

VoxSTAGE + SelfMix

speak the mix. control the stage.

User Manual

Voice-controlled digital-console bridge for live sound engineers, and SelfMix — the musician self-monitor add-on.

Covers installation, console connection, the engineer panel, the full voice-command set, the SelfMix musician app, console compatibility, licensing and troubleshooting.

Great sound is emotion. Precise sound is science.

Version	Manual v1.0 — for VoxSTAGE v1.2.x
Source	KILAVUZ_KAYNAK.md v1.1.1 (17 Sep 2026)
Publisher	Yilmaz Yenyol — yilmazyenyol.com

Contents

- 1** Product overview
- 2** Installation & first launch (Mac)
- 3** Connecting to the console
- 4** The Engineer Panel
- 5** Voice commands — engineer
- 6** SelfMix — musician self-mix
- 7** Console compatibility
- 8** Licensing, packages & cloud
- 9** Troubleshooting
- 10** What this manual will not claim


1. Product overview

VoxSTAGE is a Mac desktop bridge that lets a live-sound engineer run a digital console **by voice and from a panel**. It runs on a laptop, connects to the console over the network, and opens a control panel in any phone or tablet browser on the same network.

SelfMix is an add-on inside VoxSTAGE (separate license) that lets each musician adjust their **own in-ear / monitor mix** from their own phone, by voice or by touch. No app install: scan a QR code and it opens.

Languages. Engineer commands are understood in 9 languages (EN, TR, DE, FR, ES, IT, PT, NL, RU); spoken confirmations come back in 11 (Piper voices: EN, TR, DE, ES, FR, IT, PT, NL, RU, JA, KO). On SelfMix, each musician picks their own language on their own phone — one musician can run in German while another runs in Turkish on the same stage; confirmations come back in whichever language that musician chose.

Speech engine priority. With an active license, speech recognition runs through VoxCloud first; if there is no internet, it falls back to the local offline engine (whisper.cpp). The recognized text is always tried against the built-in grammar first — instant, fully offline. Only if the grammar does not understand a sentence, and a license is active, does that sentence go to VoxCloud, where a model picks exactly one function to call (through the same confirmation/safety path as any other command). This is what lets free-form phrasing work ("turn the lead vocal down a bit") without leaving the microphone dead while the ~480 MB offline model downloads on first launch. No API key ever sits on the customer's machine — the license file is the credential.



The VoxSTAGE engineer panel — console selected, DEMO mode on, console connected and powered, voice engine ready (green status row), and the SelfMix QR generator at the bottom.

2. Installation & first launch (Mac)

1. Drag **VoxSTAGE-<version>-apple-silicon.dmg** (M1/M2/M3/M4) or **-intel.dmg** into Applications. The app is signed and notarized (Developer ID: Yilmaz Yenyol). A Windows build does not exist yet.
2. The first launch takes about 9 minutes: the offline engine files download in the background — the whisper **ggml-small.bin** model (481 MB) plus a 63 MB Piper voice for the reply language. If the license is active, voice already works immediately through VoxCloud; without a license, only typed commands work until the download finishes. The panel log shows "first run: downloading..." then "offline speech engine ready". The panel's ENGINE field reads **LOCAL** (whisper), **CLOUD** (VoxCloud), or **TEXT** (no engine ready yet).
3. Settings and logs live in ~/Library/Application Support/VoxSTAGE/ (config.yaml, license.txt, certs/, vendor/models, vendor/voices). Updates install on top; this folder is preserved.
4. Open the panel at <https://localhost:8443/panel> (the VoxSTAGE window itself), or scan the QR code from a phone on the same network. The certificate is self-signed, so a first-run browser warning is normal — choose Advanced → continue anyway.
5. License: click the badge at the bottom of the panel (**TRIAL 14d** / **LICENSED** / **EXPIRED**) and paste the contents of your license.txt. The 14-day trial is fully featured, SelfMix included. A license is tied to

the person, not to one machine.

6. Updates: when a new version ships, an orange **UPDATE v1.x** badge appears on the panel (red = mandatory). Click it, the DMG downloads, install over the old copy. There is no silent auto-install.

3. Connecting to the console

In the panel window, pick the console **profile** from the model list, type the console's **IP**, and press the power button. Auto-scan finds a matching console automatically on the same network (Yamaha, Behringer, Midas families supported — see the compatibility table in §7).

