

FROM THE WEBINAR — AI IN HR

Guidelines on **identifying your AI use case**, the **questions to ask any AI vendor**, testing if your data is ready, and more.

Here are the seven checks Nick Seewer walked through during our “Before You Bring AI into HR: A Decision Framework” webinar — and they're just as applicable to AI projects outside of HR.



Nick Seewer

Vice President of Data, Analytics, & AI
G&A Partners

Do You Actually Have a Use Case?

Flip the usual question. Instead of asking where AI could help, ask what's actually costing you time, money, or risk right now.

IS THE PAIN ALREADY SITTING THERE?

- 1 Is there a spreadsheet somebody rebuilds by hand every month?
- 2 Does the same question get three different answers depending on who you ask?
- 3 Is there a workaround everyone complains about but nobody's fixed?

⇒ A “yes” to any of these means you have a real use case. But if the idea started with the tool instead of a number you're trying to improve, that's a sign it's about looking modern, not fixing something real.

The 4 Questions to Ask Before You Talk to Any AI Vendor

These four questions take fifteen minutes. They'll spare you from buying a tool for a problem that doesn't exist.

1 What's the specific task, and who owns it? Not “AI could theoretically help somewhere.” A real task handled by a real person who feels that pain every day.

2 What does success look like in 90 days, in a quantifiable number? “It feels faster” is not an answer. Something like “we cut the time to answer a benefits question from three days to same-day” is.

3 Where does the data actually live, and who can vouch that it's accurate? This is the one everybody skips, and it can sink the project six months later.

4 What happens when the AI is wrong? Because it will be, sometimes. If the answer is “nothing changes,” dig into that carefully before you buy anything.

DIY Is Fine — Until It Isn't

Experimenting with free AI tools is how people learn what they can actually do. The risk shows up when a free, DIY tool for an experiment morphs into “the way we do things now” without anyone deciding that on purpose.

THE GUT CHECK

Does your experiment not touch any personal or confidential data, and there's only one person using it? By all means, try it this afternoon. But the moment it touches employee or candidate data, feeds a real decision like pay or hiring, or more than one person starts relying on it, get a real process behind it.

4 SIGNS A “QUICK WIN” IS QUIETLY BECOMING A BIGGER MESS

- 1 Nobody else can explain how it works.
- 2 It's touching employee, candidate, or benefits data.
- 3 Someone downstream is trusting what it produces without checking it.
- 4 It's been quietly running for months, and nobody's revisited it since launch.

Hasty experiments can keep working for a while, so that's not the signal to watch for. Ask instead: Can you say, right now, who's accountable for it?

Is Your Data Actually Ready?

Whether a tool is good and whether the data behind it can be trusted are two different topics. A great tool with clean instructions can still hand you a wrong answer if the data you've handed it isn't consistent.

BEFORE YOU ASK “IS THE TOOL GOOD,” ASK:

- 1 Is there only one definition per metric?** If “turnover” is calculated three different ways in three reports, the AI inherits whichever one it happens to touch first.
- 2 Is there a single source of truth,** or is your team still cross-referencing systems to get a real answer?
- 3 Is anyone manually correcting the data** (usually in a spreadsheet) before it's used? That cleanup step is invisible to AI, and it will work off the messy underlying data.

THE 10-MINUTE TEST, NO IT REQUIRED

Pick a number your leadership already relies on, like headcount, turnover, or PTO balance. Ask three different people, each using their usual method, to pull that same number. A match is a good sign. A mismatch means you just found your data problem.

The Minimum You Need Before You Start

None of this costs money or requires technical skill.

- 1 One definition per metric, agreed on ahead of time.
- 2 One person or team who owns the number being right.
- 3 A way to spot-check the AI's answer against something you already trust.

Measure the Right Thing

Beyond hours or dollars saved, there are four KPIs worth tracking:

- 1 Decision speed** — did this actually get you to a decision faster, not just complete a task?
- 2 Accuracy** — was the answer actually right, not just fast? A fast tool that's wrong 20% of the time has simply automated the mistake.
- 3 Business impact** — did retention improve, did a compliance risk close before it became a problem, did cost-per-hire drop?
- 4 Validation** — how much time goes into checking the AI's work, on top of the time it saved? Execs in one recent survey estimated AI saved them about 4.5 hours a week, then spent nearly 4 of those same hours double-checking it.

SOUNDS GOOD ON PAPER, DOESN'T TELL YOU MUCH:

usage rate

hours saved on its own

number of pilots launched

tool satisfaction

the AI's own confidence score

Where AI Should Never Take the Wheel


Two compliance traps catch small businesses off guard.

- 1** Assuming these rules only apply to big companies — several state AI-hiring laws apply regardless of headcount.
- 2** Treating an automated recommendation as more objective than a human one just because it's AI. AI does not eliminate bias.

THE SIMPLE RULE

Any AI-assisted decision touching hiring, pay, discipline, or termination must be reviewed by a real person. And that person needs to be able to discern why the decision was made rather than just repeat what the tool said. If you can't answer that “why,” you're exposed.

AI can draft the response and flag what's worth a second look. The decision itself still needs a person who can explain why. It prepares the material, a specialist signs off, and nothing goes out the door without that review.



A NOTE ON HOW G&A MADE THIS GUIDE

AI absolutely helped us create this guide for you. Nick Seewer and his team provided their thoughts and instructions at the start. AI put it all together for us. And those same humans reviewed, edited, and approved this. Just like we recommend for your AI project, expert human judgement should always be in the loop.