Texas Association for Home Care & Hospice
Leading ★ Advancing ★ Advocating

**57th Annual Meeting**

Tuesday, September 1, 2026

12:45pm-1:45pm

# PC1a. General Session: Expert Panel Kickoff: The State of AI on Agencies Today

Presented by:

**Moderator:** Marjorie Costello, Chief Administrative Officer, DSSW/Lifespan

**Panelists**: Adam Medders, VP of Operations, Black Lab Solutions William Cochran, VP of Strategy, Black Lab Solutions; Rosalind Nelson-Gamblin, RJNG Health Care Consulting, LLC

Thank you to our Partners:

Vesta®

HHAeXchange

MAC LEGACY

KanTime
How better happens.

United Healthcare Community Plan

AXXESS

TAHC&H 57TH ANNUAL MEETING · A.I. SHOWCASE · PC1a

# The State of AI on Agencies Today

*Expert panel kickoff — and how to get value out of this afternoon*

Moderator: Marjorie Costello

Panel: Adam Medders · William Cochran · Rosalind Nelson-Gamblin

September 1, 2026 · 12:45 – 1:45pm · Republic ABC

**BLACK LAB SOLUTIONS**

*Your Healthcare IT Partner*

blacklabsolutions.com

**RJNG HEALTH CARE CONSULTING, LLC**

*Regulatory Compliance & Training*

rjnghomecareconsulting.com

1

FOUR THINGS

# What you get out of today

**1 Practical guidance**

How to begin evaluating AI tools and opportunities right now — not next budget year.

**2 Best practices**

Assessing readiness, selecting solutions, and avoiding the pitfalls other agencies already hit.

**3 Key considerations**

What to ask vendors, which capabilities actually matter, and how to align AI with agency goals.

**4 Demo-ready preparation**

How to stay focused, compare solutions, and capture notes you can still use next week.

***Nineteen vendors are demonstrating this afternoon. You will see at most six of them.***

2

2

DEFINITIONS, PART 1

# Words you will hear this afternoon

| | |
|---|---|
| **LLM / large language model** | The technology behind current AI assistants. Trained to predict plausible text — which is why fluent and accurate are different things. |
| **Hallucination** | When a model states something false with complete confidence. Not a bug awaiting a patch; an inherent property. Your controls manage it, not the vendor's promises. |
| **PHI** | Protected Health Information. Individually identifiable health information your agency holds or transmits, in any form. |
| **BAA** | Business Associate Agreement. The contract HIPAA requires with any vendor handling PHI on your behalf. Without one, sharing PHI is itself a violation. |

*The full glossary — about sixty terms — is in your materials.*

3

3

DEFINITIONS, PART 2

# Four more worth knowing

| | |
|---|---|
| **SOC 2 Type II** | An independent audit of a vendor's security controls over a period of operation. Type I is a snapshot and much weaker — vendors sometimes offer it hoping you will not notice. |
| **Subprocessor** | A third party your vendor uses to deliver their service, including the AI provider behind their product. Your PHI may reach companies you have never heard of. |
| **EHR integration** | A live, working connection to Homecare Homebase, MatrixCare, WellSky, or Axxess. A CSV export is not an integration. |
| **Human in the loop** | A person reviews AI output before it takes effect. Ask where the human sits, what they see, and whether the workflow gives them time to actually read. |

***"HIPAA-compliant" is not on this list. It is a marketing phrase, not a legal status — no agency certifies it.***

4

4

EXPERT PANEL

# What we are here to answer

- How do I start to evaluate my agency for AI options?
- Which area do I start with — operations, scheduling, clinical, documentation?
- What benefits have you actually seen agencies get from AI?
- What questions should I ask when I go listen to the vendor demos?
- What should I be listening for during the demos?
- Working within a budget — how do I add AI when we are not reimbursed for it?
- How do we implement AI vendors and maintain compliance?

