

PRODUCT · OCTOBER 2026

Features

Grouped by who uses them.

For participants – the phone page

Feature	Detail
Join by QR code or address	The room display shows a QR code and a plain address. Both point to the room PC, not a website.
No app, no account	Opens in the phone's normal browser – iPhone, Android, tablet or laptop.
Meeting PIN	Optional. Until it is accepted, the page offers nothing to choose.
One-tap language choice	Languages listed by their native names. Audio starts on the tap.
Live caption	The sentence being spoken, in large type, synchronised with the audio – it appears when the voice starts, not before. Earlier sentences fade above it.
Right-to-left	Arabic, Urdu and Persian captions render right to left.
Switch language	From the options sheet, without re-scanning.
Mute and Skip to live	Mute the voice and read only; or drop any backlog and jump to what is being said now.
Honest delay	Shows how far behind live it is, including audio still queued on the phone.
Catches up gracefully	If it falls behind, it speeds speech up slightly first; if it must drop queued sentences, it never cuts one mid-word and says how many it skipped.
Stays awake	Keeps the screen from sleeping while listening.
Recovers on its own	Reconnects automatically if the WiFi drops; resumes when the meeting restarts.

Feature	Detail
Audio paused by the phone	If the phone suspends audio (a call, an alarm, a screen lock), a "Tap to resume audio" button appears and works on iOS and Android.
Meeting state	Says plainly when the meeting has not started yet, or has ended – calmly, not as an error.
Light and dark	Follows the phone's setting, or a manual toggle.
Remembers your language	The next time you scan, your language is pre-selected.
Nothing stored	Captions exist on the phone only while the page is open, and clear when the meeting ends.

For the room – the room display

Feature	Detail
Join card	Large QR code, the address, the PIN, and the languages running. Readable from across the room.
Live transcript (optional)	Original and translation side by side, time-stamped, with a choice of translation language. Off by default; shows only what is said after it is turned on; clears when the meeting ends.
Listener count and clock	How many people are listening.
Light and dark	Brand yellow or black.
Explains itself	If opened without the PIN it needs, it says so instead of sitting empty.

For the operator – the desktop application

Starting up

- **Guided start-up** – "Getting the interpreter ready": speech-to-text, translator and voices each show their progress. About 45 seconds; nothing to press.
- **Set the languages** – the first screen after start-up asks what the room speaks and what it needs to hear.
- **Single instance** – opening it twice brings the existing window forward.

- **Runs in the tray** – closing the window keeps interpretation running; the tray icon shows state by colour (green live, amber starting, red needs attention).

Home – "Ready for the meeting"

- **Seven readiness indicators** – hearing the meeting audio · sound reaching the interpreter · detecting speech · engines loaded · translations arriving · languages set · phones able to connect. Hover any for the detail and the fix.
- **Languages in place** – tick source and target languages directly; changes apply immediately everywhere they are shown.
- **Start meeting / End meeting** – nothing is interpreted outside a meeting.
- **Open room display** – opens it with the PIN it needs.
- **Send test phrase** – speaks a phrase to every connected headset.
- **Who is listening** – listeners per language.

Meeting live screen

- The join card for the operator's own screen, the meeting timer and listener count, and one-click access to the room display, transcript and End meeting.

Live – measurement

- End-to-end latency (median and 90th percentile), listeners, sentences delivered.
- Per-stage timings: recognition, translation, synthesis.
- **Four separate loss counters**, each pointing at a different fix: **Dropped** (system fell behind) · **Not transcribed** (heard, nothing recognised) · **Not translated** · **Never spoken**.
- Counters reset at each meeting start so they describe this meeting.

Languages – automatic detection, expected languages, fixed source, up to four targets, Chinese caption script.

Audio – what to listen to: everything the PC plays, a microphone or line-in, or one application only; device, gain and a live level.

Settings

- **Profile** – on-premise or cloud.
- **Models** – replace the translation model (upload or URL); choose the voice engine per language; see the installed language packs.

- **Glossary** – fixed names and terms per language.
- **Cloud keys** – providers and credentials for the optional cloud profile.
- **Access** – listener PIN (with random generation) and administrator password.
- **Privacy** – transcript retention, export, text in logs – all off by default.
- **Tuning** – sentence segmentation and engine parameters.
- **Diagnostics** – send a test phrase; test a translation into all targets.
- **Licences** – terms of use and acknowledgements.

Built for non-technical operators

- Every problem is a sentence with the fix and a link to the right page – never an error code.
- "Saved" and "saved, but needs an engine reload" look different; a reload is an explicit action with its one-minute cost stated.
- Light and dark themes.

For IT and security

- Offline installer with every model and voice included – no download at install or run time.
- One network port (default 8760); every internal service bound to the PC only.
- No outbound connections – measured (see the Network Behaviour Statement).
- Administrator password generated at install; credentials kept apart from the shareable configuration file.
- No audio ever written to disk; transcripts off by default; logs record sentence lengths, not words.
- Built-in diagnostics and rotating logs for support.