AETHERWAVE STUDIO

Music DAW

User Manual

2026-11-14 · 11 chapters · aetherwavestudio.com

Contents

1 What is the Music DAW?

2 Getting started with the Music DAW

3 Settings and setup

4 Tracks and clips

5 Recording live audio (mic, guitar, keys)

6 The Library - samples, SFX, and AI generation

7 Mixing

8 Effects and the rack

9 Mastering

10 Export and render

11 The AI-Engineer

CHAPTER 1

What is the Music DAW?

> A multi-track audio workstation for arranging, recording, generating, mixing, and exporting music, right in your browser. Layer tracks, record live instruments, generate music and SFX to spec, add effects, ask the AI-Engineer to mix for you, and render a WAV or MP3.

What it is

The **Music DAW** (Digital Audio Workstation) is a full multi-track studio that runs entirely in your browser. It is where you take ideas, generated songs, live performances, and samples and turn them into a finished mix.

With the DAW you can:

- **Layer multiple tracks** on a timeline and trim, split, and rearrange clips
- **Record live audio** - plug in a guitar, bass, or keyboard through an interface, or use a mic for vocals
- **Generate music** directed to your project's tempo and key, and **generate sound effects** from a text description, both landing right on the timeline
- **Browse samples** from your own sample directory and a shared community library
- **Draw MIDI** in the piano roll and turn a melodic stem into MIDI notes
- **Adjust volume, pan, mute, and solo** per track, with meters and lockable faders, and draw volume and pan automation
- **Add effects** - EQ, compression, reverb, delay, distortion, and more, all undoable
- **Ask the AI-Engineer** to mix, add and dial in effects, draw automation, replace a section of a track with new AI audio, and batch-export sample packs
- **Export** the final mix as **WAV** (lossless) or **MP3** (compressed)

It runs on the Web Audio API, so there is no install, no plugin host, and no proprietary file format.

The DAW and Music Studio

Music Studio is built to generate a single finished AI song from a prompt. The DAW is built to **produce** - to combine, record, arrange, and mix. They complement each other:

You want to...	Use
Generate one finished AI song	Music Studio
Combine multiple songs into a layered mix	Music DAW

You want to…	Use
Record a live guitar or vocal over a track	Music DAW
Generate a bed at your project's tempo to build on	Music DAW
Add a sound effect or a sample to a mix	Music DAW
Re-arrange the structure of a song	Music DAW

Anything saved to your **Gallery** in Music Studio shows up in the DAW's library, ready to drag onto the timeline.

Who it is for

You do not need prior DAW experience. The defaults are sensible, and the most common workflow (drag in a track, add a reverb, export) takes about a minute. A first-run guide opens the first time you visit, and you can reopen it any time from the **Help** button in the header.

If you are coming from FL Studio, Logic, Ableton, or BandLab, the timeline plus tracks plus effects-rack model will feel familiar. If you have never used a DAW, start with the Getting Started guide.

What you can export

- **WAV** - lossless, 44.1 kHz / 16-bit stereo. Bigger file, used for archival, distribution masters, and re-importing into another DAW.
- **MP3** - compressed, smaller file. Used for sharing, streaming previews, and social posts.

Both export the full mix including all effects, automation, mute and solo state, and the master volume.

A note on desktop and accounts

- The DAW is **desktop-only**. Viewports under 900px width show a "desktop required" message with a link back to the site, because arranging and mixing need a keyboard and mouse.
- You can open and use the DAW without an account for local editing, but **generating music or SFX, browsing the community sample library, and saving a project** require a signed-in account and (for generation) credits.

What it is not

- It is not a mastering suite. Use the dedicated Mastering tool for a final polish pass after you export.

- Generated music in the DAW is tuned for raw material that fits your session (a foundation, an element, a scratch bed). For a finished standalone song, generate in Music Studio.

Where to go next

- New here? Start with Getting started
- Want to record a live take? See Recording live audio
- Need sound to build with? See The Library - samples, SFX, and AI generation
- Ready to arrange? Jump to Tracks and clips
- Want the DAW to do the work? See The AI-Engineer
- Ready to ship? See Export and render

CHAPTER 2

Getting started with the Music DAW

> A 5-minute walkthrough from "open the DAW" to "downloaded a finished WAV". Covers your first session, drag-from-gallery, transport, saving, and a single export.

Before you start

You need:

- An account (free tier is fine to try the DAW)
- At least one song in your **Gallery** - generate one in Music Studio if you have not yet
- A desktop browser. The DAW is desktop-only - viewports under 900px width show a "desktop only" message.

Open the DAW

Go to **Music DAW**. The page loads in three regions:

1. **Media library** on the left - your gallery, samples, SFX, and AI Music
2. **Timeline** in the middle - tracks and clips arranged horizontally over time
3. **Effects rack** on the right - effects you have added to the master bus or a track

Below the timeline there is a **mixer panel** with channel strips for each track. Above it, a header with transport controls (play, pause, record, loop), the time display, and the export button.

Your first 60 seconds

1. Click the **My Gallery** tab in the library if it is not already active. Your generated songs appear as a list.
2. **Drag any track** from the library onto an empty track lane in the timeline. A clip with a waveform appears.
3. Click the **Play** button at the top. You should hear the track and see the playhead move across the timeline.
4. Click **Pause** or press space to stop.
5. In the top-right, click **Export WAV**. A file downloads to your computer within a few seconds. Open it - that is your mix.

You have just used the DAW end-to-end. Everything else is iteration on this loop.

Add a second track

1. Drag another song from the library onto a different track lane (or onto an empty area below the existing tracks - a new track is created).
2. Drag the clip horizontally to align it with the first track. The grid snap will lock it to the beat.
3. Use the **volume slider** on each track header to balance the levels.
4. Press play - you are now hearing both tracks layered.

If one track is louder than the other, that is normal for AI-generated audio. Pull its volume down. Better still, add a **Compressor** effect (see Effects and rack).

Record or generate your own sound

You are not limited to your gallery. From the same session you can:

- **Record a live take** - plug in a guitar or use a mic and hit **Record** in the transport. See Recording live audio.
- **Generate a bed or an element** - the **AI Music** tab generates music matched to your project tempo, and the **SFX** tab generates sound effects from a description. See The Library.

A first-run guide opens the first time you open the DAW. Reopen it any time from the **Help** button in the header.

Save your work

Click the **Save** button (floppy disk icon in the header). The DAW POSTs your project state to the server and shows a confirmation toast. The project now lives in your account - click the **Projects** folder icon to see all your saved DAWs.

You are covered between saves too:

- **Auto-save picks up every kind of edit.** Timeline edits, clip drags and trims, effect and knob changes, piano-roll note edits, AI-Engineer actions, and generated audio all mark the project unsaved and schedule a save a few seconds later. The status readout next to Save tells you where things stand.
- **A local backup runs every 30 seconds** to your browser's storage as a second net.
- **You get asked before you lose anything.** Leaving the DAW with unsaved work - by the AetherWave logo in the top-left, or by the browser's Back button - opens a dialog with three choices: **Stay here**, **Leave without saving**, or **Save and leave**. If that save fails, you are kept on the page rather than sent away from the work, and a local backup is written. Closing the tab outright falls back to the browser's own prompt, which has no Save button, so prefer the logo when you have work in flight.
- **One DAW tab at a time.** Two tabs open on the same project would overwrite each other. If a second tab spots the first, it pauses its own auto-save and tells you - close the other tab and reload.