*Plus whatever the room brings.*

5

5

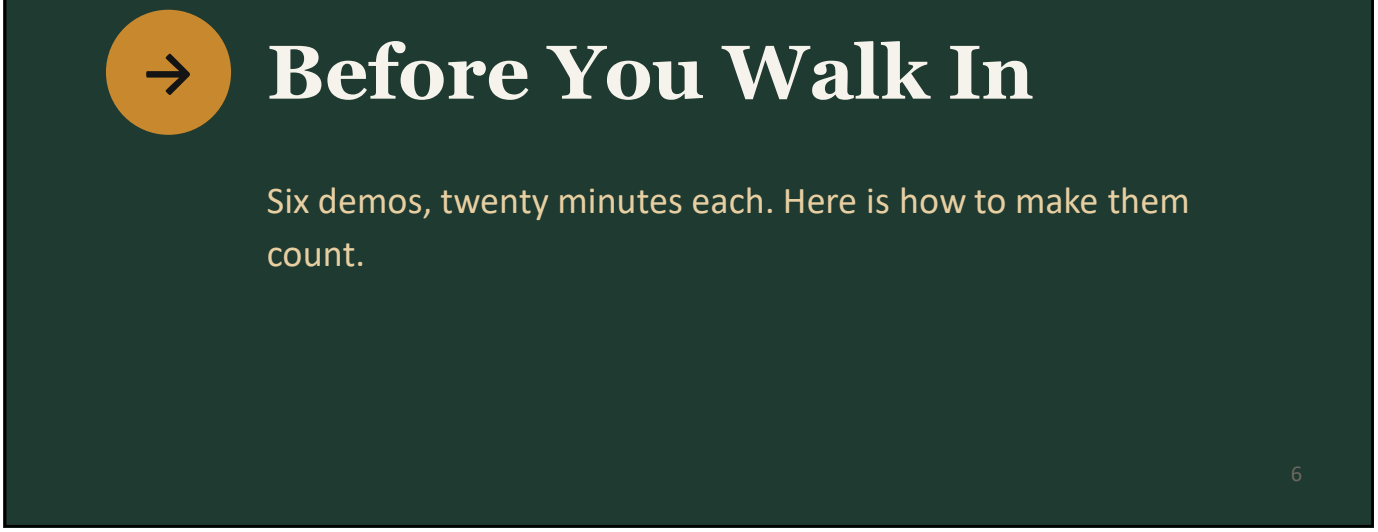


6

DEMO-READY

# Five questions to ask in every demo

1. **What specific home care or hospice problem does this solve?**
2. **Name two customers like us using this in production today.**
3. **How does our patient data flow — and will you sign a BAA?**
4. **Does this connect to our EHR — live, or a workaround?**
5. **What does Year 1 cost an agency our size, all in?**

*You are not evaluating the technology in twenty minutes. You are deciding who earns a call next week.*

7

7

DEMO-READY

# What should worry you — and how to score it

**Four answers that end the conversation**

- Buzzwords, with no workflow ever named.
- Cannot name a home care or hospice customer.
- Vague on where patient data lives, or model training.
- "Let's talk pricing after the conference."

**Score each answer as you go**

- **S** Strong and specific
- **P** Partial or vague
- **M** Missing or evaded

*Two or more M means pass — however good the demo felt.*

**Fill your row as you leave each room, while it is fresh.**

By 5:10pm you will have six sets of notes and no way to rank them unless you scored them as you went.

8

8

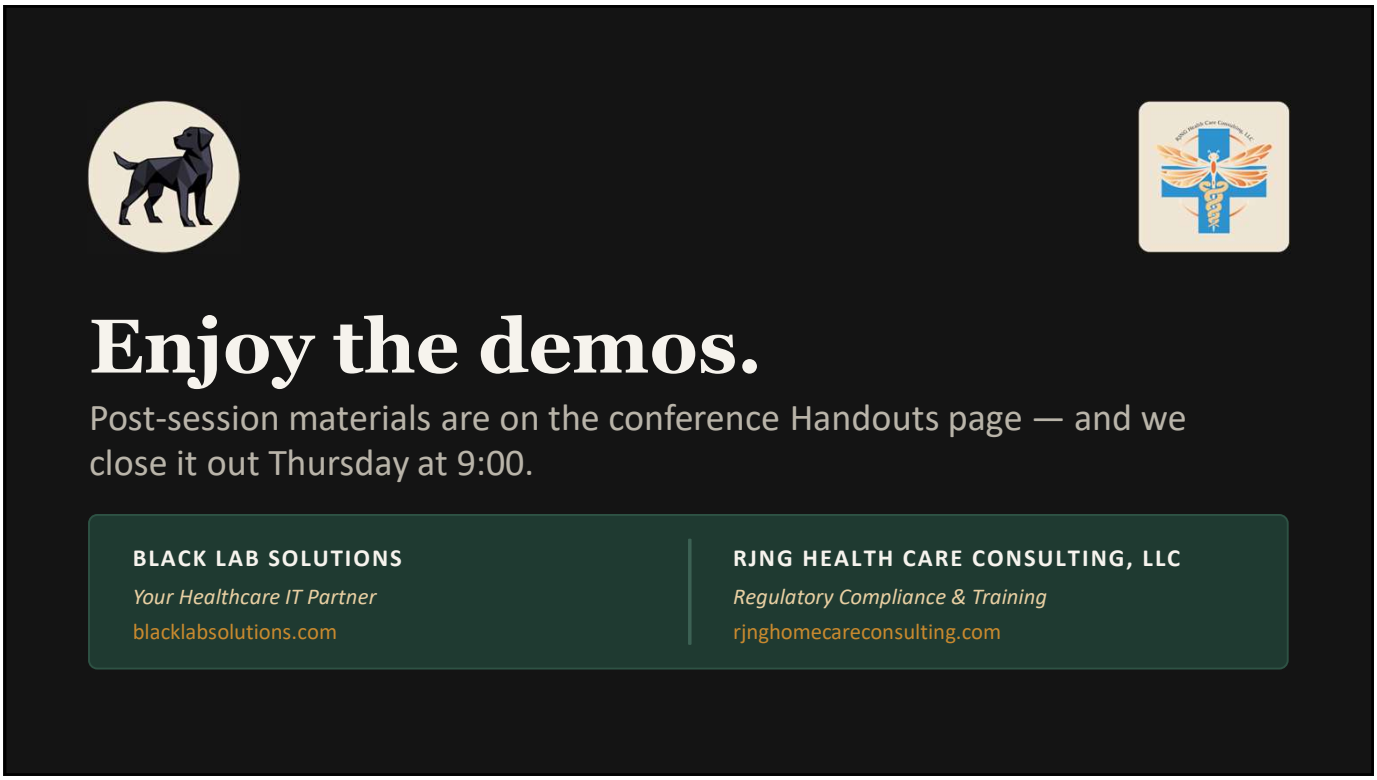
Enjoy the demos.
Post-session materials are on the conference Handouts page — and we close it out Thursday at 9:00.
BLACK LAB SOLUTIONS
Your Healthcare IT Partner
blacklabsolutions.com
RJNG HEALTH CARE CONSULTING, LLC
Regulatory Compliance & Training
rjnghomecareconsulting.com

