HINSHAW

# Deploying AI Companions in Elder Care: A Privacy Compliance Playbook

## Treat AI Companions as High-Risk Programs

**Privacy, Cyber & AI Decoded Alert | 10+ min read**

May 21, 2026

By: Jason J. Oliveri

Emotionally intelligent AI "companions" are increasingly deployed in senior living and home care settings to address loneliness and improve engagement. While these tools may offer meaningful benefits, they routinely process conversational data, behavioral patterns, and emotional inferences, exposing organizations to heightened privacy, consumer protection, and elder-care risks.

For in-house counsel and compliance leaders, the challenge is not whether to use these tools, but how to deploy them with clear limits, informed consent, and durable governance, particularly in environments where providers have a heightened duty of care to residents. Emotionally responsive AI systems are being marketed as "companions" or "engagement platforms" for older adults. They may offer genuine benefits in addressing loneliness and supporting engagement, but they also raise foreseeable risks around privacy, dignity, consent, manipulation, and security that cannot be addressed through standard IT procurement alone.

The emerging consensus in ethics and policy is not that emotionally intelligent AI should be rejected outright, but that companion use in elder care requires clear limits and strengthened oversight, particularly where vulnerable populations are involved. For in-house counsel and compliance leaders, the task is not to say "no," but to design a framework that allows the organization to safely say "yes" to tightly scoped deployments, with defined controls and accountability.

These tools are not medical treatments, and they should not be positioned or evaluated as such, even though they are sensitive to health-related and biometric data.

## What "Emotionally Intelligent" AI Means in Plain Language

For compliance purposes, these **emotionally intelligent AI systems are best understood as companionship and support tools**, rather than clinical tools. They engage residents in ongoing conversation, adapt to preferences,

and often attempt to infer mood or social state over time.[1]

They should not be used to diagnose conditions, make treatment decisions, or independently assess health status.[2] At most, they may identify concerning interaction patterns, such as apparent withdrawal or repeated expressions of distress, and route those signals to human staff according to predefined escalation rules.[3] That distinction is critical for both regulatory purposes (*e.g.*, HIPAA, state health privacy rules) and ethical design.

# Where are Emotionally Intelligent AI Companions Used?

**Operationally, emotionally intelligent companions tend to appear in three settings:**

1. in senior living communities as part of resident engagement programming;
2. in non-clinical home-care, where agencies place devices in a client's home to supplement human visits; and
3. as direct-to-consumer devices or apps that residents and families bring into facilities.[4]

**Each scenario requires a different allocation of roles:** controller vs. processor, covered entity vs. non-covered entity, and who is responsible for notices, consents, and complaint handling.

# Understanding the Data Footprint

A compliance-ready program starts with a data map. While vendors may describe these tools as "just conversational," emotionally intelligent systems typically collect and generate multiple layers of sensitive information.[5]

## Profile Data

- At the most basic level, they process direct identifiers and profile data such as names, ages, language preferences, family contacts, and stated interests.
- In a senior living or home care context, this information is inherently sensitive because it is tied to a person's residence, routines, and relationships within a protected environment.[6]

## Behavioral and Usage Data

- Beyond that, the systems generate behavioral and usage data: when residents interact, how long conversations last, which features are used, and whether engagement suggestions are accepted or declined.
- In elder care, these patterns can reveal loneliness, depression, cognitive changes, or family conflict, even if no one explicitly labels them as such.[7]

## Inference

- The most sensitive layer is inference. To function as "emotionally intelligent," these tools often infer mood, social withdrawal, anxiety, or changes in engagement over time.[8]
- These inferred insights can shape how residents are treated by staff, how families perceive them, and how the system itself interacts going forward. Inferences are powerful precisely because they may be invisible to residents, hard to explain, and not easily contestable.[9]
- From a compliance standpoint, this means companionship tools should be classified alongside other high-risk processing activities. They involve deeper insight into residents' inner lives and vulnerabilities and warrant elevated safeguards, DPIAs/PIAs, and board-level visibility where applicable.[10]

# Health Privacy Concepts at the Edges

Whether health privacy laws apply depends on how the tool is used and by whom. The same technology can be either non-clinical or health adjacent, depending on deployment choices.[11]

If conversational data feeds into clinical records, informs treatment decisions, or is used by regulated providers such as home health agencies or skilled nursing facilities, traditional health privacy and security concepts may apply.[12]

Staff who begin relying on chatbot summaries, copying notes into health records, or configuring direct integrations with electronic health record systems can inadvertently convert "engagement data" into protected health information.[13] Many senior living operators intentionally position AI companions as non-clinical engagement tools.