You can close the tab and reopen the DAW later without losing your work.

Upload local audio

If you have an audio file on your computer that did not come from Music Studio, click the **upload arrow** above the library search bar and pick the file. Two things happen at once:

- The file is uploaded to your gallery (so it lives there permanently)
- A clip is dropped onto the next available track in the DAW

Supported formats: WAV, MP3, M4A, OGG, FLAC. File size limits depend on your plan (free / starter: 10 MB, studio and up: 50 MB).

Common first-session questions

Why are there four tracks already when I open the DAW? The boot path can restore tracks from a previous session. If you do not want them, click the **Projects** button, then **New Project** to start fresh.

Why can I not hear anything? First, click anywhere on the page. Browsers require a user gesture before audio context can start. Then check the master volume slider at the bottom-right (look for the MASTER fader).

My clip is red and says "Audio Unavailable" The original audio URL has expired or is unreachable. This usually happens with old projects. Re-upload the source from your gallery or Music Studio.

The export is silent / very quiet Make sure at least one track is not muted, and the master volume fader is not at zero. Solo a single track and try again to confirm it produces audio.

Where to go next

- Tracks and clips - drag, split, trim, snap
- Recording live audio - mic and instrument recording
- The Library - samples, SFX, and AI music
- Effects and rack - EQ, reverb, the Pro lineup
- Export and render - WAV vs MP3, troubleshooting

CHAPTER 3

Settings and setup

> Every setting that shapes how the DAW behaves - project tempo and time signature, the grid and snap, audio input for recording, MIDI, and export format and destination. A one-stop reference for configuring your session.

Before you start

The DAW is **desktop-only**. Below 900px wide it shows a "Desktop Required" message with a link back to home, because arranging and mixing need a keyboard and mouse. Use a computer with Chrome or Edge for the full feature set (some export features are Chrome/Edge only).

A first-run guide opens the first time you open the DAW. Reopen it any time from the **Help** button in the header.

Project settings (the header)

- **Project title** - click it to rename. This becomes the base of your export filenames.
- **BPM (tempo)** - type a value, tap the **TAP** button to the beat, or click **Detect** to read the tempo from your first imported song (accurate, server-side). The field shows a clean whole number, while the project keeps the precise tempo internally so the grid stays locked to the song. Type a specific value to override.
- **Time signature** - 4/4, 3/4, or 6/8. Sets how the ruler and grid are divided.
- **Snap** - the grid magnet. Toggle it on or off, and pick the snap value (1 bar, 1/2 bar, 1 beat, 1/2 beat, 1/4 beat) from the dropdown. With snap on, clips lock to the beat when you drag them.
- **Overlap mode** - **Free** lets clips stack on a track; **Trim** trims the underlying clip when you drop another over it. Toggle with the mode button.
- **Metronome** - a click track for recording and timing. Toggle from the header.

Audio input (for recording)

Right-click the **Record** button to open the input settings:

- **Input Device** - which microphone or audio-interface input to record from.
- **Input Mode** - **Instrument (raw)** for guitar, bass, and keys (DSP off, keeps tone), or **Voice** for vocals (echo and noise cleanup on).
- **Input Level** - a gain trim for the incoming signal.
- **Latency offset (ms)** - fine-tune so takes land exactly on the beat.
- **Monitor input** - hear yourself through the speakers (use headphones to avoid feedback).

Full detail is in Recording live audio.

MIDI settings

Click the **MIDI** button in the header to choose a MIDI input device for playing and recording notes. MIDI tracks are edited in the piano roll (see Tracks and clips).

Export settings

- **Format** - choose **WAV** (lossless) or **MP3** (compressed) from the Export dropdown.
- **Export destination** - by default files go to your browser's Downloads folder. To send exports to a specific folder for the session, open the **AI-Engineer** panel and click **Export folder**, then pick your folder. Both the main Export button and the AI batch export then save there. Browsers do not allow a typed absolute path (security), so you choose the folder through the picker. Available in Chrome and Edge.

See Exporting for the full export workflow.

Saving and autosave

The DAW autosaves your project to local storage every 30 seconds, and you can click **Save** to persist it to your account. Reopen saved projects from the **Projects** button. See Getting started for the full save flow.

Where to go next

- Recording live audio - the input settings in depth
- Exporting - format and destination
- Tracks and clips - snap, zoom, and the grid

CHAPTER 4

Tracks and clips

> How to add tracks, drop clips on the timeline, drag and trim them, use snap and zoom, and work with MIDI tracks.

Track types

The DAW has two track types:

- **Audio track** - holds audio clips dragged from your gallery, uploaded files, or AI Music generations
- **MIDI track** - holds MIDI notes you draw or import. Click a MIDI clip to open the **Piano Roll** for editing.

Add a track by clicking the **+ Add Track** button above the timeline and choosing Audio or MIDI.

Adding clips

You have four ways to get a clip onto an audio track:

1. **Drag from My Gallery** - your previously-generated songs
2. **Drag from Samples / SFX** - curated content
3. **AI Music tab** - generate a new track inline and it lands on a track
4. **Library upload** - the up-arrow icon above the search picks files from your computer, uploads them to your gallery, and drops them on the next available track

Drop position respects the **Snap** setting. Default snap is 1 beat. Toggle Snap off to drop at any time position.

Moving clips

- **Drag horizontally** to reposition. Snap applies if it is on.
- **Drag between tracks** to move a clip from one track to another. The track must be of the same type (audio cannot move to a MIDI track).
- **Hold Shift while dragging** to bypass snap temporarily.

Trimming and splitting

- **Drag the right edge** of a clip to shorten or extend it. Extending past the original length loops the source audio.
- **Drag the left edge** to trim the start.

- **Right-click a clip** for the context menu - split, duplicate, delete, rename, change color.
- **Shift + S** while a clip is selected splits it at the playhead position.

The transport bar

Control	Keyboard	What it does
Play / Pause	Space	Toggle playback
Stop	(no shortcut)	Stop and return to position 0
Record	(no shortcut)	Records live audio onto an armed or new track. Right-click for input settings
Loop	(no shortcut)	Toggle loop mode using the loop region in the ruler

The **time display** to the right of transport shows current position in bars (1.1.1 means bar 1, beat 1, sub-beat 1) and clock time (00:00:00).

For a full walkthrough of recording a mic or instrument, see Recording live audio.

Snap and zoom

Above the timeline there is a control bar:

- **Zoom slider** (10 to 400 percent) plus zoom-in / zoom-out / fit-to-view (F) buttons. Higher zoom shows more detail per second.
- **Snap toggle** - turns grid snap on or off. The dropdown next to it picks the snap value (1 bar, 1/2 bar, 1 beat, 1/2 beat, 1/4 beat).

A blue **playhead** line marks the current playback position. Click anywhere on the ruler to jump the playhead.