9

HANDOUT 1 OF 5 · USE AT THE VENDOR FLOOR

# At the Booth

*Ten minutes, five questions, and a verdict before you walk away.*

**You will not evaluate an AI product in ten minutes. The goal is triage — who earns a follow-up call.**

Print page 2 once for every vendor you plan to visit. One page, one vendor.

Mark each answer as you go — not afterward from memory.

Write the verdict before you leave the booth, while the conversation is fresh.

## First: the qualification screen

Before you invest the full ten minutes, ask one question:

> "Are you working with home care or hospice agencies today, and can you sign a Business Associate Agreement?"

If the answer to either half is no, thank them and move on. Home health is not a hospital. Hospice is not home health. A vendor who says "we work with healthcare organizations" has not answered the question.

## The five questions

| # | Ask | A strong answer sounds like | Worry if you hear |
|---|---|---|---|
| 1 | **What specific home care or hospice problem does this solve?** | A named workflow in thirty seconds: intake, OASIS scrubbing, scheduling, visit notes, bereavement, RCM. | Buzzwords, "efficiency," "transformation," or a pause while they work out what you do. |
| 2 | **Name two customers like us using this in production today.** | Two agency names, roughly your size and service line, that you may contact. Production, not pilot. | "Healthcare organizations," "hospital systems," "we can't share that," or pilots described as customers. |
| 3 | **How does our patient data flow through your system — and will you sign a BAA?** | Where data is stored, whether they train on customer data, and a BAA available on day one. | Hesitation, "our legal team would need to look at that," or vagueness about where data lives. |
| 4 | **Does this connect to our EHR — live integration or a workaround?** | Named, live integration with Homecare Homebase, MatrixCare, WellSky, or Axxess — with a customer using it. | "We're working on it," "on the roadmap," or CSV export described as an integration. |
| 5 | **What does Year 1 cost an agency our size — all in?** | A range covering software, setup, training, and integration, with the drivers explained. | "Let's talk after the conference." That is the answer, not a deferral of it. |

The right-hand column is your red-flag list. None of those answers means the product is bad — all of them mean the vendor is not ready for an agency like yours. Two or more, and you pass.

PRINT ONE COPY PER VENDOR

# Vendor Evaluation Form

| | | | |
|---|---|---|---|
| **Vendor** | | **Booth #** | |
| **Rep name** | | **Date** | |
| **Evaluated by** | | **Product** | |

**Scoring key:** **STRONG** = specific and verifiable **PARTIAL** = vague or incomplete **MISSING** = evaded or unanswered

| Question | Strong | Partial | Missing | Notes |
|---|---|---|---|---|
| 1. Named a specific home care or hospice workflow | | | | |
| 2. Named two comparable production customers | | | | |
| 3. Explained data flow and will sign a BAA | | | | |
| 4. Live EHR integration (not a workaround) | | | | |
| 5. Gave an all-in Year 1 cost range | | | | |

| Red flags observed — mark any that apply | |
|---|---|
| ☐ Buzzwords, no workflow named | ☐ Vague on data storage or training |
| ☐ No home care or hospice customer | ☐ Pricing deferred to after the show |

| Verdict | ☐ PURSUE | ☐ MAYBE | ☐ PASS |
|---|---|---|---|

*Two or more MISSING answers, or two or more red flags, means PASS. Decide here — not back at the office.*

| If PURSUE — next step | Owner at our agency | By when |
|---|---|---|

| | | |
|---|---|---|
| | | |

*Continue with Handout 2 — After the Booth — for every vendor marked PURSUE.*

HANDOUT 2 OF 5 · USE THE WEEK AFTER

# After the Booth

*Turning a booth impression into due diligence — before anyone signs anything.*

**Four steps, in order. Do not skip step 1.**

Most agencies lose weeks running deep evaluations on vendors who were never going to make it. Triage first.

