

AI Meeting Assistant With No Bot: How Botless Capture Works (2026)

How botless AI meeting capture works: what a bot in your call really costs, the three ways a tool can hear you, and when a bot is still the better choice.

Published by Julian Pscheid · June 11, 2026

[Read this article online: https://www.hedy.ai/post/ai-meeting-assistant-no-bot/](https://www.hedy.ai/post/ai-meeting-assistant-no-bot/)



A woman wearing a headset alone at an office desk, speaking and gesturing during a video call, her laptop screen turned away and a face-down iPhone beside it

Quick answer A botless AI meeting assistant captures audio on the device you are already using instead of dialing a separate participant into your call. On macOS and Windows it captures the computer's own audio output, so it hears every person on the call. On iOS and Android it uses the microphone, which covers in-person conversations. Nothing appears in the participant list, and no calendar access is needed to start.

Most people arrive at this question from the same experience. You are on a call with a client or a candidate, and a participant appears in the list with a vendor's name on it. Sometimes it posts a message in the chat too. Everyone sees it, and you did not get to decide how that landed.

Most AI meeting tools were built bot-first because it was the easiest way to reach a meeting: get a calendar invite, dial in as a guest, take the audio the platform hands to attendees. Several mainstream tools have since added botless modes alongside that default, sometimes with features held back on the quieter path. So "does it use a bot" is no longer a question about the product. It is a question about which capture path you are on, and what that path can and cannot reach.

There are three paths. This guide compares those, not vendors.

What "a bot joined your call" actually means

A meeting bot is a real participant. The vendor operates an account that accepts your calendar invite, joins the meeting as a guest, and receives the same audio stream every other attendee gets. That is the whole mechanism, and each part of it has a consequence.

The bot is visible. It shows up in the participant list, usually carrying the vendor's product name, so your client sees which tool you bought. Your candidate sees it in the first thirty seconds of an interview, before they have said anything. On some platforms it also drops a line into the meeting chat. You cannot control the label, because the account is the vendor's and not yours.

It usually needs your calendar. Something has to know a meeting exists and when to join it, which in practice means granting a third party read access. That is not only meeting times. It is titles, attendee lists, and often the notes and links in the invite body, including meetings the tool will never join.

And it changes the room, which is the part that does not show up on a feature page. Once a named tool is visibly present, people phrase things more carefully and the half-formed idea stays unsaid. In a negotiation or a first call with someone who does not know you yet, that is the real cost of the bot, and it lands on the conversation rather than on your budget.

That does not make bots wrong. It makes them loud, and loudness is a choice you should get to make per meeting rather than inherit from your tool's architecture.

Three ways a tool can hear your conversation

Every AI meeting assistant reaches your audio through one of three paths. They differ in what they can physically hear, which is the part that decides whether a tool fits your work.

How it captures | What it can reach | What it costs you

A bot joins the call | Every participant on the call, whether or not you are there. Produces a recording that belongs to the meeting rather than to your device. | A visible participant carrying the vendor's name, calendar access so it knows when to join, and only works on calls that admit guests.

System audio capture | Everyone the computer is playing, on any conferencing app, plus imported audio and video. No participant list entry, no chat message. | Runs only on the desktop you are attending from. Does not hear the room around you unless you add the microphone.

Microphone capture | The room. In-person conversations, and calls played aloud on speaker from another device. | On a headset video call it reliably gets your side only, so it is the wrong path for a remote meeting unless the audio is in the air.

System audio capture is worth understanding properly, because it is the one that surprises people. Both desktop operating systems now expose the audio a machine is playing to an app that asks for it. On macOS that is a CoreAudio process tap, available on macOS 14.2 Sonoma and later. On Windows it is WASAPI loopback. Neither requires a virtual audio device, a driver, or any setup step from you. If you can hear the other person, so can the app, and the meeting platform is not involved in the arrangement at all.

Microphone capture is the older path and still the only one that works when there is no call. A doctor's appointment or a coffee-shop interview has no dial-in for a bot to accept. It has a room, and the room is what a phone microphone hears.

So pick your path from where the conversation physically happens rather than from a marketing label. Remote call on a laptop, use system audio. Conversation in a room, use the microphone. Meeting you are not attending, use a bot.

When a bot is the right answer

A bot does one thing no user-side capture tool can do.

If you need a meeting captured that you are not attending, you need a participant that can be there on its own. System audio and microphone capture both run on your device, in your session, while you are present. Turn off the laptop and the capture ends. No setting fixes that, because it is what the architecture is.

The same applies to a shared team record of a call you do not host. When a recording is meant to be the team's artifact rather than one person's notes, and it has to exist identically for everyone whether or not any particular person showed up, a bot that joins as its own attendee is the cleaner mechanism. Botless capture gives you your view of the conversation, on your device, tied to your attendance.

So if your job is running a call library that a manager reviews, or capturing meetings across a team you personally do not sit in, a bot-based tool is the right purchase. Hedy will not do that job, and Hedy Teams is a waitlist rather than a shipping product today.

Worth checking before you accept the bot: do you actually need those properties, or did you inherit them from a tool that only offered one way in? Most individual professionals capturing their own conversations do not need a meeting-owned recording. They need their own notes from a meeting they were in.

Botless does not mean consent-free

When a bot joins, the tool announces itself for you. Remove the bot and that announcement disappears. Your obligation does not, and the quiet architecture makes it easier to forget.

Recording and wiretapping laws vary by jurisdiction, and the variation is not cosmetic. Some places treat one party's consent as sufficient, which usually means your own. Others require every party to a conversation to agree before it can be captured, and getting that wrong is a legal problem rather than an etiquette problem. Cross-border calls can put two different rules on the same conversation. If your work routinely spans jurisdictions, this is a question for your own counsel and not for a vendor's blog post.