Loop region

To loop a section of your timeline, drag the loop region handles in the ruler. The shaded area is the loop. Click the **Loop** transport button to toggle whether playback follows the loop.

Track header controls

Each track has its own header with:

- **Track name** - click to rename

- **Volume slider** - 0 to 100, where **75 is unity** (0 dB, no change). Above 75 is boost, up to +12 dB at the top of travel. The tinted zone on the fader shows you where that line is.
- **Lock** - freezes this track's level so a stray drag cannot move it
- **Pan slider** - left / right balance
- **M** - mute
- **S** - solo (only soloed tracks play)
- **R** - record-arm
- **A** - automation lane (volume / pan envelopes)

Working with MIDI

When you add a MIDI track, double-click any MIDI clip to open the **Piano Roll** modal. Inside:

- Draw notes on the grid
- Set velocity in the lane below
- Quantize to 1/4, 1/8, 1/16, or 1/32 notes
- Import or export MIDI files (top-left and top-right buttons)

Close the modal with the X or by clicking outside it. Edits commit when you click **Apply**.

Turn a stem into MIDI

Select an audio clip of a melodic part (bass, keys, guitar, or a vocal line) and use **To MIDI** to detect its notes and place them on a new MIDI track, ready to edit in the piano roll or re-voice with a different instrument. Note detection loads the first time you use it. Drums and other unpitched material are skipped, since there is no pitch to detect.

Saving your arrangement

The DAW auto-saves to local storage every 30 seconds, and you can click **Save** at any time to persist the project to your account. See Getting started for the full save flow.

Where to go next

- Effects and rack - shape your sound
- Export and render - render the final mix

CHAPTER 5

Recording live audio (mic, guitar, keys)

> Record a live microphone or instrument straight onto a track. Covers arming, input device and mode, monitoring, latency compensation, and lossless capture, plus recording over a backing track.

What it does

The DAW records live audio from your computer's microphone or audio interface directly onto a track. Plug in a guitar, bass, or keyboard through an interface, or use a USB or built-in mic for vocals, and the take lands on the timeline as a normal audio clip you can trim, move, and mix.

Recording captures **lossless raw audio** (uncompressed PCM), so a clean direct instrument keeps its full tone and dynamics.

Record in three steps

1. Position the playhead where you want the take to start.
2. Click the red **Record** button in the transport bar. If no track is armed, the DAW creates a new audio track for the take. Grant the browser microphone permission the first time.
3. Play or sing, then click **Record** again to stop. The take appears as a red clip on the track.

Playback starts automatically when you record, so you can perform along to whatever is already on the timeline.

Input settings

Right-click the Record button to open the input panel:

- **Input Device** - pick which microphone or interface input to record from. Useful when you have both a built-in mic and an external interface connected.
- **Input Mode**:
  - **Instrument (raw)** - for guitar, bass, and keys. The browser's echo cancellation, noise suppression, and auto-gain are turned **off** so the signal is captured exactly as it comes in, preserving tone and dynamics.
  - **Voice** - for vocals and spoken word. Re-enables the browser's echo, noise, and gain cleanup for a cleaner spoken result.
- **Input Level** - a gain trim for the incoming signal if it is too quiet or too hot.

- **Latency offset (ms)** - see below.
- **Monitor input** - hear yourself through the speakers while recording. Off by default to avoid feedback. If you turn it on, **use headphones** so the mic does not pick up the playback.

Latency compensation

Every computer has a small round-trip delay between the sound you play and the sample the browser records. Without correction, a note you play on the beat would land slightly late on the timeline.

The DAW **measures your input and output latency automatically** and shifts each take earlier so it lands where you played it. For an exact fit to your specific interface, nudge the **Latency offset (ms)** field in the input panel:

- If your recordings still sit a touch **behind** the beat, increase the offset (pulls the take earlier).
- Most audio interfaces land somewhere between 5 and 30 ms.
- Record a few bars against a metronome or a drum loop, then adjust until the take locks to the grid.

Recording over a backing track

A common workflow is to lay down a bed first, then perform over it:

1. Drag a song from your gallery, or generate an instrumental in the **AI Music** tab, onto a track.
2. Turn on the **Metronome** if you want a click, or just play to the bed.
3. Arm a new track and hit **Record**. You will hear the bed while you play.
4. Swap the AI bed out later if it was only a scratch guide.

See The Library: samples, SFX, and AI generation for generating a bed to record against.

Arming a specific track

Each track header has an **R** (record-arm) button. Arm a track first if you want the take to land on that exact track rather than a new one. Only one track records at a time.

Troubleshooting

"Microphone access denied" The browser blocked mic access. Check the site permissions (the padlock icon in the address bar) and allow the microphone, then reload.

"No microphone found" No input device is connected or the browser cannot see one. Check that your mic or interface is plugged in and selected as an input in your operating system.

My take lands late (or early) Adjust the **Latency offset (ms)** in the input panel. Increase it to pull takes earlier, decrease it if they land ahead of the beat.

I hear an echo or howl while recording Monitoring is on and the mic is picking up the speakers. Switch to headphones, or turn **Monitor input** off.

The recording is very quiet or distorted Raise or lower the **Input Level** trim, and check your interface's gain knob. For instruments, make sure you are on **Instrument (raw)** mode.

Where to go next

- Tracks and clips - trim and arrange your take
- The Library: samples, SFX, and AI generation - generate a bed to play over
- Effects and rack - shape the recorded sound

CHAPTER 6

The Library - samples, SFX, and AI generation

> The Library panel is where sound comes from. Browse your gallery, your own sample directory, and the shared community library, generate sound effects, and generate music that matches your project tempo, all draggable onto the timeline.

The Library panel

The panel on the left side of the DAW is where you bring sound in. It has four tabs:

- **My Gallery** - your generated songs and uploaded audio
- **Samples** - your own uploads plus a shared community sample library
- **SFX** - generate sound effects on demand
- **AI Music** - generate music directed to your project

Every item in every tab is **draggable onto a track**, and has a preview button so you can audition it before you commit.

Finding things: search, sort, and view

Under the library search box there is a small toolbar with two controls that work across every tab:

- **Sort** - `Newest` (the natural order the tab loads in), `A to Z`, or `Z to A`. Sorting is case-insensitive and number-aware, so "mango 2" comes before "mango 10" instead of after it.
- **View** - switch between the **list** view (one row per item) and the **tile** view (a two-up grid with the artwork on top). Tiles are easier to scan when you know a sample by its cover.

Both settings stick, so the library opens the way you left it next time.

My Gallery

Your Music Studio generations and any audio you upload live here. Drag any item onto a track to add it as a clip.

Add to Library (the upload button above the search) uploads audio files from your computer to your gallery without dropping them on the timeline, so you can build up a personal library to drag from later. Supported formats: WAV, MP3, M4A, OGG, FLAC.

Samples

Sound in the DAW comes from three separate directories, and they never mix:

Directory	Where it lives	Yours to delete?
Gallery	The My Gallery tab - your generated songs and your uploads	Yes, from the Gallery
Your Samples	Top of the Samples tab - your own personal sample directory, plus the audio you have uploaded to your gallery	Yes, for files in your personal directory
Community	Below it, in the Samples tab - a shared, curated pool every signed-in user can browse, grouped by category (drums, FX, and more)	No, never

Drag any sample onto a track to use it. This is the place to grab loops, one-shots, and textures to build with.