In those cases, HIPAA may not apply directly because the operator is not a covered entity, and the vendor may not be a business associate.[14] Even so, risk can arise if staff blur the line between "engagement" and "care," or when marketing and disclosures imply health-related functions.[15]

# Compliance Checklist for Deploying AI Companions

**A compliance-oriented deployment should:**

1. Document whether the tool is within or outside clinical workflows;
2. Prohibit casual copying of conversational content into health records; and
3. Address, in contracts and internal documentation, whether any subset of data triggers business associate obligations.

For most non-clinical companionship deployments, the more realistic exposure comes from consumer protection and state privacy law. Commentators and advocates are increasingly focused on transparency, fairness, and protection of older adults from deceptive or exploitative practices in technology-mediated services. You can bet the same for regulators.

# Consumer Protection, Elder Fraud, and Manipulation

Independent of HIPAA, AI companions raise unfair and deceptive practices and elder protection concerns. Older adults are disproportionately targeted by fraud, and emerging tools like AI-enabled voice cloning and persuasive chatbots can exacerbate those risks.[16] These risks typically originate outside the care environment but can intersect with AI-enabled tools if not properly controlled.

Because emotionally intelligent companions can be experienced as human, residents may not fully appreciate that they are interacting with software, that conversations may be recorded, or that transcripts may be reviewed.[17] This is particularly acute in memory care and other settings involving cognitive impairment. In that environment, incomplete or confusing disclosures can look deceptive, especially if the system is positioned as a "friend" or "care partner."[18]

Anthropomorphic systems also invite emotional attachment. In vulnerable populations, that attachment can become over-reliance, particularly if the system presents itself as a confidant or authority figure. Commentators have raised concerns that some AI companions in consumer settings can steer vulnerable users toward harmful real-world behavior when boundaries are unclear, even though documented cases in elder care facilities remain limited.[19] A realistic compliance position is to treat these as foreseeable risks that warrant controls, rather than as hypotheticals to be ignored or as already litigated issues.

## AI Companion Compliance Requirements

**For compliance teams, this translates into the following concrete requirements:**

- no targeted advertising or upselling through companionship tools;
- no scripts that encourage secrecy or suggest withholding concerns from staff or family;
- review and approval of vendor content libraries and prompt templates for manipulative patterns or high-risk advice; and
- escalation paths when staff see the tool being used in ways that increase residents' vulnerability.

Federal consumer protection authorities have publicly emphasized the scale of fraud and manipulation affecting older consumers and highlighted how AI technologies such as voice cloning can exacerbate those risks.[20] Organizations that deploy human-like AI systems around older adults should expect scrutiny around disclosures, consent, data use, and safeguards against manipulation.[21]

# Transparency and Resident Expectations

From a compliance lens, transparency is not just a best practice; it is central to avoiding deception claims and to satisfying many privacy law notice requirements. Residents and, where appropriate, families should understand that:

- the system is AI and not a human being;
- conversations may be recorded or transcribed;
- certain data may be accessible to staff, family members, and vendor personnel;
- data may or may not be used to train or improve models, and what choices they have; and
- they can pause, mute, or permanently stop interactions.

This is particularly important in memory care and other settings involving cognitive impairment, where capacity to understand and remember disclosures may fluctuate.[22] Compliance programs should mandate recurring, plain-language disclosures (for example, brief on-screen reminders or staff prompts) rather than relying solely on admission packets or vendor terms of service.

# Consent and Decision-Making Capacity

## Consent in elder care should not be a one-time event.

Capacity may fluctuate, and residents who initially agree to use a tool may later find it confusing or distressing. Families and legal agents may have authority in some cases, but day-to-day reality often involves staff responding to informal preferences and behavior.[23]

## Good programs treat consent as an ongoing process.

They provide simple opt-out mechanisms, revisit consent periodically, and apply heightened safeguards in memory care and similar settings. Physical mute buttons, voice commands to stop listening, and staff procedures for disabling devices are not optional conveniences – they are part of the control environment.

## Consent for cognitively vulnerable populations must be ongoing and supported.

Ethics and digital health literature increasingly emphasize that consent for cognitively vulnerable populations must be ongoing and supported, rather than presumed to be indefinite based on a single initial agreement. For compliance teams, this suggests:

- documenting how consent is obtained and refreshed;
- training staff on when to revisit consent; and
- building audit-ready logs where feasible, without turning the process into something residents experience as bureaucratic.

# Surveillance, Autonomy, and Dignity

AI companionship can quietly drift into continuous monitoring through always-on microphones, ambient listening, or integrated sensors. In environments that residents consider their homes, this can risk undermining autonomy and dignity if not carefully constrained. While these tools can support independence and safety, they also require careful calibration to ensure monitoring remains proportionate and respectful.

**A workable compliance framing is "dignity-by-design."** Monitoring should be:

- proportional to the stated purpose (for example, mood support or engagement),
- not a broad surveillance tool;
- understandable to residents and families, including clear explanations of what is and is not being captured; and
- meaningfully controllable by the resident or their representatives through simple settings rather than opaque defaults.