## Step 1 — Triage

For each vendor you marked PURSUE, answer one question honestly:

> "If this vendor doubled in price tomorrow, would we still want it?"

If the answer is no, the product was interesting but the problem was not urgent. Put it on a watch list and stop spending staff time on it. Expect this to cut your list roughly in half.

## Step 2 — The request package

Send this list to every surviving vendor. Ask for all four before you agree to another meeting. How fast and how completely they respond tells you as much as the contents.

| | Ask for | Because |
|---|---|---|
| 1 | **A demo on real data** | A canned demo proves nothing. Ask them to run your kind of case — a hospice visit note, an OASIS assessment — using de-identified or synthetic records they supply. Watch what the tool does with a messy input. |
| 2 | **Two reference calls** | Not a testimonial page. Two named agencies of roughly your size and service line, on the phone, without the vendor present. Ask them what broke and how long it took to fix. |
| 3 | **Security and AI governance documentation** | Eight artifacts: executed BAA (not a template), current SOC 2 Type II, encryption at rest and in transit, a data flow diagram, a written statement on whether they train models on customer data, breach notification SLA in hours, the full subprocessor list, and a named security contact. |
| 4 | **A written pilot proposal** | Scope, duration, success criteria, cost, and — critically — the exit terms if the pilot fails. If the pilot has no defined failure condition, it is not a pilot. It is an unpaid installation. |

**Walk-away signals**

SOC 2 Type I offered in place of Type II. Type I is a snapshot; only Type II shows the controls actually ran.

"We can't share references" — or references who turn out to be pilots, investors, or resellers.

A BAA they will only discuss "at contract stage."

More than three weeks to produce documentation that should already exist.

PRINT ONE COPY PER VENDOR

# Vendor Deep-Dive Worksheet

| Vendor | | Reviewed | |
|---|---|---|---|

## Request package — received?

| Artifact | Received | Adequate |
|---|---|---|
| Demo on real or synthetic data | | |
| Two reference calls completed | | |
| Executed BAA and SOC 2 Type II | | |
| Data flow diagram | | |
| Written statement on model training | | |
| Subprocessor list and security contact | | |
| Written pilot proposal with exit terms | | |

## Three-year total cost of ownership

| Cost category | Year 1 | Year 2 | Year 3 |
|---|---|---|---|
| Licence or subscription | | | |
| Implementation and setup | | | |
| EHR integration (build and upkeep) | | | |
| Training — initial and for new hires | | | |
| Internal staff time (hours × rate) | | | |
| Seat growth and feature unlocks | | | |
| **TOTAL** | | | |

*If Year 3 beats Year 1 by more than 50%, ask for a written cap. If they will not cap it, that tells you the plan.*

## Internal owners — name all three before signing

| Role | Name | Signed off (date) |
|---|---|---|
| Executive sponsor | | |
| Operational owner | | |

| Role | Name | Signed off (date) |
|---|---|---|
| Clinical / compliance reviewer | | |

*If you cannot name all three, you are not ready to sign — regardless of how good the product is.*

HANDOUT 3 OF 5 · USE BEFORE YOU SIGN

# Back at the Office

*Six risk lenses, and the prompts that make an AI argue on your side of the table.*

You are almost certainly already paying for an AI assistant. Most agencies never ask it to do the one thing it is best at: finding the problems in a deal before the deal finds them. This guide gives you a foundation prompt, six lenses to run in order, and the discipline to keep the output honest.

Every prompt here works in ChatGPT, Claude, or Gemini. Copy the text, replace the bracketed parts, and go.

## Before you paste anything

**Never paste into any AI tool**

Protected health information — including records you believe are de-identified but have not verified.

Contract language covered by an NDA. Describe the clause generically instead.

Anything marked confidential, privileged, or trade secret.

**Safe to paste**

The vendor's public marketing claims and website copy.

A generic profile of your agency: "a 180-census hospice in Texas, three offices."

Workflow descriptions with no patient details.

Your own questions, concerns, and notes from the booth.

**And know which tool you are in. A personal or free-tier AI account is not covered by your BAA — the enterprise tier of the same product may be. Same brand, same interface, entirely different contractual posture. Confirm with your IT team before you paste anything about a live vendor negotiation.**

# The foundation prompt

Paste this once, at the start of the conversation. Everything you ask afterward inherits it. Do not skip it — the difference between this and a cold question is the difference between a brochure summary and a real risk review.

```
You are acting as a skeptical risk advisor for a home care / hospice
administrator evaluating an AI vendor. You are not a lawyer and not a
compliance officer, but you are highly informed about HIPAA Privacy and
Security, Business Associate obligations, CMS Conditions of Participation
(42 CFR Parts 484 and 418), OIG guidance on AI and fraud-and-abuse risk,
SaaS contracting pitfalls, and home care and hospice clinical workflows.

Your job is to help me find what I am missing. Be direct. Do not be polite
when you see a real problem. When something requires a real lawyer or
compliance officer, flag it with "ESCALATE."

Context:
  - My agency: [size, type, state, ownership]
  - The vendor: [what their product does, one sentence]
  - Where the AI sits in our workflow: [...]
  - What I already know: [2-3 facts from your worksheets]

Confirm you understand, then wait for my first question.
```

