this tool advancing our mission, or reacting to a trend?
- Do we have a shared understanding of what success looks like?
- Have we mapped where this tool fits into current workflows and staffing models?

Many AI projects fail not because the technology underperforms, but because no one defined how it actually supports the organization's goals, or who owns the outcome.

Communication and Transparency

- Have we openly communicated why this tool is being piloted and what staff can expect?
- Do we have a plan for addressing staff questions, concerns, or resistance?
- How will we explain this tool to residents, families, or regulators if needed?

Uncertainty breeds resistance. Clear, proactive communication builds the trust needed to experiment safely.

Governance and Accountability

- Do we have a process for reviewing unintended consequences or ethical concerns?
- Is there a designated owner for monitoring, evaluation, and reporting?
- What would make us pull the plug, even if the tool performs technically well?

AI governance isn't about creating red tape. It's about protecting your values, and your residents, from risk hidden behind innovation.

Even the most promising technology can become a liability without clarity, communication, and care.

True leadership in this moment isn't about rushing to adopt AI, it's about shaping a thoughtful path forward that aligns with your mission, protects your people, and positions your organization to adapt, not react. The goal isn't to be first. It's to be ready.

Closing the Gap Between Innovation and Risk

AI in long-term care is no longer a question of if, but how well.

Done right, it can alleviate staffing burdens, streamline operations, and support better outcomes. Done poorly, it can introduce hidden risks that degrade trust, documentation, and care quality. The difference lies not in the tools themselves, but in how leaders prepare their organizations to use them.

This guide wasn't meant to make you an AI expert. It was built to help you ask sharper questions, plan smarter pilots, and lead with foresight, not just enthusiasm. If you're already having internal conversations about AI, use this as a structure to strengthen them. If you're just getting started, use this to get aligned.

But don't do it alone.

Our healthcare risk and insurance team is actively tracking AI trends, regulations, and real world adoption across long-term care. We help organizations design proactive strategies that balance innovation with governance and protect the people who depend on you most.

Let's Keep the Conversation Going

We're building a living library of tools, case studies, and readiness frameworks for AI in care environments.

If you'd like access, or want a partner to help think through your organization's next steps, reach out directly.

There's no one size fits all answer to AI in long-term care. But with the right guidance, there is a path forward.

Let's Connect

Drew Colwell

Healthcare Risk Advisor | INSURICA

Drew.Colwell@INSURICA.com

406-333-6309

[LinkedIn.com/in/drewcolwell](https://www.linkedin.com/in/drewcolwell)

