Who owns what?

Varun Choraria · 2026-05-17 · https://www.varunchoraria.com/who-owns-what/

Everyone these days is asking the same question in the age of agentic AI: what should an AI do vs. what should a human do?

Some background first.

I remember last year (2025), in a previous org, I championed the use of customGPTs to break down knowledge and skill gaps within the product marketing team. The problem statement by leadership was simple: writing is core to product marketing, but not everyone is good at it.

The level 1 of solving that problem was to fix the technical writing side of things: it's hard to memorize use-cases, JTBDs of every new feature and enhancement released. Knowledge base / feature documentation is not always hygienic, and even if it is — there's a lag between when the latest and greatest is actually published.

Product marketing or even marketing is a frontline engine, and therefore needs to be the first to create the buzz in the market. Not having access to the latest technical docs can directly hinder the freshness of any collateral produced.

I owned training the customGPT via .md files because it could ingest a lot more data. Remember, this is a time when even the team plan on ChatGPT had limits to how much you could load. Older models like 5.3 and 5.4 could only cache so much and context windows would be tiny.

Technical problem — mostly solved.

Now, the creative writing problem. When you're in a larger public-listed product company, typically your writing needs to go through legal before it sees light of day. Besides being messaging-aligned and keyword-infused, writing had to be legal-approved as well.

The answer was guardrails. I instructed the customGPT to read certain .md files first which had the best-practice guardrails intact. The GPT would parse those files first, ask clarifying questions to an open-ended prompt like "Write me a landing page" and work with the user to build the content. It would then end the output with checks on guardrails.

In between, the user would have the autonomy to write in a tone they preferred.

Across a 40-member product marketing org spread across 2 business units, the quality and velocity both improved in under a month.

The evergreen question

Now in the agentic AI world, where everyone is using MCPs, skills, agents and subagents — things have become more advanced as AI percolates into deeper and more complex workflows, in addition to job descriptions themselves being re-written.

The question still hasn't changed: Who owns what?

The single mental model I've shared with my team at GTM Buddy was this: do the work once, automate it the next time. Instead of starting at absolute 0% next time, try pushing the start line to 20-30%. Over time, 100% will look a lot better in terms of both quality and token cost.

The older mental model from pre-AI about a job was that it's a static list of tasks, and hiring an agency or a full-time person is a way of outsourcing it. AI will replace the task, therefore the job. Much of the world is still operating on that mental model, whether they like to admit it or not. The clearest signal is them cutting down on headcount (easiest lever to pull) and squeezing more out of the remaining few. Up until now, the few at least had access to nearly limitless tokens, but in H2 of 2026 and beyond — the "get more done with less" will also seep into your AI usage.

The better mental model to go after this understanding what the JTBD is, not just the job:

> Jobs = tasks + judgement + context + ownership + metrics

This changes the hiring equation quite a bit, because EQ becomes a signal you look out for in candidates, and give lesser weightage to technical abilities. In the pre-AI world, most people who climbed the corporate ladder where inherently great at the people equation: change management, relationship building, taking solid judgement calls, and stepping up when things break.

The evergreen answer

As anyone in the management side of the business, you need to stop asking yourself this question:

> Can AI do the job?

And continually ask yourself:

> Does a task require context that can exist only between 2 humans? Eg. managing an executive relationship between a seller and a prospect.

If the answer is:

- Yes → Human or human + AI owns it
- No → AI owns it

Every quarter, as a management person, you need to sit down with each team member and work on updating a document around each person's swimlane.

- For each swimlane, list down the entire process. Specify each task in the process, and classify it as owned by AI / Human / Both. Alongside it, also write down what's happening in each task, and what to do / what not to do.
- This needs to be visited every Q, as a larger human + AI retrospective. This is what the new coaching would look like in the agentic AI world, between a manager and a direct report. It prevents skill sprawl, using AI as a hammer to every nail, and encouraging smarter AI usage.

- For the entire team, the manager can draw out and publish a boundary map: list the entire piece of the business a team impacts, swimlanes under each, and task distribution between AI and Humans. If you’re able to tie it to growth metrics — that’s a validation that you’re truly using AI to drive real business outcomes, not just becoming AI-native for the sake of it.