

PRODUCT · OCTOBER 2026

Frequently Asked Questions

The basics

What does Ovvos Wasel do? It interprets a meeting live. People in the room listen on their own phones, in their own language, with a live caption — about two to four seconds behind the speaker.

Does it need the internet? No. On the default on-premise profile it installs, starts and runs a meeting with no internet connection, and sends nothing outside the room PC.

How many languages? It can recognise 15 and speak 10. Any four can be offered at once, chosen per meeting.

How many people can listen? Listeners cost almost nothing: the work is done once per language, not once per person, so thirty people on one language cost the same as three. The practical limit is your WiFi.

How far behind the speaker is it? Typically two to four seconds. A sentence has to finish before it can be translated, so long sentences take longer.

Conferencing platforms

Does it work with Microsoft Teams / Zoom / Webex / Google Meet? Yes — with all of them, because it does not connect to any of them. It listens to what the room PC plays, the same sound the speakers carry.

Does it join the call or show up as a participant? No. Nothing is installed in the conferencing platform, nothing joins the meeting, and no platform account, plugin or approval is needed.

Can remote participants hear the interpretation? No. It interprets what your room hears, for the people in your room — like an interpreter sitting at your table.

Can it interpret people speaking in the room, not just the call? Yes. Set the audio source to a microphone or line-in, or feed it from the room's audio mixer.

Participants

Do participants need to install anything? No. They scan a QR code and the page opens in their phone's browser.

What do they need? A phone or tablet on the room WiFi, and headphones.

Can they change language during the meeting? Yes, from the options button, without scanning again.

What happens if their WiFi drops? The page reconnects by itself.

Is anything saved on their phone? Only two preferences: their language, so it is pre-selected next time, and light or dark display. Captions exist only while the page is open.

Security and privacy

Where does the meeting audio go? Nowhere but the room PC's memory, for the seconds it takes to process. It is never written to disk, and there is no setting that would.

Is a transcript kept? Not by default. Keeping one, exporting one and writing meeting text into logs are three separate settings, each off unless an operator turns it on.

Who can listen? Anyone on the room network who has the meeting PIN, if one is set. Without the PIN, the page offers nothing to choose.

Can it be used in an air-gapped room? Yes. That is what the on-premise profile is designed for. See the Network Behaviour Statement for how this was measured and how your team can verify it.

Does it send telemetry or check for updates? No.

How is it supported if you cannot reach our PC? Your operator opens a session with your own approved tool when they choose; your staff watch throughout; the session ends when you close it. For fully air-gapped rooms: telephone support with your operator at the keyboard, or an on-site visit under escort.

Accuracy

How accurate is it? It is built for comprehension in the room. Clear audio and a known set of languages give the best results; crosstalk, distant microphones and strong dialects reduce accuracy.

Can I make it get names and terms right? Yes — the glossary fixes names, ranks, units and designations per language.

Can it replace a certified interpreter? Not where a binding rendering, a certified translation or an official record is required.

Hardware and setup

What hardware does it need? One Windows 11 PC per room with an NVIDIA graphics card (RTX 3090 24 GB minimum, RTX 5090 32 GB recommended), plus a WiFi access point for the room.

How long does it take to start? About 45 seconds from opening the application to ready.

How is it installed? From one offline package that includes everything — software, models and voices.

Can we switch the translation model? Yes. An administrator can load a different translation model file from the Models page; the voices and speech recognition are fixed by the installation.