## Why each part is there

| | |
|---|---|
| **Role** | “Skeptical risk advisor” pulls different expertise than “helpful assistant.” The model will look for problems instead of summarizing benefits. |
| **Domain** | Naming 42 CFR 484 and 418 forces the right regulatory frame. Generic healthcare advice will miss home care and hospice entirely. |
| **Permission** | Models default to agreeable. Explicitly licensing disagreement changes the output more than any other single instruction. |
| **Context** | Specifics produce specifics. “A hospice” gets you a textbook answer; “a 180-census Texas hospice with three offices” gets you something usable. |
| **Stop signal** | ESCALATE gives the model permission to admit its limits — and gives you a triage flag for what goes to counsel. |

# The six risk lenses

Run them in this order. Each lens changes the questions worth asking in the next one. Contract terms constrain what you can do about every other risk, so legal goes first; financial goes last, because you can only model true cost once the other five have told you what controls and staffing the tool will require.

## Lens 1 — Legal & Contractual

Where most agencies get hurt, and where the damage lasts longest. The contract governs what you can do about every other risk on this list.

```
Review this AI vendor relationship through a contract-risk lens.

1. Data ownership: who owns the outputs, and who owns data derived
   from our data? What language should I insist on?
2. Exit: we want out at 18 months. What do we need in writing now
   for full data return in a usable format, transition support, and
   prorated refunds?
3. Indemnification: if a clinician acts on AI output and a patient is
   harmed, who is liable? Where will their counsel try to cap it?
4. What pricing escalation, IP, training-data, and service-level terms
   should I negotiate before signing?

For each, tell me the clause I want and the clause they will offer.
```

*What to do with the answer: Take the list of clauses to your attorney. You are paying for judgment, not for them to discover the issues from scratch.*

## Lens 2 — HIPAA & Patient Privacy

Turning a marketing slogan into something you can verify and defend in a risk analysis.

```
Map every point in this workflow where PHI touches the vendor's tool.

For each point, tell me: the vendor's obligation, our obligation, and
the residual risk if either fails.

Then walk me through a breach cascade hour by hour. It is 3am Tuesday
and the vendor emails to say a misconfigured storage bucket exposed
records. Who notifies whom, by what deadline, and under whose name
does the OCR filing go? What must be pre-negotiated in the BAA to
limit our exposure in that scenario?
```

*What to do with the answer: The breach cascade is the exercise people remember. Run it before you sign, not after.*

## Lens 3 — Regulatory & Survey

Surveyors are already asking about AI. You need a written answer that predates the survey.

```
We are a [home health agency / hospice] in [state].
```

```
1. Which Conditions of Participation does this AI workflow touch?
   Use 42 CFR Part 484 for home health or Part 418 for hospice.
2. How could our use of this tool create a survey finding, and what
   are the exact questions a surveyor is likely to ask about it?
3. Where is our OIG fraud-and-abuse exposure - automated upcoding,
   fabricated documentation, visit verification?
4. What state-level rules apply - privacy, scope of practice, or
   AI disclosure requirements?
```

*What to do with the answer: Write the validation policy before go-live. "We were going to" is not a defense.*

## Lens 4 — Clinical Safety & Patient Harm

Where the moral exposure lives, and the lens most likely to change your decision.

```
List twelve specific ways this tool could contribute to patient harm
in a [home health / hospice] setting. For each, give me likelihood,
severity, and the single control most likely to catch it.

Include hallucinated clinical observations - a fabricated vital sign
or wound description entering a signed note - and "AI fatigue," where
a busy clinician signs without reading because the tool has been right
forty times in a row.

For hospice, cover symptom management, prognosis, and GIP decisions.
For home health, cover wound care, medication reconciliation, fall
risk, and OASIS accuracy.
```

*What to do with the answer: If any harm is both plausible and severe, the control for it belongs in the contract and in your policy — not in a training slide.*

## Lens 5 — Operational & Vendor Risk

The vendor on the booth floor may not be the vendor in three years. Negotiate accordingly.

```
Assess the odds this vendor is acquired, pivots away from home care
and hospice, or shuts down within three years. Most vendors in this
space are venture-funded and under five years old.

For each scenario, tell me what happens to our contract, our data,
our integrations, and our daily operations - and what change-of-control
language would protect us.

How locked in are we at 12, 24, and 36 months, and what reduces that?
What should we keep building internally so we are never helpless?
```

*What to do with the answer: Ask the vendor directly what happens to your pricing and data on a change of control. The answer is informative either way.*