Same-network rule. The laptop, the console and every musician's phone must be on the same LAN/Wi-Fi. If the network changes, the QR codes on the panel refresh to the new addresses.

DEMO mode lets you try the entire interface without a real console attached — a simulated desk stands in, which is exactly how the screenshots in this manual were produced. It is meant for training and sales demos.

Allen & Heath dLive needs a little extra setup in config.yaml: `midi_base_channel` (read it off the console at Utility > Control > MIDI), `aux_map` (e.g. ["S1","S2","M1"] — the order of stereo/mono auxes; this is filled in automatically if the console answers a name query, otherwise set it by hand), and `socket_map` (channel → preamp socket, 1:1 by default).


4. The Engineer Panel

The panel has two faces: the **VoxSTAGE window** itself (console/IP/profile, voice-engine status, network addresses + QR codes, license badge, version/UPDATE badge, SelfMix QR generator, DEMO switch, QUIT — this is the screen shown in §1) and the **mixer screen** opened from a phone or tablet browser, described tab by tab below.

Tab / card	What it does
Channel list (blocks of 12)	pick a channel, see its name, mute state; folders group bus / matrix / DCA / main
SELECT (channel strip)	FADER, MUTE, analog GAIN (head-amp; on an RF channel this reaches the Shure receiver instead), DIGITAL GAIN (if the desk has it), PHANTOM/PHASE (if supported), PAN on the stereo bus, MASTER ROUTING (ST/MONO, where available), INPUT PATCH (Dante/Omni, console-dependent), COMP/GATE threshold + on/off, and the channel's mix sends (level, on/off, pre/post)
SENDS	every mix / matrix send for the selected channel, as sliders with level, ON/OFF and PRE/POST
OUTPUT	master fader, on/off and name for every mix / matrix / DCA / main output, with routing and OSC-assign
OSC	the console's test-tone oscillator: pink / white / sine / burst, level, target bus
LOG	what was heard, what was sent, what came back — the developer/diagnostic log
Hold-to-talk	press and hold, speak, release — the command runs, with a spoken and written confirmation
CONFIRM / CANCEL bar	appears for bulk or risky actions before they run
UNDO	reverts the last action (up to 50 steps)

4.1 Channel strip — SELECT tab


This is the main working view once a channel is selected: fader value, analog and digital gain (with quick $\pm 1/\pm 6$ buttons and a rotary readout), pan on the stereo bus with quick L/C/R and ± 10 buttons, master routing (ST/MONO), Dante/Omni input patch, and the channel's compressor and gate threshold with their own on/off switches.



SELECT tab for channel 12 — analog gain +42.0 dB, digital gain +6.0 dB, pan L8, routed to ST, patched to OMNI 12, comp threshold -18.0 dB and gate threshold -45.0 dB both engaged.

4.2 Channel list — INPUTS tab


Lists every named input in blocks of 12, with a live search box and collapsible MIX SENDS / MATRIX SENDS shortcuts above it. Blocks with no named channels show "0 named" and stay collapsed.



INPUTS tab, channels 1-12 of a drum-heavy demo session — Kick In/Out, Snare T/B, Hi Hat, Toms, overheads, Bass DI/Mic.

4.3 Sends — SENDS tab


For the selected channel, every mix bus appears as its own row: a level slider, an ON/OFF toggle and a PRE/POST tap switch.



SENDS tab — eight monitor mixes (Drummer Mon through FOH Aux) with individual level, on/off and pre/post state.

4.4 Outputs — OUTPUT tab


Grouped by MAIN/MONO, MIX BUSES and MATRIX. Each row shows the output's routing, level (with quick +/- buttons), on/off, and whether the oscillator is currently assigned to it.



OUTPUT tab — main L/R and mono fill, all eight mix-bus masters, and two matrix outputs, each with routing and level.

4.5 Oscillator — OSC tab

The console's built-in test-tone generator: on/off, level, waveform (pink / white / sine / burst) and, for sine, the frequency. Bus assignment for the oscillator's output is done from the OUTPUT tab's own OSC-assign switches; this tab controls the tone itself.



OSC tab — oscillator off, level -20, four waveform choices, sine-frequency stepper.

4.6 Monitor Matrix

A separate live view (opened from the panel, or directly at /matrix) showing every channel against every mix as a grid: level bar and pan readout in each cell, and each mix's own master level in its column header. Useful for scanning the whole monitor picture at a glance instead of tab-by-tab.