Organizations should work with vendors to avoid configurations that could function as continuous listening devices without clear resident awareness. Where feasible, they should default to activation in response to clear resident cues (such as a wake word) rather than constant listening, especially in private spaces, and document those configuration choices in the AI or privacy inventory.

# Compliance Safeguards to Prevent Manipulation and Over-Reliance

Anthropomorphic systems invite emotional attachment, particularly when they use human-like voices, avatars, or self-disclosures. In vulnerable populations, that attachment can lead to over-reliance on the AI system for comfort or advice.[24]

**For compliance, the question is what guardrails are in place. Practical safeguards include:**

- prohibiting the system from making urgent recommendations involving money, medication, or safety decisions;
- ensuring that scripts do not discourage residents from speaking with staff or family about concerns;
- defining what categories of advice are off-limits; and
- training staff to recognize signs of over-attachment or distress and to adjust use – or turn off the tool – when necessary.

These controls should live in written policy and vendor standards, not just in informal expectations.

# Security and Access Control

Always-on audio, cloud storage, and staff or family portals expand the attack surface for sensitive information. Weak access controls or unclear retention practices can expose deeply personal conversations and emotional profiles.[25] From a compliance perspective, companionship tools that perform emotion inference should be risk-ranked alongside other high-impact systems.

**Baseline expectations include:**

- role-based access, with only appropriate staff able to view conversational summaries or usage dashboards;
- audit logging for access and configuration changes;
- clear retention schedules and secure deletion, with contractual commitments from vendors;
- encryption in transit and at rest consistent with other high-sensitivity systems; and
- defined incident response and breach notification procedures that account for the reputational impact of compromised conversational data.

Given the documented prevalence of impersonation and fraud targeting older adults, protection of audio and transcript data is especially critical, both to guard against identity-related harm and to prevent potential misuse in social engineering or voice cloning schemes.[26]

# Governance: Building an Operating Model That Holds Up

For in-house compliance teams, the central task is to integrate AI companions into existing risk and governance structures, not to run a one-off pilot. Programs should be classified in your AI or high-risk processing inventory, assigned a clear business owner, and subject to periodic review.

Vendors update models, add features, and change training practices; staff and resident populations evolve; and laws are shifting.

Programs should be reviewed at least annually, and sooner when material changes occur. Metrics such as privacy complaints, opt-out rates, incidents involving misuse or over-reliance, and staff workarounds can provide early warning signs that adjustments are needed. AI governance and safety reports recommend monitoring user complaints and adverse incidents as leading indicators in high-risk deployments.[27]

# Vendor Due Diligence and Contracting

From a compliance standpoint, much of the risk is determined before the system ever goes live – at the vendor selection and contracting stage. If the system generates resident-facing content, scripts, messages, or summaries, contracts should address human review expectations and ownership of outputs that incorporate resident data.

**Key compliance-driven points include:**

- clear restrictions on secondary use of resident data, including model training and marketing;

- commitments around content safety guardrails and escalation paths for harmful outputs;
- alignment with applicable privacy, security, and consumer protection laws, including state privacy acts and general unfair/deceptive practice standards;
- security requirements proportionate to the sensitivity of conversational and inference data; and
- audit and termination rights tied to privacy or safety failures.

Ethics and policy work on AI in elder care emphasizes that design alone will not solve dignity and safety issues; organizations must have enforceable standards, audits, and, where appropriate, certification or internal approval processes for AI systems in care settings.[28] Those same principles can be embedded into vendor due diligence and ongoing monitoring.

## Designing for Dignity and Human Connection

One of the most important questions a compliance-oriented review can ask is **whether the tool is designed to facilitate human connection rather than replace it.** Systems that prompt residents to call family, attend activities, or engage with staff align with that principle. Systems that quietly substitute for human interaction or encourage residents to spend most of their time in solitary AI conversations cut against this goal.[29]

Involving residents and families in rollout planning can materially improve outcomes. Demonstrations, plain language FAQs, and feedback channels build trust and surface concerns early, before they become complaints or headlines. Documenting that engagement and feedback process provides evidence that the organization considered dignity and autonomy in its deployment decisions.

## Closing Thoughts for Counsel and Compliance Leaders

Emotionally intelligent AI companions sit at the intersection of innovation and vulnerability. They may offer real benefits in addressing loneliness and supporting engagement, but they also pose foreseeable, governable, and increasingly visible risks to regulators and advocates. Senior living and home care providers are uniquely positioned to implement these tools in ways that enhance, rather than replace, human care.

**For in-house counsel and compliance leaders, the path forward is intentional design and governance by:**

- treating these tools as high-sensitivity programs,
- demanding clarity and control from vendors,
- embedding consent and dignity into operations, and
- measuring success not only by engagement metrics, but by whether residents remain respected, autonomous, and connected to real people.

