

PRODUCT · OCTOBER 2026

Product Overview

— Ovvos Wasel 1.0

What it is

Ovvos Wasel is a live interpretation system for meeting rooms. It listens to a meeting, recognises what is said, translates it into up to four languages at once, speaks each translation in a natural voice, and delivers the audio — with a live caption — to participants' own phones over the room WiFi.

It runs on **one Windows PC in the room**. On its default on-premise profile it needs no internet connection and sends nothing about the meeting outside that PC.

Who it is for

Organisations that hold meetings across languages and cannot, or would rather not, send meeting audio to a cloud service or bring interpreters into the room:

- Government ministries, defence and security organisations
- Diplomatic and protocol offices hosting foreign delegations
- Boards and executive committees with international members
- Critical-infrastructure and regulated-industry operators
- Any room where meeting content must stay inside the walls

The problem it solves

Sensitive meetings already cross languages. The usual answers each carry a cost:

- **Human interpreters** — cleared linguists are scarce, expensive and booked weeks ahead. Each additional language is another person who hears everything.
- **Cloud translation services** — fast and inexpensive, but the meeting audio goes to a third party, outside the facility and often outside the jurisdiction.

Ovvos Wasel is machine interpretation with the deployment posture of the customer's own infrastructure: no linguist in the room, no packet off the site.

How a meeting works

Remote participants speak (Teams, Zoom, Webex — any platform, or a microphone)

room PC plays the call on its speakers

Ovvos Wasel listens at the Windows audio layer (never joins the call)

detect sentence → recognise → translate (×4) → speak → deliver

room WiFi → each participant's phone → their headset + caption

1. **The operator** opens the app on the room PC. About 45 seconds later it is ready and asks which languages the room needs. They press **Start meeting** and open the **room display** on the wall screen.
2. **Participants** scan the QR code on the room display, enter the meeting PIN if one is set, and tap their language. Audio starts in their headset at once, with the current sentence written underneath.
3. **During the meeting** nobody touches anything. The operator's window shows latency, listeners and four loss counters if they want to watch.
4. **At the end** the operator presses **End meeting**. Interpretation stops at once on every phone and the room display clears.

Headline facts

Languages it can hear	15 — detected automatically or restricted to a set
Languages it can speak	10 — any 4 at once, chosen per meeting
Delay behind the speaker	Typically 2–4 seconds
Participants	Their own phones or tablets — no app, no account, no receiver
Cost of adding listeners	None — load scales with languages, not people
Internet	Not required
Conferencing platforms	Any — it never joins the call
Hardware	One Windows 11 PC with an NVIDIA GPU, per room

What makes it different

1. **On-premise and air-gap capable** – every stage runs on the room PC; no connection leaves it (measured – see the Network Behaviour Statement).
2. **Platform-independent** – captures sound at the Windows audio layer, so it works with any conferencing platform and needs no integration or approval.
3. **Nothing for participants to install** – a web page served by the room PC.
4. **Fan-out per language** – thirty listeners in one language cost the same as three.
5. **Built for operators who are not technicians** – plain-language readiness checks that name the fix, one-click meeting start and end.
6. **Honest about its limits** – designed for comprehension in the room, not as a certified translation or a record of proceedings.

Two profiles

	On-premise (default)	Cloud (optional)
Recognition	On the room PC	Azure Speech, Deepgram or OpenAI
Translation	On the room PC	Azure Translator, DeepL, OpenAI or Claude
Voices	On the room PC	Azure Neural TTS or ElevenLabs
Audio leaves the building	No	Yes, to the chosen provider
Internet needed	No	Yes
Running cost	None	Per language, per hour, from the provider

The cloud profile is off by default and does nothing until an operator enters API keys. For sovereign or classified environments it should stay off.

Where it is not the right tool

- A **certified or binding translation**, or an official **record of proceedings** – use a qualified interpreter or translator.
- **Interpretation back into the video call** for remote participants – Ovvos Wasel serves the people in the room.

- **Heavily accented or dialect-heavy speech in a noisy room** – accuracy drops; good audio in (a line feed, a close microphone) matters more than anything.