Kanal \ Mix	Drummer Mon -2 dB	Guitar Mon -4 dB	Bass Mon 0 dB	Vocal IEM +1 dB	Side Fill +2 dB
Kick In ch 1	-10 C	-3 R20	-24 R40	-13 L40	+2 L20
Kick Out ch 2	-3 R20	-30 R40	-17 L40	-4 L20	+9 C
Snare T ch 3	+4 R40	-23 L40	-10 L20	+3 C	-24 R20
Hi Hat ch 4	-29 L40	-16 L20	-3 C	-30 R20	-17 R40
OH L ch 5	-32 L20	-9 C	+4 R20	-23 R40	-10 L40
OH R ch 6	-15 C	-2 R20	-29 R40	-16 L40	-3 L20
Bass DI ch 9	+6 L40	-21 L20	-8 C	+5 R20	-22 R40
Gtr Elec ch 11	-20 C	-7 R20	+6 R40	-21 L40	-8 L20
Lead Vocal ch 12	-13 R20	0 R40	-27 L40	-14 L20	-1 C
BGV 1 ch 13	-6 R40	-7 L40	-20 L20	-7 C	+6 R20

Monitor Matrix — ten channels across five monitor mixes, with per-cell level and pan.

5. Voice commands — engineer

The grammar's core is English; other languages are translated to the same English command set through a lexicon. A channel is addressed **by number or by its name on the console** ("channel 5", "kick", "lead vocal"). Numbers may be spoken as digits or words. An unrecognized command returns "did not understand" — it never guesses and picks the wrong channel.

5.1 Channel commands

Command (EN)	Example (TR)	Result
channel 5 fader minus 3 / down 3 / up 3	kanal 5 fader eksi 3	relative fader move (± 12 dB cap per command)
channel 5 fader to minus 10	kanal 5 fader eksi 10	absolute/relative fader (TR/DE "to/auf" may read as relative — prefer "unity/zero" and $\pm$ phrasing)
channel 5 unity / zero	kanal 5 sıfır	sets 0 dB
mute channel 5 / unmute 5 / open 5	kanal 5'i sustur / aç	mute / unmute
mute channels 2 through 12	2 ile 12 arasını sustur	bulk mute (an explicit range runs without confirmation)
channel 5 gain up 3 / gain to 30	kanal 5 kazanç artı 3	analog head-amp gain (on a Shure/RF channel this reaches the receiver; a digital-source channel gets a warning)
channel 5 digital gain minus 6	kanal 5 dijital gain eksi 6	digital gain, where the console offers it
channel 5 phantom on/off · phase invert/normal	kanal 5 fantom aç/kapat, faz ters	48V, polarity, where supported
channel 5 comp threshold minus 20 · gate threshold ...	kanal 5 comp eşik eksi 20	comp/gate threshold
channel 5 comp on/off · gate on/off	kanal 5 comp kapat	dynamics on/off, where present
select channel 5 / show kick	kanal 5'i seç, kick'i göster	brings the panel to that channel

5.2 Send / monitor commands

Command	Result
channel 5 mix 3 minus 6 · channel 5 to mix 3 at minus 6	send level (absolute/relative)
channel 5 mix 3 on/off	send on/off
channel 5 mix 3 pre/post · channels 1 through 12 mix 1 pre · mix 1 pre for channels 1-12	pre/post fader tap, single or bulk (an explicit range runs without confirmation; "all" asks for one)
mix 4 is fred's wedge he is on channel 12 · set up mix 6 for the drummer on channels 1 through 8 · monitor 3 own channel 12	monitor setup shortcut (Yılmaz's house rule): pre-fades that channel range into the mix and opens the owner's own channel in it; asks for confirmation

Command	Result
mix 3 fader minus 2 · matrix 1 fader to minus 6 · main fader zero · dca 2 fader up 3	output masters (mix / matrix / main / DCA)
mix 3 on/off · main off · dca 3 off	output mute
mix 3 name "Fred wedge"	output naming — backed up to the cloud with a license; always works from the panel

5.3 Oscillator, scenes, naming, patching

Command	Result
oscillator on/off · osc level minus 20 · osc to mix 5 / pink noise on mix 5	oscillator (waveform/frequency are panel-only; voice control is limited)
recall scene 12	scene recall — cloud-backed with a license; always works from the panel
channel 5 name "BD IN"	channel naming (8 characters on Yamaha, 12 on X32, 16 on WING) — cloud-backed; always works from the panel
channel 1 patch omni 3 · patch channel 4 to none	input patching — see §7: CL/QL cannot write Dante (Omni/None only); unverified on DM7/X32
undo · confirm / "yes, ok" · cancel / "no"	undo / confirm / cancel