Deleting your own samples. Files that live in your personal sample directory carry a **trash button** on the row. Click it, confirm, and the file is gone from your directory and disappears from the list. Nothing else in the Samples tab can be deleted: the Community pool is shared and permanently read-only, and gallery-backed uploads are removed from the **Gallery** instead, where they actually live. Deleting a personal sample cannot be undone, so make sure you have the file elsewhere if you want it back.

SFX generator

The SFX tab generates sound effects from a text description using AI:

1. Describe the sound you want, for example "vinyl crackle", "riser into a drop", or "deep impact hit".
2. Set the **length** with the slider (1 to 20 seconds). The credit cost updates as you change the length.
3. Click **Generate SFX**. When it finishes, the effect is saved to your gallery and appears under **Your SFX**, ready to drag onto a track.

Generated SFX are not dropped on the timeline automatically, so you can place them exactly where you want.

AI Music generator (directed)

The AI Music tab generates music shaped by structured controls rather than a bare prompt, so the result fits the track you are building:

- **Genre** - pick from a large grouped list
- **Mood** and **Energy**
- **BPM** with a **Match project tempo** toggle. When it is on, generation is locked to your session's tempo so the result lands on the grid. Turn it off to set a custom BPM (each genre suggests a

typical one).

- **Key**
- **Instrumentation** and **Exclude** terms
- **Instrumental** or vocals (with a vocal gender option)
- **Extra notes** for anything the controls do not cover

A live **prompt preview** shows exactly what will be sent. Click **Generate Music** and the finished audio drops onto the timeline.

What you get back. Generation costs **20 credits**, the same on every plan, and one run returns **two takes of the same prompt**. Each take lands on **its own track**, named **AI Music 1** and **AI Music 2**, so you can solo them against each other, keep the one that works, and delete the other. Nothing is dropped on top of the tracks you already had.

While it runs. The panel reports real progress as it goes, not just a spinner: **"Lyrics created! Generating audio..."**, then **"First track ready! Finishing up..."**, then a toast telling you how many tracks were added. A full run takes a couple of minutes.

Because it matches your project tempo, AI Music in the DAW is best for **starting a foundation, filling in an element, or making a scratch bed to record over**, rather than producing a finished standalone song. For finished songs, use Music Studio.

Signing in and credits

Browsing the community sample library and all generation (SFX and AI Music) require a signed-in account and credits. If you are not signed in, each of these shows a **Sign in** prompt. Generation shows the credit cost before you run it, and you are only charged on success.

Where to go next

- Recording live audio - play over a generated bed
- Tracks and clips - arrange what you bring in
- Export and render - render the final mix

CHAPTER 7

Mixing

> Balance your tracks - volume, pan, mute, solo, meters, lane locks, and automation - using the track headers and the mixer. Includes gain-staging tips for AI-generated audio and how the AI-Engineer can mix for you.

The mixer and track headers

Every track has the same set of controls, available both on the track header and in the mixer panel below the timeline:

Control	What it does
Volume	Track level. Unity (0 dB) sits at 75 on the fader and the top of the travel is +12 dB.
Pan	Left / right placement in the stereo field.
M (mute)	Silences the track.
S (solo)	Plays only soloed tracks (see below).
R (record-arm)	Arms the track for recording.
A (automation)	Opens the automation lane for volume and pan.
Lock	On the track header. Freezes the track's level so it cannot be dragged by accident (see below).

The **mixer panel** shows a channel strip per track for a classic mixing-board view. You can **pop the mixer out** into its own window (handy on a second monitor).

Meters

Every track has a **level meter in its track header**, right under the volume fader, as well as the vertical pair on its mixer channel strip. Both are driven from the same measurement, so there is one truth about how loud a track is and the two never disagree. The head meter is the quick glance while you are arranging; the mixer strip is the detailed read while you are balancing.

Two things to know:

- Meters only move while **audio is actually playing**. Silence reads as nothing, which is correct.

- They only animate while the **DAW tab is in the foreground**. Browsers pause animation in background tabs, so a meter that looks frozen after you tabbed away is not a bug. Come back to the tab and it picks up again.

The master meter sits on the master channel at the right of the mixer. That is the one to watch for clipping.

Fader safe zones

Both faders, on the track head and on the mixer strip, are **tinted to show you where you are**:

- **Below the white tick** is the normal operating range. This is where a mix should live.
- **The white tick** is unity, 0 dB. The track plays at exactly the level of the audio on it, no gain added or taken away.
- **Above the tick**, tinting warms into amber and red. That is **boost**, and it is headroom being spent rather than level being gained.

The top of the fader is +12 dB, which is deliberately generous so you always have room to reach for. But the guiding principle is that **you should never have to max a fader to hit your target level**. If one track needs the last few dB of boost just to be heard, the real problem is that everything else is too loud: pull the rest down instead, and let the master carry the final level.

Lane lock

Once a track is dialled in, click the **lock button** on its track header. That freezes the track's level, so a stray drag while you are working elsewhere cannot move a fader you already got right.

The lock covers **both surfaces at once**: the head fader and the matching mixer strip dim and stop responding to the pointer together, so there is no back door into the level from the mixing board. Click the lock again to release it. Locking affects the level control only, so mute, solo, pan, and effects all keep working.

Solo logic

When no track is soloed, each track plays according to its own mute state. The moment you solo one or more tracks, everything **not** soloed is silenced, so you can focus on a part. Un-solo to return to the full mix. Solo and mute are independent, so a soloed-but-muted track still stays quiet.

Balancing a mix

AI-generated tracks often come in at different levels, which is normal. A quick way to a clean balance:

1. Play the full song and listen for the loudest and quietest parts.

2. Pull loud tracks down with the volume control rather than pushing quiet ones way up (leave headroom).
3. Pan elements apart (for example, double-tracked guitars left and right) to open up the center for vocals and bass.
4. If one track jumps around in level, add a **Compressor** to even it out (see Effects and the rack).

Keep the master from clipping: if the mix is too hot, bring the whole thing down a little rather than fighting individual tracks. The master fader is at the bottom right.

Automation

Click **A** on a track to open its automation lane. Draw volume or pan envelopes over time to fade a part in, duck it under a vocal, or sweep the pan. Automation is applied on playback and included in your export.

You do not have to draw it by hand. The AI-Engineer can write the curve for you from a plain description: "fade the intro in", "duck the pad under the vocal", "automate a filter sweep into the drop". It automates volume, pan, or any parameter of an effect already on the track, and the breakpoints land in the lane where you can drag them afterward.

Let the AI-Engineer mix

You can drive all of the above by asking the AI-Engineer in plain language: "mute the drums", "turn the bass down to 70", "pan the guitar left", "add a compressor to the vocals". Every change is undoable with Ctrl+Z.