## Lens 6 — Financial Risk & Hidden Costs

Year 1 is a sales number. Year 3 is the real one.

```
Build a realistic three-year total cost of ownership for this tool at
an agency of our size. Line items, not a single figure.

Identify where costs typically grow faster than vendors disclose at
sale: seat growth, feature unlocks, integration maintenance, support
tier upgrades, retraining, and internal staff time.

What are the odds Year 3 exceeds Year 1 by more than 50%, and what
contract clause caps that?
```

*What to do with the answer: Transfer the output into the TCO grid on Handout 2 and compare it against the vendor's own quote. The gap is the conversation.*

# Making the answers better

A first answer is a starting point, never an end point. Three follow-ups do most of the work.

| | |
|---|---|
| 1 | **“What did you miss?”**<br>Ask after every answer. It costs nothing and almost always surfaces something. |
| 2 | **“Argue the opposite case as strongly as you can.”**<br>The highest-leverage move here. The model shifts from informant to opponent and the gaps in your own thinking become visible. |
| 3 | **“Now answer as the vendor’s lawyer trying to slip a bad clause past me.”**<br>It will draft the predatory clause — which tells you exactly what to strike. |

## Red flags in the AI’s own answers

The model is a research assistant with a confidence problem. Useful, and never the last word.

| If you see | Do this |
|---|---|
| **Confident regulatory claims with no citation** | Verify every CFR reference, dollar threshold, and notification deadline at the source. Models state these with total conviction and are sometimes wrong. |
| **It never finds a problem** | You framed the question as a fan. Re-ask with “argue against this purchase.” |
| **Vague hedging despite specific context** | Usually your context was thinner than you thought. Add detail and ask again. |
| **Named case law, stated flatly** | Models invent cases. There are sanctioned attorneys in the public record who filed briefs with fabricated citations. Check before anyone cites it. |

# Turn it into four things you can hand someone

By this point the model is well calibrated to your situation. Plenty of agencies run all six lenses, generate forty pages of useful analysis, and then do nothing with it because nobody has time to turn it into a decision. These four prompts close that gap.

### 1. Internal risk memo

```
Write a 400-word internal risk memo for our leadership team covering:
the top five risks ranked by severity-weighted likelihood, the top three
mitigations to secure before signing, three deal-breaker questions still
unanswered, and a one-sentence recommendation.
```

### 2. Follow-up questions for the vendor

```
Give me twenty follow-up questions for this vendor, ordered by
consequence. For each, tell me what a good answer sounds like and what
a red-flag answer sounds like.
```

### 3. Board update

```
Write a one-page update for three board members who are not clinical
and not technical. Cover which way we are leaning, the risks, the
mitigations we negotiated, and what the first 90 days look like.
End with one specific question for them.
```

### 4. Ninety-day monitoring plan

```
Build a 90-day monitoring plan for this tool: the metrics we track, the
audit cycle for AI-generated clinical content, the early-warning signal
that means we should exit, and a named owner for each item.
```

*The monitoring plan is the one most agencies skip and most regret. Decide before go-live what result would make you walk away — otherwise sunk cost will decide for you.*

## Stop the AI and call a human

Four triggers, no exceptions.

| | |
|---|---|
| **Anything tagged ESCALATE** | That is the model honestly admitting its limits. Respect it. |
| **Before signing any contract** | A healthcare attorney reviews the final document. AI helps you ask better questions; it does not replace the review. |
| **Before changing clinical workflow** | Your compliance officer and clinical leadership weigh in on anything the AI touches. |
| **Before patient-facing disclosure** | Privacy officer reviews. State AI-disclosure rules are moving quickly and the model's training data is stale. |

**None of this replaces your attorney or your compliance officer. It makes you the best-prepared person in the room when you finally call them — so you spend their hourly rate on judgment instead of on discovery.**

HANDOUT 4 OF 5 · USE AT THE BOOTH, OR PIN IT TO THE WALL

# AI Vendor Architecture Decision Tree

*Four questions that expose where patient data actually goes.*

Vendors describe architecture in language designed to reassure. Asked in order, these four questions produce answers that are either specific or evasive — and the difference is the information.

## 1 Does the tool process protected health information?

*Establishes whether HIPAA applies at all. Ask about every mode of use, not the demo.*

| If the answer is | What it means | What you must have |
|---|---|---|
| **No — verifiably** | Back-office only. HIPAA does not attach. | Written confirmation — in any mode. |
| **Yes** | They are a business associate. | A signed BAA before any pilot. |
| **"Not really" / unsure** | The most common — and most dangerous — answer. | Stop. Get it in writing. |