Naming a channel by voice

Works fully offline through the grammar. Examples: "channel 33 kick", "name channel 33 kick", "rename 12 to lead vox", "call channel 4 e gt1", "channel 7 named bass" (Turkish: "kanal 33 kick", "kanal 12 lead vokal diye adlandır"; German: "kanal 33 heisst kick"). Spoken names are converted to Yilmaz's own channel-naming standard — for example kick out/in become BD OT / BD IN, snare top/bottom become SN TP / SN BT, lead vocal becomes LEAD, and so on for guitars, keys, percussion and vocals. A name outside the standard is written exactly as spoken (capitalized) — nothing is invented. Saying "double"/"stereo"/"çift X" writes a stereo pair; "keyboard" alone still writes KEY L, KEY R.

A single sentence can name a run of consecutive channels — "channel 12 hi hat and 3 tom and 2 overhead" writes, starting at 12: HI HAT, TM 1, TM 2, TM 3, OH L, OH R. If no starting channel is given, it starts at the currently selected channel and reads the list back for confirmation first. There are also ready-made drum templates ("4-channel drums", "6-channel drums", "8-channel standard drums", "10-channel 3-tom drums", "12-channel drums", "14-channel drums") that fill in a sensible standard layout; the template is always read back before it is written, and a single "undo" reverts the whole set.

Safety behavior

- A single command moves at most **±12 dB**; a channel or mix that does not exist returns an explicit error, never a guess.
- Anything that touches **every** channel ("mute everything", "mix 1 all pre") asks for confirmation first; an explicit range ("2 through 12") runs without asking.
- If the console swallows a write (gain, patch), the value is read back and the mismatch is **stated out loud** — there is no silent false "done".

- Every write goes onto the UNDO stack (50 steps).
- Gain ownership: Rio/Omni inputs → the console's own head-amp; Shure RF → the receiver (TCP 2202); an unknown DVS/MTX/Dante source has no analog gain at all (digital gain is suggested instead). The panel shows this under the channel.
- **Scene recall by voice now always asks for confirmation** ("recall scene 7 — confirm?") and cannot be undone once run; `safety.voice_scene_recall: off` in `config.yaml` disables voice scene recall entirely (recommended during a show — recall from the desk instead). SelfMix (a musician's phone) can never touch scene recall under any circumstance.

6. SelfMix — musician self-mix

6.1 Setup flow

- 1.** In the panel, under **SelfMix > generate QR**: enter the musician's name, their aux/mix number, their channel range (a whitelist), and optionally a role (bass / guitar / keys / brass / bvox / lead / perc — used for the Premix preset). This step is PIN-protected. On a brand-new install there is no PIN yet — type 4–8 digits into the panel's "No QR PIN yet" field and press SET PIN (this writes to config.yaml). QR generation is disabled until a PIN is set, and changing an existing PIN requires the old one.
- 2.** The musician scans the QR with their phone camera, which opens SelfMix in the browser (a PWA — it can be added to the home screen, no app-store install). The token is locked to that one musician and that one aux; it cannot reach any other mix.
- 3.** At the top: a language picker (their own choice), a lock button  (guards against accidental taps), and the **MY MIX** meter for their own aux output, where the console provides it.
- 4.** Each channel row shows an instrument icon, its name, a **level bar** (drag to adjust, 0.5 dB steps) with the incoming signal meter shown behind it, and a **pan bar** (± 6 -degree center detent).
- 5. Hold & speak:** press and hold the microphone button, talk, release. The confirmation comes back spoken and written, in whichever language that musician picked.



The screen a musician sees after scanning their QR code. Every row is one channel from their whitelist: an instrument icon (matched automatically from the console's channel name), the live input meter as a green-to-red bar, a lighter overlay marking the current fader position, and a pan indicator underneath. Tapping **MY MIXES 1/2/3** recalls a saved mix; holding one for about half a second saves the current mix into that slot (see §6.6).