Where to go next

- Effects and the rack - shape each track's tone
- Mastering - the final polish on the master bus
- Exporting - render the finished mix

CHAPTER 8

Effects and the rack

> How the effects rack works, the difference between master-bus and per-track effects, the free vs Pro lineup, how every rack change is undoable, and how each effect contributes to the final mix.

What the rack is

The **Effects Rack** is the right-side panel of the DAW. It holds a chain of audio effects that process whatever the **rack target** is set to:

- **Master Bus** (default) - effects apply to the whole mix on the way out
- **A specific track** - effects apply only to that track

Click any channel strip in the mixer to point the rack at that track. The strip highlights and the chain shown updates to match. (The old target dropdown is gone.)

Adding an effect

1. Click **+ Add Effect** at the top of the rack
2. The menu lists effects grouped by category: Dynamics, EQ and Filters, Time, Modulation, Distortion, Spatial, Utility
3. Free effects open immediately. Pro effects are marked with a `PRO` badge and require an upgrade.
4. The effect appears in the rack with default parameters and is active right away

Effects process in **order from top to bottom**. Drag them to reorder.

Per-effect controls

Every effect card has:

- A **bypass toggle** (eye icon) - turns the effect off without removing it. Useful for A/B comparison.
- A **remove button** (X) - removes the effect from the chain
- **Knobs and sliders** for each parameter

Right-click an effect for a context menu (reset, copy settings, paste settings).

Whole-chain controls

Once a channel has inserts, two controls sit in the rack title row:

- **ACTIVE / BYPASSED** - switches the entire chain off in one click, so you can A/B the processed channel against the dry signal. Switching it back on restores exactly what was on before: if you had deliberately bypassed a single insert, it stays bypassed.
- **CLEAR CHAIN** - removes every insert on the channel. It always asks first, and Ctrl+Z brings the whole chain back with every dialled parameter intact.

Neither appears on an empty channel.

Starter Chains

The **Starter Chain** button loads a ready-made chain. The menu has two parts:

- **My Chains** - anything you saved yourself with *Save this channel's chain*
- **Starter Chains** - the factory chains that ship with the DAW

Factory chains cover eight genres (Rock, Pop, Hip-Hop, Electronic, R&B, Acoustic, Metal, Podcast) across nine instrument roles, and are identical for every user on every plan - "Rock Bass" is the same known-good starting point on any account.

The menu reads your **track name** and shows only what fits. A track called `Bass`, `Bass Gtr`, `808` or `Sub` gets the eight bass chains; `Lead Vocals` gets the eight vocal chains. Each entry leads with what it sounds like, with the chain underneath in short codes:

> **Hip-Hop Vocals** - Cut low, compressed hard, right up front - `Flt-EQ-Cmp-Rev-Dly`

Recognised roles: vocals, drums, bass, guitar, piano, keys, synth, strings, fx. A track whose name matches none of them (`Track 1`, `Audio 3`) shows no starter chains, and neither does the Master Bus - rename the track to turn them on. You can always add effects by hand either way.

Saving a chain under a genre's name keeps the factory version intact underneath, so customising "Rock" never overwrites the original.

Everything in the rack is undoable

Ctrl+Z works on the whole rack. Adding an effect, removing one, bypassing one, and every knob move all go on the same undo stack as the rest of your edits, so you can push a sound somewhere strange and walk it straight back.

Two details worth knowing:

- **One undo step per knob gesture.** Dragging a knob from 20 to 80 is a single undoable move, not one step per pixel, so a couple of Ctrl+Z presses will not leave you halfway through a sweep. Undoing a knob move restores just that parameter and disturbs nothing else in the rack.

- **Undo restores the exact chain order.** Remove a reverb from the middle of a chain, undo, and it comes back where it was, not appended to the end. Order matters to the sound, so the undo preserves it.

Rack changes also count as unsaved work, so they mark the project dirty, get picked up by auto-save, and trigger the save prompt if you try to leave with them unsaved.

The free lineup

These effects are available on all plans and are exported correctly when you render the mix:

Effect	Category	What it does
3-Band EQ	EQ	Boost or cut low, mid, and high frequencies. The most common mix tool.
Compressor	Dynamics	Tames loud peaks and brings up quiet sections. Add this to vocals, drums, or any track that feels uneven.
Reverb	Spatial	Adds room ambience. Wet / dry mix and decay control how present the reverb is.
Delay	Time	Echoes the signal at a set time. Sync to BPM for tempo-locked echoes.
Filter	EQ	Lowpass, highpass, bandpass, notch. Useful for sweeps and DJ-style transitions.
Distortion	Distortion	Adds grit. Drive and tone shape the character.

The Pro lineup

These require a paid plan. They appear with a `PRO` badge in the menu and produce no sound on render unless your account is on a qualifying tier.

- **Limiter** - hard cap on output level, used at the master bus to prevent clipping
- **Multiband Compressor** - independent compression on three frequency bands
- **De-Esser** - reduces harsh "s" sounds on vocals
- **8-Band Parametric EQ** - surgical frequency control
- **Gate** - silences audio below a threshold (cleans up bleed)
- **Tape Saturation** - warm analog character
- **Convolution Reverb** - real-room impulse responses
- **Vocoder** - turn instrumental into a robotic voice
- **Pitch Correction** - tune vocals to a key
- **Stereo Width** - widen or narrow the stereo image
- **Glitch** - rhythmic stutter and chop

For the full menu and what each does, browse the **Effects Store** (the store icon in the rack header).

Master bus vs track effects

Use master bus for...	Use track effects for...
Final compression / limiting	Cleaning up a single source
Overall EQ shaping	Vocal-specific de-essing
Stereo width on the whole mix	Putting one element in a different room
Reverb on a wide send	Tightening a single drum hit

A typical mix uses a mix of both. Most tracks get one or two per-track effects (compressor, EQ), and the master bus gets a final compressor or limiter to glue everything together.

Effects and export

When you click **Export**, the renderer rebuilds the same effect chains in an offline audio context and processes the mix faster than realtime. Free effects render exactly as you hear them. Pro effects are skipped if your plan does not include them.

A 2-second tail is added to the export so reverb and delay can decay naturally.

Browse the Effects Store

Click the **store icon** (top-right of the rack) to open the **Effects Store**. You can browse the full Free and Premium lineups, see thumbnails and descriptions, and tap **Back to DAW** to return.

The Store is a discovery surface today. Premium effects unlock automatically with the matching plan - there is nothing to install.

Add effects by asking

You can also add an effect to a track or the master bus by asking the AI-Engineer: "add a compressor to the drums", "put a reverb on the vocals". It is undoable like any effect you add by hand.

Where to go next

- Mixing - balance the tracks the effects sit on
- Mastering - the master-bus chain
- Export and render - render the final mix with effects baked in

CHAPTER 9

Mastering

> The final polish - use the master bus to glue and control your mix, hit a sensible loudness, and know when to reach for the dedicated one-click Master Audio tool instead.

What mastering means here

Mastering is the last step: taking a finished, balanced mix and giving it cohesion, controlled loudness, and a polished top end so it holds up next to commercial tracks. In the DAW you master on the **master bus** - the effects chain that processes the whole mix after the individual tracks.