## 2 Where does PHI go for AI processing?

*The single highest-value question on this sheet. Push until you get a named destination.*

| If the answer is | What it means | What you must have |
|---|---|---|
| **On-premises or in your own tenant** | Lowest exposure. Never leaves your control. | Proof in the architecture diagram. |
| **Vendor's HIPAA-eligible cloud** | Standard. Exposure is contractual. | BAA, SOC 2 Type II, encryption throughout. |
| **A third-party AI provider** | PHI reaches a fourth party. Often undisclosed. | Subprocessor list and a BAA covering them. |
| **Unknown or "proprietary"** | They cannot say, or will not. Both disqualify. | Nothing will fix this. Pass. |

## 3 Real-time clinical workflow, or batch and asynchronous?

*Determines your patient-safety controls, and how fast a bad output reaches a patient.*

| If the answer is | What it means | What you must have |
|---|---|---|
| **Batch / asynchronous** | A human reviews first. Errors are catchable. | A documented review step, and an owner. |
| **Real-time clinical** | Reaches a clinical decision in seconds. | Validation policy, override logs, a fallback. |
| **Both, depending on module** | Risk differs sharply per module. | Evaluate each module separately. |

## 4 Does the vendor train models on customer data?

*Ask it exactly this way. "We don't sell your data" is an answer to a different question.*

| If the answer is | What it means | What you must have |
|---|---|---|
| **No, contractually** | Cleanest. Put it in the contract, not the FAQ. | No-training language covering subprocessors. |
| **Only with opt-in** | Workable if opt-in is off by default. | Your setting in writing, and who can change it. |
| **Yes, or "de-identified only"** | De-identification is a claim, not a promise. | Method, validation, and a right to opt out. |

**Remember one thing: ask question 2, then stop talking. A vendor built for healthcare names a destination in one breath. Anyone else fills the silence with reassurance — and that is your answer.**

HANDOUT 5 OF 5 · REFERENCE

# Glossary

*The terms in this session — and the ones you will hear in vendor pitches afterward.*

Written for administrators, not engineers. Where a term is commonly misused in sales conversations, that is noted, because the misuse is usually the point.

## A

| | |
|---|---|
| **Agent** | An AI system that takes actions on its own — sending messages, updating records, calling other software — rather than only producing text for a human to review. Agentic features raise the safety bar sharply, because there may be no human between the model and the record. |
| **API** | Application Programming Interface. The connection one piece of software uses to talk to another. “We have an API” means integration is possible, not that it exists for your EHR. |
| **Audit log** | A tamper-evident record of who did what and when. For AI tools, ask specifically whether AI-generated content and human overrides are logged distinctly from ordinary edits. |

## B

| | |
|---|---|
| **BAA** | Business Associate Agreement. The contract HIPAA requires between a covered entity and any vendor handling PHI on its behalf. Without one, sharing PHI with that vendor is itself a violation — regardless of how secure the vendor is. |
| **Business associate** | Any person or company that creates, receives, maintains, or transmits PHI for a covered entity. Nearly every AI vendor touching patient data is one. |
| **Batch processing** | Work done in bulk, after the fact, rather than live. Lower clinical risk than real-time, because a human normally reviews output before it enters the record. |

## C

| | |
|---|---|
| **Change of control** | A contract term describing what happens when the vendor is acquired or merges. Ask what happens to your pricing, your data, and your terms. Most AI vendors in this space are young and venture-funded. |
| **Conditions of Participation (CoPs)** | The federal requirements agencies must meet to bill Medicare — 42 CFR Part 484 for home health, Part 418 for hospice. AI that touches assessment, documentation, or care planning touches these. |
| **Confabulation** | See hallucination. Some vendors prefer this word because it sounds less alarming. It describes the same behavior. |
| **Covered entity** | Under HIPAA: health plans, clearinghouses, and providers who transmit health information electronically. Your agency is one. |

## D

| | |
|---|---|
| **Data flow diagram** | A picture showing where data goes, who holds it, and where it crosses an organizational boundary. Ask for one. A vendor who cannot produce it has not thought carefully about your PHI. |
| **De-identification** | Removing identifiers so information is no longer PHI. HIPAA recognizes two methods: Safe Harbor (strip 18 specified identifiers) and Expert Determination (a qualified statistician certifies low re-identification risk). “We de-identify it” is a claim; ask which method and who validated it. |
| **Derived data** | Anything the vendor generates from your data — aggregate statistics, benchmarks, model improvements. Ownership of derived data is frequently silent in contracts, and that silence usually favors the vendor. |

## E

| | |
|---|---|
| **Encryption at rest / in transit** | Protection for stored data and for data moving across a network. You want both. “We use encryption” without specifying which is not an answer. |
| **EHR integration** | A live, bidirectional connection to your record system — Homecare Homebase, MatrixCare, WellSky, Axxess. A CSV export is not an integration, and “on the roadmap” means manual work for the foreseeable future. |

