

A human behind the agent is fine. Saying so is the product.

Reuters reported on 22 September that Meta tested a "human concierge" for Muse: contractors quietly taking some of the calls the agent placed. Here is what that means for anyone who delegates work to an agent.

22 Sep Date Reuters published the report, citing internal company posts	8 Sep When Meta launched the Muse personal agent, weeks earlier	3 Verified details: the work happened, it was quiet, Meta calls it a test
---	---	---

PROBLEM	WHY IT MATTERS	WHAT GOOD LOOKS LIKE
Undisclosed handoff	The other side agreed to a call with a bot or a person — not to a swap mid-sentence	Introduce the colleague joining the call
Data spread	Whoever hears your call learns what your agent knows about you	A stated access and retention policy for anyone listening
Expectation drift	A human-polished launch makes your unattended results look worse	Judge an agent on unattended performance only

THE SINGLE QUESTION

For any agent you use: **who is in the loop, and would you know if that changed?** Voice is the surface where a mistake is heard by a stranger, so start there with calls where a wrong turn is cheap.

THE PATTERN TO WATCH

Every consumer agent will need humans somewhere — the hardest surface, the rarest task, the call that must go right. The labs that succeed will be the ones that disclose it.

- 1 Ask before you delegate a phone call**
Voice is the surface where your agent has the least room to recover, and where a mistake is heard by a stranger.
- 2 Make disclosure a default**
An agent that identifies itself changes how the other side behaves — and the outcome is usually better.
- 3 Write down what it may say out loud**
Spoken details leak in a way typed ones do not. Scope access before you give it a voice.

Sources: Reuters (22 September 2026), BNN Bloomberg, KSL. Full analysis at wolffi.sh/blog/meta-muse-human-concierge