Master well by finishing the mix first. Mastering cannot fix a mix that is unbalanced or clipping; sort those out in Mixing before you polish.

The master bus

Open the **Effects Rack** and switch it to the **Master Bus** (click the master channel strip). Effects you add here process the summed mix. A typical, restrained master chain:

1. **Compressor** (gentle) - a little glue, low ratio, just a dB or two of reduction on peaks.
2. **EQ** - small, broad moves. A touch of high shelf for air, a gentle low cut if it is muddy. Subtractive first.
3. **Limiter** (at the end) - catches peaks and raises overall level without clipping. This is the loudness stage.

Less is more on the master. If you find yourself doing big moves here, go back to the mix.

Loudness

Aim for a level that is competitive but not crushed. Streaming platforms normalize to roughly **-14 LUFS**, so pushing far past that mostly just squashes dynamics without getting louder on playback. Use the limiter to control peaks and set a sensible ceiling (around -1 dBFS true peak is a safe target), and trust your ears against a reference track.

The one-click Master Audio tool

If you want a fast, hands-off master, AetherWave has a dedicated **Master Audio** tool (separate from the DAW) that masters a track with a chosen preset. Use it when:

- You want a quick, reliable master without building a master chain by hand.
- You have already exported a mix and want to polish that file.

Use the **DAW master bus** instead when you want full manual control over the master chain, or want the master processing baked into your DAW export.

What the Master Audio tool costs

20 credits per track. If you are short, the tool tells you how many you need rather than half-running.

It is **included at no credit cost** on the **Artist, Producer, Mogul and Ultimate** plans, so on those plans you can master as often as you like without watching the balance.

On the free plan, your **first master is free**. Once you have used it, further masters are 20 credits each. The banner offering a free master stops appearing after you have claimed yours.

There is no separate "remaster" product, so if you are looking for the price of remastering a track, this is it: run the same Master Audio tool again on the file, at the same 20 credits (or free on the plans above). Each run is charged on its own, so trying a second preset on the same track is a second master.

Mastering a track again does not undo or replace the earlier master. Each run saves a new file to your gallery and leaves the original untouched.

Export at the right level

Master with the export in mind: your master-bus chain is applied when you export, so the WAV or MP3 you render already carries the master. Export lossless **WAV** for distribution or further work, and keep a little headroom under 0 dBFS. See Exporting.

Where to go next

- Mixing - get the mix right first
- Effects and the rack - the master-bus effects in detail
- Exporting - render the mastered mix

CHAPTER 10

Export and render

> How the DAW renders your final mix, the difference between WAV and MP3, what gets included, and what to do when an export fails.

What "export" means

When you click **Export WAV** (or pick MP3 from the dropdown chevron next to it), the DAW rebuilds your entire project in an **offline audio context** - a faster-than-realtime renderer that processes the full timeline, every track, and every effect, then writes the result to a downloadable file. For control over format and quality (bit depth, sample rate, bitrate, loudness, stems), use **Render options...** in the same dropdown - see Render options below.

The render covers:

- All audio clips on all tracks
- Track volume, pan, mute, and solo state
- Volume and pan automation
- Track-level effects (chains attached to specific tracks)
- Master-bus effects
- Master volume

Plus a 2-second tail at the end so reverb and delay can decay naturally.

WAV vs MP3

	WAV	MP3
Quality	Lossless, 16-bit / 44.1 kHz stereo	Compressed (lossy)
File size	Larger (about 10 MB per minute)	Smaller (about 1 MB per minute)
Processing	Encoded in your browser	Sent to a server, encoded, sent back
Speed	Fastest	Slower (network round-trip)
Best for	Archive, distribution masters, bringing into another DAW	Sharing, social posts, streaming previews

WAV is recommended for any "real" use of the audio. Use MP3 when file size matters more than fidelity. The quick **Export WAV** button gives you a 16-bit / 44.1 kHz file; to render higher-resolution WAV (24-bit, 48/96 kHz), FLAC, or a normalized master, use **Render options** below.

The export flow

1. Click **Export WAV** in the header
2. The button shows a "Rendering..." spinner
3. Your browser opens a save dialog within a few seconds (longer for long projects)
4. Open the file in any audio player to verify

For MP3:

1. Click the small chevron next to Export WAV to open the format menu
2. Pick MP3
3. The button label changes to "Export MP3" and an export starts immediately
4. The first part renders to WAV, then uploads to the audio service for compression
5. The MP3 downloads when conversion is complete (typically 5 to 15 seconds for a 3-minute song)

Render options (format, quality, stems)

For full control, open the export dropdown (the chevron) and choose **Render options...**. A dialog lets you set:

- **Format** - **WAV** (lossless), **FLAC** (lossless, compressed to about half the size), or **MP3** (lossy).
- **Bit depth** (WAV/FLAC) - **16-bit** (CD quality), **24-bit** (studio and distribution standard), or **32-bit float** (maximum headroom, WAV only). Higher bit depth means more dynamic-range detail and a larger file.
- **Sample rate** - **44.1 kHz** (CD / most streaming), **48 kHz** (video and modern distribution), or **96 kHz** (high-resolution). The mix is rendered at the rate you pick.
- **MP3 bitrate** - 320, 192, or 128 kbps (shown only for MP3). 320 is the highest MP3 quality.
- **Loudness** - **None** (leave levels as mixed), **Peak -1 dBFS** (raise so the loudest sample sits just below the ceiling), or **Normalize to -14 LUFS** (streaming loudness target) / **-9 LUFS** (louder). LUFS normalization matches the perceived-loudness targets that Spotify, YouTube, and Apple Music use.
- **Stems** - tick "Export each track as its own file" to render every track separately (its own fader, pan, and effects, without the master bus), for mixing elsewhere or building sample packs. Set an export folder first so the stems land together.

FLAC, MP3, and LUFS-normalized renders finish on the server (the same pipeline the Mastering tool uses); plain WAV renders entirely in your browser. A 24-bit / 48 kHz WAV is the safe default for a distribution-ready master.

You can also just ask the AI-Engineer: "render a 24-bit 48k WAV", "bounce a 320 MP3", "export FLAC stems", or "render and get me to -14 LUFS". See The AI-Engineer.

What if export fails

"No audio clips in project" - the timeline is empty. Drop at least one clip on a track and try again.

Spinner hangs / never finishes (MP3 only) - the audio service is unreachable. Try WAV instead, or wait a minute and retry. WAV is always client-side and reliable.

File downloads but is silent - check that:

- At least one track is not muted
- Master volume fader is above zero
- No track is soloed except a silent one
- Effects are not bypassing all signal (a wet-only delay with no dry signal can look like silence at start)

Audio cuts off too early - the renderer measures duration from the latest clip end. If your last note has a long reverb tail, extend the clip itself or add a silent clip at the end so the measured duration is long enough.

Distorted / clipping audio - reduce master volume, or add a Limiter on the master bus (Pro effect). Many AI-generated tracks are already loud, and stacking multiple tracks pushes the sum past 0 dBFS.

One track sounds different in the export than in playback - usually a Pro effect that does not render on your plan. Check the effects on that track and bypass any with a `PRO` badge to confirm.