Done well, AI companionship can support human care without eroding it. Done casually, it can create privacy, dignity, and trust failures that no amount of innovation will justify.

[1] Jeena Joseph, "Designing for Dignity: Ethics of AI surveillance in older adult care," *Frontiers in Digital Health* (2025), https://pmc.ncbi.nlm.nih.gov/articles/PMC12411420/.

[2] LeadingAge, *Comments on AI in Clinical Care* (Feb. 22, 2026), https://leadingage.org/wp-content/uploads/2026/02/LeadingAge-Comments-on-AI-in-Clinical-Care-RFI-FINAL-022326.pdf.

[3] *Id.*

[4] LeadingAge, LeadingAge Offers HHS Input on AI's Use in Clinical Care (Mar. 11, 2026), https://leadingage.org/leadingage-offers-hhs-input-on-ais-use-in-clinical-care/.

[5] H. Abdollahi & M. H. Mahoor, Artificial Emotional Intelligence in Socially Assistive Robots for Older Adults: A Scoping Review, *IEEE Trans. Affective Comput*. (2022), https://pmc.ncbi.nlm.nih.gov/articles/PMC10569155/.

[6] *Id.*

[7] J. Chu *et al.*, Enhancing Elderly Care Services Through Integrated Sentiment Analysis and Knowledge Reasoning (2025), https://www.semanticscholar.org/paper/Enhancing-Elderly-Care-Services-through-Integrated-Ai-Chu/.

[8] Joseph, supra note 1.

[9] Saadati Gottschlich *et al.*, AI in Elderly Care: Understanding the Implications for Independence, Privacy, and Dignity, *AI Tech. & Behav. Soc. Sci*. (2024), https://journals.kmanpub.com/index.php/aitechbesosci/article/view/2741.

[10] *Id.*

[11] Am. Hosp. Ass'n, AHA Response to HHS RFI on AI in Health Care (Feb. 22, 2026), https://www.aha.org/lettercomment/2026-02-23-aha-response-hhs-rfi-ai-health-care.

[12] NIST, *supra* note 6.

[13] AHA, *supra* note 13.

[14] *Id.*

[15] LeadingAge, *supra* note 2.

[16] Fed. Trade Comm'n, Fighting Back Against Harmful Voice Cloning (Apr. 7, 2024), https://consumer.ftc.gov/consumer-alerts/2024/04/fighting-back-against-harmful-voice-cloning.

[17] Sarah Barrington *et al.*, People Are Poorly Equipped to Detect AI-Powered Voice Clones, *Sci. Reports* (2025), https://www.nature.com/articles/s41598-025-94170-3.

[18] Blanka Klimova *et al.*, Ethical Considerations of AI Use by the Elderly, *Int'l J. Hum.–Computer Interaction* (2025), https://www.tandfonline.com/doi/full/10.1080/10447318.2025.2531274.

[19] Elena Portacolone *et al.*, Ethical Issues Raised by the Introduction of Artificial Companions to Older Adults with Cognitive Impairment, 21 *BMC Med. Ethics 1* (2020), https://pmc.ncbi.nlm.nih.gov/articles/PMC7437496/.

[20] Fed. Trade Comm'n, Preventing the Harms of AI-enabled Voice Cloning (Nov. 15, 2023), https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/11/preventing-harms-ai-enabled-voice-cloning.

[21] International AI Safety Report 2026, Extended Summary for Policymakers (2026), https://internationalaisafetyreport.org/publication/2026-report-extended-summary-policymakers.

[22] Brooke Wolfe *et al.*, Caregiving Artificial Intelligence Chatbot for Older Adults and Their Caregivers: Mixed-Methods Study, *J. Med. Internet Res*. (2025), https://pubmed.ncbi.nlm.nih.gov/40080043/.

[23] Wolfe *et al.*, *supra* note 29.

[24] Portacolone *et al.*, *supra* note 23.

[25] Abdollahi & Mahoor, *supra* note 7.

[26] Fed. Trade Comm'n, Fighting Back Against Harmful Voice Cloning, *supra* note 20.

[27] Int'l AI Safety Report Expert Panel, International AI Safety Report 2026 (extended summary) (2026), https://arxiv.org/abs/2602.21012.

[28] Joseph, *supra* note 1.

[29] Chu et al., *supra* note 9.

---

## Related People

**Jason J. Oliveri**

Partner

212-471-6237

## Related Capabilities

Aging Services

Data Privacy, AI & Cybersecurity

Healthcare

Healthcare Regulation, Compliance & Licensing

Regulatory & Compliance

## Related Insights

Jason Oliveri Unpacks Cybersecurity and Healthcare Compliance Risks With AI Agents

From “Find a Doctor” to “Call a Regulator:” Why Hospital Websites are the Next Privacy Scandal