

stems

FIELD GUIDE · STABLE AUDIO 3 · V1

you play your track; stems writes the band around it — drums, bass, keys, pad, lead — in whole bars at your tempo and in your key, on this mac.


this guide covers setting up, the plug-in, and how to talk to the model. it is short on purpose: most of it is in your ears.



left: a take is ready — your track above (120 bpm · 16 bars · a minor), the part it wrote below, take 2 of 3 in a minor. right: writing, step 3 of 8 — the hatch is where the part will land.

install, open stems, play.

what you need	an apple silicon mac (m1 or later), macos 14 or newer, about 8 gb free, and logic – or any host that loads audio unit effects.
install	open stems.pkg. it puts stems in applications and opens it when it's done.
open stems	the stems app sets the model up once: stable audio 3 medium (~7 gb) and a small brain that keeps it ready from login. paste your free key for the fast download (or skip – it comes from hugging face instead, checked the same way). stop it any time: it picks up where it left off.
already have stems loop?	then stems uses its model as it is – nothing to download, about a minute.
the plug-in	in logic: a track > audio fx > audio units > s7ems > stems.
the first-run card	four steps: the brain, write, feel, takes. if the brain is off it says whether it was never set up or just isn't running, and offers open stems. settings (the slider key, top right) > show the intro brings it back.
your track	put stems on the track you want a band for and press play in logic – stems listens while it plays (listen again re-captures, up to two minutes). or load a file, or drop one on the window.



the stems app. left: setting up – each file with its size and percent, from the stems cdn. right: ready, with the path to the plug-in and the brain's state.

one key. two lanes.

top to bottom it reads like the stems loop stage: what you played, what the model wrote, and the one key that always does the next thing.

status bar	the wordmark, the session tempo, and the brain: its state in words and a lamp – steady when ready, blinking while it works. click it to wake the brain. the slider key next to it is settings.
your track	the captured audio: length, bpm, bars and the key stems hears in it, and a three-band waveform – ash is the low end, amber the mids, paper the highs. play · listen · load · clear.
the part	what the model wrote, from your track to the part. machine-made audio carries red and cyan ghosts at its edges, so you always know what's real. the line under the title: bars, seed, take, and the key the take landed in.
the big key	write → writing (the key fills as it renders, with the step it's on) → play it. before you press, it says what it will write – and, for long parts, about how long it takes. shift-click cancels while writing, and re-rolls once a take is ready.
takes	‹ take 2/3 · d minor › flips between every take of this part; re-roll writes another with the same words and a new seed. nothing is thrown away until you close the session.
parts	every part you've written, as thin lanes: play, mute, drag it onto a track in logic, or remove it. a part drags out as written – 128 bars come out as 128 bars.

say it, then write.

part	drums · bass · keys · pad · lead. there's no voice: the model makes vocal-like textures, never words.
style	one tap sets the session's genre: indie rock, hip hop, house, techno, lo-fi, funk, ambient, trap, disco, dub. every part you write afterwards shares it, so they sound like one band.
prompt	a few words about this part – what it plays and how it sounds. return writes.
the line under it	the exact words the model receives, built for you, e.g. TrackType: Instrument, Bass groove loop, solo bass, round fingered bass, house style. Genre: House. Instruments: Bass. Key: A minor. BPM: 120
bars	4 · 8 · 16 · 32 · 64 · 128. the part comes back in whole bars at logic's tempo: stems asks for one bar more, finds the first hit, cuts so it starts on the downbeat, and trims to the exact length.
long parts	the model writes up to about 6:20 of music. 128 bars fit from about 82 bpm (in 4/4); slower, stems writes the most whole bars that fit and says so first – 128 bars is 7:19 of music at 70 bpm – writing 109. a part longer than your track plays whole, with your track looping under it.
feel	off · low · mid · high. off writes from words alone. turn it up and the part starts from your track itself, so it sits on your groove – high follows it closest.
takes	1 · 2 · 3 written per press. 8 bars take about 6 seconds each, so three is a good habit; 128 bars take about 50 seconds, so 64 and 128 drop to one take at a time (turn it back up if you like).
key	auto · any of the 24 keys · off. auto sends the key stems hears in your track; pick one to override it, or turn it off. drums never get one.
steps	8 · 6 · 4. the model is built for 8; fewer is faster and rougher.

short, concrete, sound words.


what it hears	stable audio 3 reads genre, instrument and rhythm, then mood and energy, then how it's recorded. stems already supplies the part, the genre, the key and the tempo – your words are the middle: the playing and the sound.
good words	dusty breakbeat with ghost notes · round fingered bass, sparse, laid back · wurlitzer stabs, offbeat, warm tape · wide analog pad, slow swells · plucky arp, bright, short delay
keep it short	under a dozen words. the whole prompt is kept under 45 words; long prompts drift. your words give way first – the key and tempo always stay.
the key is a hint	the model has no key input, but Key: D minor. in the prompt steers it: in our tests it landed on the asked key or a fifth away every time, and exactly on it about a third of the time. so: it lands on your key or a fifth away; re-roll if it clashes. every take shows the key it landed in, so you can keep the one that sits.
don't ask for	singing, lyrics or vocals (it can't), or 'no ...' (negative words don't do anything on this model).
same words, new take	re-roll keeps your words and changes the seed. takes of one prompt are often close cousins – change a word to get somewhere new.

it runs on this mac.

the model	stable audio 3 medium, stability ai's open-weight model, in its fast 8-step form, running on apple's mlx. 44.1 khz stereo out, never clipped. powered by stability ai.
how it runs	a small server on 127.0.0.1:7432, installed by the stems app in ~/library/application support/stems and started from login. it's lean: under 1 gb while idle, about 4.3 gb while it writes.
speed	on an m3 pro: 8 bars at 120 bpm in about 6 seconds; 128 bars (4:16 of music) in about 50; a part that follows your track takes a little longer.
the states	not set up (open stems) · off (open stems, it starts it) · asleep (wakes on your first write, or click the status) · waking · ready · writing.
the stems app	open it any time: the brain's state, start it, repair (checks everything, fetches only what's gone), credits and the model's licenses, uninstall.
nothing leaves	what you play and what stems writes stay on this mac. the only downloads are the model, once, and its runtime.

fixes.

the brain isn't running	click open stems on the card (or open stems from applications): it starts the brain, and from then on it starts with your mac.
stems isn't in logic	quit and reopen logic so it rescans its plug-ins – it's under s7ems › stems.
the part is off the beat	check logic's tempo matches the song – stems writes at the tempo logic reports. then re-roll: the downbeat finder works best on parts with a clear first hit.
the part is in another key	re-roll – each take says its key; keep the one that matches. if auto heard your track wrong, pick the key yourself on the key stepper.
it sounds like the whole band	say what it plays alone – solo bass, drum loop – and keep feel low; high feel pulls your whole track in.
setup stopped	open stems and set up again: it resumes, and nothing is downloaded twice. repair re-checks every file.
where the logs are	~/Library/Logs/stems/ (brain.log, setup.log) · the plug-in's own log: settings › reveal log.



the card with the brain not set up yet: what it means, and open stems.