Filename and where it saves

Files save with the pattern `{project-title}_{timestamp}.{ext}` (for example, `My Mix_1715281234567.wav`). By default they go to your browser's Downloads folder.

Export folder: to send exports to a specific folder for the session instead, open the **AI-Engineer** panel and click **Export folder**, then pick your folder (for example a "My-Sample-Pack" folder). After that, both the main Export button **and** the AI batch export save straight into that folder. Browsers do not let a page write to a typed absolute path for security, so you choose the folder through the picker. This is available in Chrome and Edge; on other browsers exports use Downloads.

Batch export for sample packs

Instead of exporting one mix at a time, you can export several full-song WAV mixes at once by asking the AI-Engineer. For example:

> "Export full WAV mixes of: drums muted, bass and drums only, and vocals only."

Each mix mutes or solos the right tracks, renders through this same engine, and saves as its own named WAV, then your live mix is restored. There is also a one-click **Foundational Samples** preset (Full Mix, Drumless Mix, Vocal Mix, Bass and Drums, plus each track as its own stem) that confirms the list with you before exporting. Set an export folder first so the whole pack lands together. See The AI-Engineer.

What is NOT in the export

- Sample previews (clicking the play icon on a library item) - those are just for auditioning
- Sample tabs / SFX tabs that you have not dragged onto a track
- The metronome click track
- Looping - the export renders one full pass of the timeline, not a loop

What to do with the file

- Upload to **SoundCloud, Bandcamp, or YouTube** (WAV gets re-encoded by their servers, so there is no quality penalty)
- Re-import into the DAW as a new track if you want to layer further
- Run through the Mastering tool for a final polish pass. The Mastering tool is separate from this DAW export: it runs a Pedalboard processing chain plus loudness normalization, and its output is a **24-bit / 48 kHz WAV**. That is a full studio-master pass; the DAW's own **Render options** can already give you a 24-bit / 48 kHz WAV (with optional LUFS normalization), while Mastering adds the dedicated processing chain on top.
- Generate a **Track Reel** (social Shorts of your song) from the Track Reels page
- Pin it as a **Soul Forge band identity** if it represents an artist persona

Where to go next

- Effects and rack - shape the sound before exporting
- Tracks and clips - timeline mechanics

CHAPTER 11

The AI-Engineer

> The AI-Engineer is a chat assistant built into the DAW that knows your project and can do real work - answer questions, generate music and SFX, mix hands-free (mute, volume, pan, effects, automation), replace a section of a track with new AI audio, and batch-export sample packs. Everything it changes is undoable.

What it is

The **AI-Engineer** is a conversational assistant that lives in the DAW's right panel, on its own tab next to the **Effects Rack**. It understands the DAW and it is aware of your current project (tempo, tracks, levels), and it can take real actions on your mix. You talk to it in plain language.

It is **free** to chat with (generating music or SFX still costs credits, the same as doing it by hand), and every change it makes to your project is **undoable with Ctrl+Z**.

Opening it

In the right panel, click the **AI-Engineer** tab (next to Effects Rack). The effects rack stays put behind the other tab. A quick-start greeting appears the first time, along with a few suggestion chips.

What it can do

1. Answer questions and read your project

- How-to: "How do I add reverb?", "How do I record my guitar?" (grounded in this manual)
- Project facts: "What tempo is this?", "What tracks do I have?", "Which track is loudest?"

2. Generate music and sound effects

- "Make a lo-fi beat at this tempo" - generates an instrumental and drops it on the timeline, defaulting to your project tempo so it fits the grid
- "Generate a riser into a drop" - creates a sound effect in your SFX tab

A live progress card shows when a real generation is running. Generation needs a signed-in account and credits. See The Library.

3. Mix and edit, hands-free (all undoable)

Refer to tracks by their name and just ask:

- **"Mute the drums"**, **"turn the bass down to 65"**, **"pan the guitar left"** - volume is 0 to 100 on the fader, where **75 is unity** (no change) and 100 is the top of travel at +12 dB. Pan is -100 (left) to 100 (right).
- **"Add a compressor to the drums"**, **"put a limiter on the master"**
- **"Add a vocal track"**, **"rename track 1 to Drums"**, **"set the tempo to 120"**
- **"Delete the Vocals track"** - deleting asks you to confirm first with a Cancel / Delete button

If a track name does not match, it tells you instead of guessing. It cannot move or trim individual clips yet; for that, use the timeline controls (see Tracks and clips).

You can also give it a goal and let it work out the details across every track. For example, after loading a stem-separated song whose tracks are named like "Leverage wav Guitar mp3", "Leverage wav Bass mp3", and so on, just say:

> "Rename the tracks with only the instrument name."

It renames them all in one go - Guitar, Bass, Vocals, Drums, Piano, Other - each with its own confirmation, and the change shows in the track headers and the mixer. This is handy for tidying up imported stems before you build a sample pack.

A starting mix in one step. Ask for a genre or style and the AI-Engineer sets sensible starting levels across every track at once:

> "Please set the track levels for a Pop style mix."

It reads each track's role, makes the lead vocal the reference so it stays present, and sets every other track relative to and below it. Crucially, it leaves **headroom** - keeping the loudest element off the ceiling so the combined mix peaks below 0 dB rather than clipping. It applies as a single undoable step (one Ctrl+Z reverts the whole balance). This is a starting point to build from, not a finished master, so tweak to taste from there. It works for other styles too ("a lo-fi mix", "a rock mix").

Save your own mix presets. The style mix above comes from the AI-Engineer's general knowledge. Once you have tuned the balance the way *you* like it, you can save it as your own named preset and recall it exactly on any project:

> "Save this as my Pop mix."

That snapshots every track's current volume and pan under the name you give it. Later, on the same or a different song, just say:

> "Set my Pop mix."

and it applies your saved levels to the matching tracks in one undoable step - your exact numbers, not a fresh guess. This is deterministic: what you saved is what you get. The AI-Engineer knows which presets you have saved and will use them when you refer to "my" mix by name; ask for a plain style with no saved match and it falls back to building fresh starting levels. You can also say "delete my Pop preset" to remove one. Presets are stored in your browser, so they stay with the machine you save them on.

Save and recall whole mix versions. A preset is a reusable level template; a **version** is a complete snapshot of *this song's* mix - every track's volume, pan, mute, and solo *and* all the effect chains, on the tracks and the master. Use versions to try alternate mixes and jump between them:

> "Save this as Mix A." ... (you try a brighter, wetter take) ... "Save that as Mix B." ... "Go back to Mix A."

Saving captures the current mix exactly; recalling restores the whole thing - levels, mute/solo states, and every effect with its settings - in one undoable step. Leave the name off ("save a version") and it auto-numbers Version 1, 2, 3. Versions are tied to the project you are in, so one song's mixes never bleed into another, and they are stored in your browser. Say "delete Mix B" to remove one. This is the safe way to experiment: snapshot where you are, push the mix somewhere new, and step back any time.