6.2 Voice command types

Type	Examples	Notes
level	"bring the kick up a bit", "guitar's too loud", "can't hear the keys", "snare louder", "turn everything down", "bring the master up"	direction: louder/quieter; intensity: normal / more ("much more")
pan	"guitar to the left", "a bit to the right", "center it"	full / half / center
premix	"premix", "set up the premix"	snaps the whole mix to role-based preset values
self	"open me up"	opens the musician's own instrument (target ≈ -6 dB)
group	"@drums", "@vocals", "@band" (everything but vocals), or a musician-defined LINK (e.g. "the guitars")	a LINK is a channel group the musician defines on their own phone
confirm/cancel	"yes/ok", "no/cancel"	answers a pending question

Type	Examples	Notes
save / recall	"save my mix", "save my second setting", "recall my mix", "recall my third setting", "my old setting"	the musician's own mix saved to one of 3 slots (also the on-screen MY MIXES 1-2-3 buttons — tap to recall, hold to save); stored server-side, so it survives a phone change; a channel that's no longer permitted is silently skipped on recall
alias	a musician can rename a channel on their own screen (e.g. "acoustic" for A GT2)	Turkish suffixes/softening are recognized ("turn up the acoustic")

6.3 Safety rules (fixed, not configurable per show)

- A musician can only touch **their own aux** and the **channel list given in their QR**.
- Step size follows a logarithmic curve: 3-6 dB at -24 dB, 1.5-3 dB at -10 dB, 1-2 dB at 0 dB, 0.8-1.5 dB at +10 dB. Even a spoken absolute dB value never jumps past the band ceiling for one command. "Open it up" on a muted channel goes to its premix target (or -30 dB if none).
- Floor is -40 dB (nothing meaningful below it). Pan runs -100...+100. A touch gesture goes through exactly the same safety gate as a spoken command.
- Console-side work — patching, naming, scene recall, the oscillator — is **not available** in SelfMix; that stays the engineer's job.

6.4 Premix (visual balance / LPR)

A role-by-source matrix decides the premix: a musician's own instrument sits at -6 dB, everything unlisted sits at -18 dB, and some cells are muted outright. The guitar role also distinguishes amped vs. direct sources (config: `amped_channels`). The musician can override which role's balance they use from SelfMix's own settings (⚙ → "Premix profile") — the QR's role is only the default, so a drummer can choose to hear a "bassist's balance" instead. Changing the profile only changes the premix table; the aux and channel permissions still come from the token and never change. The drummer profile is currently simple (own kit at -6, everything else at -18) pending a fuller set of values.

6.5 Saved mixes (MY MIXES 1-2-3)

Under the MY MIX meter on the phone: three slots, **1 • 2 • 3**. Tap to recall a slot; hold for about 0.6 s to save the current mix into it (or use the voice phrasing in §6.2). The save lives on the server under that musician's identity, so it survives a phone swap or a cleared browser. What is saved is exactly the levels and pans on that musician's own aux, including which channels are muted.

7. Console compatibility

"documented" means it follows the manufacturer's protocol spec but has not been tried on real hardware yet; "verified" means it has been tested at a real desk. A checkmark only ever means what the table below says — nothing more is promised anywhere else in this manual.

Feature	Yamaha CL/QL (verified: CL5)	Yamaha DM7/DM3/ TF/Rivage (doc)	X32/M32 (verified: X32C)	Behringer WING (doc)	A&H dLive (doc, sim)
fader / mute / name	✓	✓	✓	✓	✓
analog gain / 48V	✓ / ✗ (not in RCP)	✓ / DM7 ✓	✓ / ✓	✓ / ✓	✓ / ✓ (per socket)
digital gain	✗	✓	✓ (trim)	✓	✗
phase	✗	DM7 ✓	✓	✓	✗
send level / on / pre-post	✓ / ✓ / ✓	✓	✓ / ✓ / ✓	✓	✓ / ✓ / ✗
send pan (stereo aux) → SelfMix pan	✓	✓	✗ (not yet)	✗ (not yet)	✗ (not in protocol)
channel pan (stereo bus)	✓	✓	✗ (not yet)	✗	✗
ST/MONO assign	✗ (not in RCP)	DM7 ✓ (doc)	✗ (not yet)	✗	✓ (main assign)
comp/gate threshold + on/off	✓	✓	✓ threshold	✓	✗
input meters → SelfMix	✗ (RCP doesn't give it); MY MIX aux meter ✓	doc	✗ (not yet)	✗	✗ (not in protocol)
oscillator	✓	✓	✓	✓	✗
scenes	✓	✓	✓	✓	✓ (1-500)
input patch	Omni/None ✓, Dante ✗ (firmware swallows it)	doc (untried)	✗ (not yet)	✗	✗
Shure RF gain (receiver)	✓	✓	✓	✓	✓

A&H SQ / Avantis support is planned but not yet available. Notes: "not yet" means the protocol supports it but the driver doesn't write it yet; "not in protocol" means the manufacturer's own protocol has no path for it at all. On dLive, SelfMix works for level + mute + send, but the pan bar and meters stay empty. On CL/QL, digital gain, ST/MONO, phase, 48V and input meters are absent because Yamaha's RCP protocol itself doesn't expose them — the separate (proprietary) CL Editor protocol does, and a future MIDI bridge may add it.

