

Four claims to be sceptical about

THE LEAK REPORT

ISSUE 13 · 28 OCTOBER 2026 · ATINNOVATORS.IN



Buy the boring,
checkable
version.

From someone who builds these systems for a living

"It learns your business" ask what exactly gets better, and how you'd know

No visible source a wrong answer looks exactly like a right one

Fully autonomous do you want to hear about mistakes from a customer?

"It replaces a person" price the task, not the salary

Contents

03

Editor's note

The pitches aimed at our own industry land on us too

04

The four claims

And the question that exposes each

06

Four claims

And the question that exposes each one

07

Why each question works

What it forces the vendor to show

09

Turn it on us

The same four questions, asked of this company

10

Honest limits

What scepticism does not replace

11

How to run this

Ask, then listen for the specific answer

12

What it costs

Published, because a quote request wastes your time




The Leak Report is published by AT Innovators, an AI development studio building voice agents, retrieval assistants and workflow automation for owner-led service businesses.

Words and research **Ramanjit Singh**. Every figure in this issue is either published on our pricing page or stated as a target rather than a result.

atinnovators.in
Photography Campaign Creators, Flipsnack, Hiroshi Kimura, Kelly Sikkema, MJ Duford, Marcos Paulo Prado, Mick Haupt, Mina Rad, Paper Textures, Scott Graham, Sebastian Schuster, Vitaly Gariev, Wafiq Raza, engin akyurt via Unsplash.

We build these systems, and the pitches still land on us.



Four claims, and the question that exposes each one.

We build AI systems for a living, which means the sales pitches aimed at our own industry land on us too.

Four claims come up often enough to be worth naming plainly.

For each, the useful thing is not the objection — it's the specific question that exposes it, including the version you should be asking us.

None of this means don't buy. It means buy the boring, checkable version: the one where you can see the source, see the failure mode, and see the arithmetic.

The through-line, if there is one: scepticism is not the same as refusal. Most projects that survive these four questions are worth doing.

Ramanjit Singh Founder, AT Innovators




THE FEATURE

Four claims, four exposing questions

A vendor who cannot answer the question attached to their own claim has not built what they're describing — including us, on any of these four.

Turn every one of these on the source's own industry. Including the source.

One survives the question. One doesn't.



"It learns your business"

Retrieval over your own documents is real and useful — an assistant that answers from your policies, your pricing, your catalogue.

A model that quietly absorbs how your company works and improves on its own is not what most vendors are actually selling.

Ask what gets better, and how you'd know.



Fully autonomous, customer-facing

The technology can do it. Whether it should turns on one thing: do you want to hear about mistakes from a customer, or from a review step you control?

For most service businesses, draft-prepared-human-sends captures nearly all the saving with none of the exposure.

Ask what happens the first time it's expensively wrong.

Neither claim is false on its face. Both become misleading the moment the specific, checkable mechanism behind them goes unanswered.

The claim, and the question

01

"It learns your business"

Ask instead: what exactly gets better, and how would I know if it didn't?

02

Fully autonomous, customer-facing

Ask instead: what happens the first time it gets something expensive wrong, before it happens?

03

Anything that can't show its source

Ask instead: show me a wrong answer, and show me how I would have caught it.

04

"It replaces a person"

Ask instead: price the task, not the salary. What does this specific task cost to automate, and what does it save?

Turn every one of these on us. If we can't show a source citation, name a failure mode, or price a task instead of a headcount, that's a reason to say no to us specifically.



The question forces a specific, checkable mechanism.

LEARNING

"What exactly gets better" cannot be answered with a phrase — it needs a described process: what data, what change, what you'd observe.

If the answer stays vague after being asked twice, the phrase was marketing, not architecture.

Retrieval over your documents is the real, useful version of this claim — worth distinguishing from the vaguer one.

AUTONOMY

"What happens the first time it's expensively wrong, before it happens" cannot be answered after the fact.

If the review step only exists in hindsight, it was never designed — it was discovered.

Draft-prepared, human-sends is the version that survives this question for almost every customer-facing use case.



The question forces the vendor to describe the failure path.

SOURCING

"Show me a wrong answer, and show me how I would have caught it" is not answerable in words — it has to be demonstrated.

A vendor who cannot produce this has not built citation into the architecture, whatever the pitch deck claims.

This is the same argument as issue 12, applied as a due-diligence question rather than a design principle.



The question forces the vendor to produce, not describe.



The question exposes which arithmetic the vendor is actually using.

PRICING

"Price the task, not the salary" — what does automating this specific task save in hours, and what does it cost to build and run?

If a pitch leans on headcount reduction for a 12-person business, that is a sign the pricing is aimed at a narrative, not the actual arithmetic.

In a team of twelve nobody is doing exactly one job — removing a task usually gives someone their afternoons back, not a smaller payroll.

Ask us these four, specifically.



| A vendor who answers well is necessary. It is not sufficient.

Every claim above is one you should ask when we quote you a project — not just when someone else does.

1

What exactly gets better, and how would you know?

If our answer is a phrase rather than a mechanism, that's the same vague claim from the other side of the table.

2

What happens the first time it's expensively wrong?

We should be able to describe the review step before it's needed, not after.

3

Can you show me a wrong answer, and how I'd have caught it?

If we can't produce this, we haven't built what we're describing.

4

Are you pricing the task, or the headcount?

If our pitch leans on reducing your team size, that's worth questioning directly.

Good answers to due-diligence questions are necessary, not sufficient — they don't replace a proper scoping conversation about your actual workflow.

Scepticism isn't the same as refusal.

The point is buying correctly, not avoiding buying.

Most projects that pass all four questions are worth doing — this is a filter, not a rejection of the category.

Good answers are necessary, not sufficient.

A vendor can answer all four well and still build the wrong thing for your specific workflow. These questions don't replace scoping.

We have the same incentive every vendor has.

Read this issue the way you'd read any due-diligence checklist from a party being diligenced — useful, and not neutral.

Ask the question. Listen for the mechanism.



| A vague answer to a specific question is the answer.

1

Ask the exposing question, not the objection

"How would I know if it improved" gets further than "I don't believe it learns."

2

Listen for a mechanism, not a reassurance

"Trust us" and "we're transparent" are not mechanisms. A described process is.

3

Ask for a demonstration, not a description

Especially on sourcing — show me a wrong answer, don't tell me the system is accurate.

4

Ask us these same four questions

Before you buy from us. If we dodge one, that's the answer.

Buy the boring, checkable version. Send us the workflow, and we'll tell you honestly whether it needs AI at all.



WHAT IT COSTS

Published, because
“contact us for a quote”
wastes your time.

AUDIT	Free	One workflow, mapped. Often concludes you should not build yet.
PROTOTYPE	Contact for pricing	A working slice you can put in front of one real user.
PILOT	Contact for pricing	A system a real team uses in production, with a run log.

Send me the workflow, and I'll tell you if it needs AI at all.

The useful first message is one sentence: what gets missed, repeated, delayed, or is hard for you to see.

ATINNOVATORS.IN