Beyond the law there is the ordinary professional case. People find out later. A summary you share, or a quote you repeat back with suspicious precision, can reveal after the fact that something was capturing, and the reaction is to the surprise rather than to the tool. Asking at the top of the call takes one sentence and removes the whole problem.

Removing the bot removes a notification, not an obligation. If your only disclosure was the vendor's name in the participant list, that was a side effect rather than a consent practice.

Two things help. We wrote word-for-word consent scripts (</post/ask-permission-to-record-meeting-consent-scripts/>) for the actual sentence to say in a client call, an interview, and a group meeting, because the reason people skip the ask is usually that they do not know how to phrase it. And if the question underneath yours is where the data goes rather than who knows about it, AI meeting privacy settings explained (</post/ai-meeting-privacy-settings-explained/>) walks through the settings that decide that.

Note also that "botless" and "private" are separate claims. A tool with no bot can still upload every second of your audio and every line of your transcript to someone else's servers. The bot question is about who knows. The processing question is about where the conversation goes, and you should answer both. Our guide to using AI meeting notes privately (</post/how-to-use-ai-meeting-notes-privately/>) works through the second one step by step.

How Hedy captures without a bot

Hedy has no bot in any flow. There is no vendor account that joins your meetings, because there is nothing to join with.

On macOS and Windows , Hedy captures system audio: the CoreAudio process tap on macOS 14.2 Sonoma or later, WASAPI loopback on Windows. It hears everyone on a Zoom, Teams, or Meet call without touching the meeting itself. No participant appears, no message lands in the chat, and no calendar access is required to start a session. You press start when the conversation starts. The Windows app has been generally available since 2025 with the same feature set as the Mac app, so this is not a platform where one side gets the quiet version.

On iOS and Android , Hedy captures microphone audio, which covers in-person conversations and a call played aloud on speaker from another device. One limit worth stating plainly: Hedy cannot capture a phone call on the same handset. iOS reserves the audio session for the call, and Android keeps call audio behind a system-only permission, so this is an operating system boundary rather than a missing feature. The supported live workaround is taking an iPhone call on a Mac through Apple's Continuity "Calls from iPhone" feature, where Hedy on the Mac captures both sides as system audio. The recording phone calls (</help/recording-phone-calls/>) article covers the setup.

While the session runs, Automatic Suggestions (</help/automatic-suggestions/>) push questions worth asking and points worth clarifying to you in the moment, which is the part a post-meeting summary cannot do at all. The free plan gives you the first 30 minutes of Automatic Suggestions per session, Pro gives you 120.

On processing: speech recognition runs on-device by default on every platform. AI analysis runs in the cloud by default, which keeps Hedy usable on hardware that cannot host a language model. Local AI Processing (</help/local-ai-processing/>) is the switch that moves the analysis step on-device too, on machines capable of running the models, offline included. It is opt-in and off by default, so nothing changes until you turn it on. Hedy also completed an independent SOC 2 Type I examination and an independent HIPAA assessment in April 2026, documented here (</post/hedy-soc-2-type-i-hipaa-complete/>) , alongside GDPR compliance and EU data residency options.

Requirements and supported platforms are listed in system requirements (</help/hedy-system-requirements/>) . Hedy is free at \$0 for 5 hours a month with no credit card, and Pro is \$12.99 a month or \$99.99 a year, with a \$299 one-time lifetime option. Same price on every platform, listed on the pricing page (</pricing/>) .

Frequently asked questions

What does "no bot" mean for an AI meeting assistant?

It means no extra attendee joins the meeting. Instead of dialing a vendor-operated participant into your call, the app captures audio on the device you are already using: the computer's own audio output on a desktop, or the microphone on a phone. Nobody sees a new name in the participant list, and nothing gets posted to the meeting chat.

Can an AI meeting assistant capture a call without joining it?

Yes, on a desktop. macOS and Windows both let an app capture the audio the computer is playing, which is every remote participant you can hear, on any conferencing app. Hedy uses a CoreAudio process tap on macOS 14.2 Sonoma or later and WASAPI loopback on Windows. Neither needs a virtual audio device or any setup from you.

Does removing the bot mean nobody knows the meeting is being captured?

That is exactly the risk. Losing the bot removes a notification, not an obligation. Recording and consent laws vary by jurisdiction and some require every party to agree before a conversation is captured. Say what you are running, at the top of the call, in your own words.

Can Hedy capture a phone call on the same phone?

No. iOS reserves the audio session for the call itself, and Android keeps call audio behind a system-only permission, so no third-party app can reach it. The supported live workaround is taking an iPhone call on a Mac through Apple's Continuity "Calls from iPhone" feature, where Hedy on the Mac captures both sides as system audio.

Does botless capture work for in-person conversations?

Yes, and it is the only architecture that does. A bot can only join a meeting that has a dial-in. For a conversation happening in a room, the microphone is the capture path, which is how Hedy works on iOS and Android. The same applies to a call played aloud on speaker from another device.

When is a meeting bot still the better choice?

When you need a meeting captured that you are not attending, or a shared team recording of a call you do not host. A bot is a participant in its own right, so it can sit in a meeting without you. A user-side capture tool runs on your device and stops when you do, so it structurally cannot cover either case.

Hedy AI · Live AI Coaching for Important Conversations

Try Hedy free: <https://www.hedy.ai/downloads/>

<https://www.hedy.ai/post/ai-meeting-assistant-no-bot/>