8. Licensing, packages & cloud

The license file is an Ed25519-signed text block (starting VOXSTAGE-LICENSE-1) and verifies fully offline. A time-limited license carries an expires date; once it passes, the badge reads EXPIRED and

the cloud responds 401.

Pricing model (pass-based): VoxSTAGE \$25/day, SelfMix +\$12/day on top, a show pass (VoxSTAGE + SelfMix, 8 seats, 24h) \$29, a 3-day pass (12 seats) \$59, a 30-day tour pass (16 seats) \$99; roughly 40% regional discount for TR/MX/BR via FastSpring. An annual or subscription plan does not exist yet.

Cloud (VoxCloud). Automatic for a licensed user; quota is 120 requests/min and 4,000/day per license. Without internet, the app falls back to the offline engine — the customer never enters an API key of their own.

Privacy. Audio clips are sent to the server only for the duration of a command and are not retained; logs carry an order number, not audio.

9. Troubleshooting

Symptom	Cause / fix
Panel stuck on "starting..." / channel range won't resolve	console isn't connected yet (the app fails closed by design) — set the IP and press ON; confirm same network
No meters in SelfMix	the console doesn't provide an input meter (CL/QL, dLive) — expected; MY MIX still works where the console supports it
"did not accept ... still reports"	the console swallowed the write (gain on a digital source, or a patch attempt) — the panel explains what happened
Weak recognition in German / accented speech	check the locked reply language; check whether cloud mode is active (LICENSED); add a lexicon/alias entry if a word is missing
"BRIDGE IS AN OLD BUILD"	quit and reopen the app — a leftover old process is still holding the port
Certificate warning	self-signed cert — choose Advanced → continue. For a permanent fix, run "mkcert -install"
No voice / microphone didn't respond	if the panel's engine badge reads TEXT, whisper hasn't finished downloading yet (first-launch download); otherwise check the mic permission
Panel shows "VoxCloud UNREACHABLE"	the cloud server is down, or the network is blocking it — the only fallback at that moment is the local model (run setup_local_voice.sh, or copy the model file in by hand)
Model download says "No route to host"	installs after 18 Sep 2026 already ship the model, so this only affects older packages: the network is blocking huggingface.co (or has no IPv6 route). Try a hotspot, point VOXSTAGE_MODEL_BASE at your own mirror, or drop ggml-small.bin directly into ~/Library/Application Support/VoxSTAGE/vendor/models/
SelfMix channel moves on the phone but not on the desk	the same channel name exists twice on the desk (e.g. BD on both channel 1 and 33) — regenerate the QR with an explicit channel range (ch1-ch2); resolution then searches that range first. A QR generated without a range writes to whichever matching channel it finds first
Panel shows SelfMix "musician speech: OFF"	that musician's speech can't be resolved: add a license (VoxCloud) or run bash setup_local_voice.sh; if the network blocks huggingface.co, a license is the only remaining path
Phone says "no sound reached the mix"	the browser didn't grant the microphone, or the tab went to the background — grant the permission and press the button once more
Can't generate a SelfMix QR ("admin_pin is not set")	in the panel, type 4-8 digits into SelfMix's "No QR PIN yet" field and press SET PIN

10. What this manual will not claim

- "Every feature on every console" — the table in §7 is the actual, current truth, nothing more.
- There is no automatic updater, no Windows build, and no iPad app.
- Pan and meters do not work on dLive; Dante patching, digital gain, ST/MONO assign and input meters do not work on CL/QL.

- There is no "100% recognition" claim — the cloud engine runs at roughly 90%+ on a noisy stage; the offline grammar covers a deliberately limited command set.

VoxSTAGE + SelfMix — Yilmaz Yenyol · yilmazyenyol.com/voxstage · manuals: yilmazyenyol.com/manuals