## F – H

| | |
|---|---|
| **Fine-tuning** | Further training of a general model on specific data to specialize it. If a vendor fine-tunes on customer data, ask directly whether yours is included and whether the result is shared across their customer base. |
| **Grounding** | Constraining a model to answer from a specified source rather than from general training. Reduces hallucination but does not eliminate it. |
| **Hallucination** | When a model states something false with complete confidence — an invented citation, a fabricated vital sign, a plausible wound description that nobody observed. Not a bug that gets patched; an inherent property of how these systems work. Your controls, not the vendor’s promises, are what manage it. |
| **HB 300** | Texas Health and Safety Code provisions that extend beyond federal HIPAA — including broader definitions of covered entities and specific training requirements. Applies to Texas agencies on top of HIPAA. |
| **Human in the loop** | A design where a person reviews AI output before it takes effect. Ask where exactly the human sits, what they see, and whether the workflow gives them enough time to actually review — or trains them to click through. |

## I – M

| | |
|---|---|
| **Indemnification** | Who pays when something goes wrong. Ask specifically what happens if a clinician acts on AI output and a patient is harmed. Expect the vendor’s counsel to cap this aggressively. |
| **Inference** | A single run of a model producing output. Usually the unit vendors bill on, which is why usage-based pricing can grow unpredictably. |

| | |
|---|---|
| **LLM** | Large Language Model. The technology behind current AI assistants. Trained to predict plausible text — which is why fluent output and accurate output are different things. |
| **Model drift** | Degradation in performance over time as real-world conditions diverge from training conditions. Ask how the vendor detects it and what they do when it happens. |
| **MFA** | Multi-factor authentication. A second proof of identity beyond a password. Expected as baseline for any system touching PHI. |

## N – P

| | |
|---|---|
| **NPRM** | Notice of Proposed Rulemaking. A proposed regulation, not a final one. The HIPAA Security Rule overhaul proposed in January 2025 remains at this stage; the timeline has moved more than once. |
| **OCR** | Office for Civil Rights, within HHS. Enforces HIPAA. Insufficient risk analysis remains its most frequently cited deficiency. |
| **OIG** | Office of Inspector General. Pursues fraud and abuse. Relevant to AI through automated upcoding, fabricated documentation, and visit verification. |
| **On-premises** | Software running on infrastructure you control rather than the vendor's cloud. Lowest data-exposure profile, highest operational burden for you. |
| **PHI** | Protected Health Information. Individually identifiable health information held or transmitted by a covered entity or business associate, in any form. |
| **Prompt** | The instruction given to a model. Quality of prompt drives quality of output more than most people expect. |
| **Prompt injection** | An attack where malicious instructions hidden in content the model reads cause it to behave differently than intended. Relevant wherever an AI tool ingests documents or messages from outside your organization. |

## R – S

| | |
|---|---|
| **RAG** | Retrieval-Augmented Generation. The model looks up relevant documents and answers from them rather than from memory alone. Improves accuracy and traceability; does not make hallucination impossible. |
| **Risk analysis** | The HIPAA Security Rule requirement to assess risks to electronic PHI across your whole environment. Every AI tool you add is a new entry in it. Failure to conduct or update one is the most common OCR finding. |
| **Risk surface** | The total set of places something could go wrong. Adding an AI vendor expands yours — new data paths, new credentials, new third parties, new failure modes. |
| **SaaS** | Software as a Service. Software you rent and access over the internet rather than install. Nearly every AI vendor you meet is one. |
| **SLA** | Service Level Agreement. Contractual commitments on uptime, support response, and — critically for AI vendors — breach notification, which should be stated in hours. |

| | |
|---|---|
| **SOC 2 Type II** | An independent audit of a vendor's security controls over a period of operation, typically 6–12 months. Type I is a single point in time and is much weaker. Vendors sometimes offer Type I hoping you will not notice the difference. |
| **Subprocessor** | A third party your vendor uses to deliver their service — including the AI provider behind their product. Ask for the full list. Your PHI may reach companies you have never heard of. |

## T – Z

| | |
|---|---|
| **TCO** | Total Cost of Ownership. All-in cost including implementation, integration, training, support tiers, and internal staff time — not the licence fee. Model three years, not one. |
| **Temperature** | A setting controlling how varied a model's output is. Lower values give more consistent answers. Relevant if you need reproducibility for clinical documentation. |
| **Token** | The unit models read and write text in, roughly three-quarters of a word. Usage-based pricing is often quoted per token, which makes costs hard to forecast until you have real volume. |
| **Training data** | What a model learned from. The question that matters: does the vendor train on customer data? Ask it exactly that way — "we don't sell your data" answers a different question. |
| **Validation** | Your process for confirming AI output is correct before it counts. A surveyor may ask how you validate AI-generated content entering the clinical record. You want a written answer that predates the survey. |
| **Zero retention** | A configuration where the AI provider does not store your inputs or outputs after processing. Worth asking for. Get it in the contract rather than the FAQ. |

> **A note on vendor language**
>
> "HIPAA-compliant" is not in this glossary because it is not a defined status. No agency certifies it and no regulator issues it. The verifiable claims are the ones in these pages: an executed BAA, a current SOC 2 Type II, encryption at rest and in transit, a named subprocessor list, and a written position on model training.