Effect recommendations. Beyond levels, the AI-Engineer can suggest a full effects chain tailored to your tracks and your style. Ask:

> "What effects should I add for a Pop mix?"

and it walks each track with specific, reasoned moves - a low-cut and presence boost plus compression and a short plate on the lead vocal, glue compression and EQ on the drums, tightening on the bass, a light room on the guitar, gentle glue and tape saturation on the master bus, and so on. It then offers to apply them, and it does not just insert the effect: it **dials in the settings** it recommended - the EQ presence lift, the compressor ratio and threshold, the reverb size, the delay time - so the chain lands as a real starting patch rather than defaults. Say "yes, add them all" and it inserts each effect with its settings into the right rack, or pick and choose ("just the vocal compressor and the master glue").

You can also fine-tune any effect afterward by asking, and it adjusts the existing effect in place:

> "Shorten the vocal reverb to about 1.2 seconds." / "Back the delay feedback off to 15." / "A touch more compression on the drums."

It keeps every value in the effect's valid range, refreshes the knobs in the rack, and, as always, it is undoable with Ctrl+Z. The dialed-in values are a musical starting point - open the rack and nudge them to taste.

Automation, drawn for you. Anything that needs to *change over time* is automation, and the AI-Engineer can draw the curve rather than leaving you to click breakpoints into a lane. Ask for the move:

> "Fade the intro in." / "Automate a filter sweep into the drop." / "Duck the pad under the vocal." / "Pan the guitar left to right." / "Ride the vocal level up in the chorus."

It automates a track's **volume**, its **pan**, or **any parameter of an effect already on that track** (add the effect first, then ask for the sweep). It writes the breakpoints into the track's automation lane, so you can open the lane and drag them afterward, and both playback and your export follow the curve. Values are clamped to the parameter's real range. Ask it to clear a curve and it deletes that parameter's automation. See Mixing for the lane itself.

4. Replace part of a track with new AI audio

The **Replace Audio** chip re-renders a *section* of a track instead of the whole thing. Use it to redo a chorus, rebuild a weak bridge, or drop a brand new section into an arrangement.

Before you click it, mark what you want changed: drag a loop range on the ruler, or click the clip you want to work on. The card reads that selection and the track it belongs to, and shows both back to you so you can check the target before spending anything.

Click **Replace Audio** in the AI-Engineer panel and pick one of three operations:

Operation	What it does
Replace this part	Regenerates the selection, using the audio on either side as context so key and tempo stay continuous. This is the default when something is selected, and the strongest of the three.
Restyle it	The same part reimagined. Keeps the underlying structure and timing, changes the style, energy, or instrumentation.
New section	Brand new material with nothing to preserve. This is the only one that does not need a selection.

Then describe what you want in the text box ("harder-hitting drop with a heavier snare") and click **Go**.

It costs **60 credits per render on paid plans**, stated on the card before you commit, and a render can take a couple of minutes.

How the result lands. Nothing is overwritten. The DAW duplicates the source track, carves the selected region out of the duplicate, drops the new audio into that hole, and places the duplicate **directly below the source**. That means you can mute one against the other and A/B the new section *in full context*, with the rest of the arrangement still playing. Keep the take you prefer and delete the other. Like everything else here, it is undoable with Ctrl+Z.

If you pick Replace or Restyle with nothing selected, it tells you to select a region first rather than guessing.

5. Batch-export mixes and sample packs

This is a highlight for sample-pack creators. Ask for several full-song WAV mixes at once and each downloads as its own file:

> "Export full WAV mixes of: drums muted, bass and drums only, and vocals only."

Each mix mutes or solos the right tracks, renders through the real export engine (your faders and effects included), and saves a WAV named `project_mix.wav`. Your live mix is restored afterward.

It can also **render with quality options** - format, bit depth, sample rate, MP3 bitrate, loudness, and stems - the same choices as the Render dialog. Just say what you want:

> "Render a 24-bit 48k WAV." / "Bounce a 320 MP3." / "Export 24-bit FLAC stems." / "Render and get me to -14 LUFS."

Defaults, if you do not specify, are a 24-bit / 48 kHz WAV with no loudness change. See Export and render for what each option means.

Setting an export folder

By default, exports go to your browser's Downloads folder. To send them to a specific folder instead (great for keeping a sample pack together):

1. Click the **Export folder** button in the AI-Engineer panel.
2. In the folder picker, navigate to and select (or create) your folder, for example a "Leverage-Sample-Pack-144bpm" folder, and allow write access.
3. The button lights up with the folder name. For the rest of the session, all exports - the AI batch export **and** the main Export button - save straight into that folder.

Browsers do not let a web page write to a typed absolute path for security, so you choose the folder through the picker rather than typing it. This uses the File System Access API and is

available in Chrome and Edge; on other browsers exports fall back to Downloads.

Foundational Samples

Foundational Samples is a one-click sample-pack preset. Click the gold **Foundational Samples** chip (or type "export Foundational Samples"), and the AI-Engineer proposes this set, mapped to your actual track names:

- **Full Mix** - the whole song
- **Drumless Mix** - drums and percussion muted
- **Vocal Mix** - vocals only
- **Bass and Drums** - bass and drums only
- **Each track as its own stem**

It **confirms first**: it lists exactly what it will export and waits. Reply **"yes, all of them"** to export the batch, or exclude items, for example **"exclude the piano stem"**, and it drops those and exports the rest. Set an export folder first so the whole pack lands in one place.

Name your engineer

Click the **Name** chip and give the assistant whatever name you want, up to 24 characters. The name replaces "AI-Engineer" on the panel tab and in the chat input placeholder, so the DAW reads as *your* studio.

This is a **subscriber feature**. On the free tier the chip shows a Subscribe link instead of the rename card. Save an empty name to go back to the default. The name is stored in your browser, so it stays with the machine you set it on.

Talk to it, and let it talk back

Two buttons sit next to the chat input:

- **Mic** - speak instead of typing. Your words fill the input as you talk, and the message sends once you stop speaking, so you are not firing off half a sentence.
- **Speaker** - reads replies aloud. Toggle it on or off; the setting is remembered between sessions.
- **Voice** - the small chevron next to the speaker opens a voice picker. Click any voice to hear a sample and set it, whether or not replies are switched on.

Both use the browser's built-in speech engine, so there is nothing to install and nothing extra to pay for.

Browser support. Reading replies aloud works in every current browser. Voice **input** needs **Chrome, Edge or Safari**; it is **not available in Firefox**. Where it is missing the buttons are simply not shown, so nothing breaks.

About the voices. The list comes from your browser and operating system, not from us, so it differs per machine. Voices marked with an asterisk are network voices and are usually the more natural ones. The DAW picks the best available English voice for you rather than the system default, which on Windows is a noticeably flat robotic voice. Pick a different one any time.

Tips

- Refer to tracks by name. If your stems have unusual names, glance at the proposed list before confirming a batch.
- Everything it changes is undoable - experiment freely, then Ctrl+Z.
- Chat is free; only generation and Replace Audio spend credits.
- Select the region *before* you open Replace Audio, so the card can confirm the target back to you.
- For a clean sample pack: set an export folder, then run Foundational Samples.

Where to go next

- Mixing - the controls the AI-Engineer drives
- The Library - what it generates
- Exporting - the export formats and folder