

MOD-07 / VISUAL

SUPERNOVA


VISUAL

FREE • A6PLV3

VST3 • AU • STANDALONE

SUPERNOVA turns any image or video into a living cloud of **262,144 GPU particles** that move, glow and burst with the audio. The differentiator: it ships as a **VST3 + AU** plugin inside your DAW and as a **standalone app** that captures your Mac's **system audio with zero drivers** — point it at Spotify, YouTube, Serato or Ableton and it just reacts. No BlackHole, no routing. And the audio itself is **bit-exact pass-through**: SUPERNOVA never colours your sound.

// ANATOMY



- Source** — where the sound comes from: the DAW track, this Mac's system audio, or a mic
- In & level** — the input meter and its trim: how hard the audio drives the visuals
- Syphon & fullscreen** — the picture out to OBS / Resolume / VDMX, or onto any display
- Mic / MIDI** — MIDI-learn on any knob, CC control, and notes that fire triggers
- Load** — drop in an image, a folder of images (photo sequence) or a video
- Export** — record the visuals to MP4 / H.264 with the audio muxed in sync
- Worlds** — ◀ name ▶ walks the 50 factory worlds; ⌘ opens the browser grid
- The field** — 262,144 GPU particles sampling your image and moving with the sound
- Media strip** — the whole session in order: click a tile to cue it, drag to reorder
- Intensity** — the master "life" fader: from a calm shimmer to a full storm
- Matter** — what the particles are: density, scatter, size, trails, outout, links
- Camera** — where you watch from: rot X/Y, orbit, rotate, form, depth, kaleido
- Colour** — the look: palette ramp, amount, hue, saturation, hue cycle, glow

// CONTROLS — COMPLETE REFERENCE

SOUND

LOAD & WORLDS — Drag in a photo, a photo sequence or a video, then browse **50 WORLDS** — presets that swap the physics, geometry and colour in one click.

INTENSITY — The master **LIFE** fader. One knob takes the field from a calm shimmer to a full storm — how hard the whole system reacts.

THE FOUR DOMAINS — Every control lives in one of four groups — **MOVEMENT** (flow, chaos, gravity), **MATTER** (size, trails, plexus, glyphs), **CAMERA** (3D depth, rotate, orbit) and **COLOR**.

COLOR LAB & 3D — 17 curated look ramps baked in OKLab, plus a real 3D cutout of your subject you can rotate — order-independent additive, no z-sort.

TRIGGERS — MIDI notes fire events — a note explodes the field or shoots a directional ray. Play the visuals from your keyboard, in time.

LFOs — Four LFOs breathe any of **12 TARGETS** — hue, size, chaos, orbit, trails. **21 BEAT DIVISIONS**, each with its triplet and dotted, or a free-running **HZ** mode that keeps moving with the transport stopped. Per LFO: bi/unipolar, phase and retrigger.

OUTPUT & SYPHON — Send the frame to a window, fullscreen on any display, or straight into OBS / Resolume via **SYPHON** (macOS).

EXPORT WITH SOUND — Record your visuals to **MP4 / H.264** with the audio muxed **IN SYNC** — four aspect presets and any length from **1 TO 600 S**. What drives the clip is the last **30 S** you played: a longer export repeats that window, picture and sound together, crossfaded over 10 ms so the loop point does not tick.

GLOBAL & I/O

AUDIO IN — In the DAW it's a **BIT-EXACT READ-ONLY TAP**; the standalone app hears your whole Mac with **ZERO DRIVERS** — Core Audio taps on macOS 14.2+, ScreenCaptureKit on 13–14.1. No BlackHole.

VISUAL OUT — Send the frame to a window, fullscreen on any display, or into OBS / Resolume / VDMX through **SYPHON** (macOS).

MIDI — MIDI-learn on any knob, MIDI CC control, and notes that fire triggers — play the visuals from your keyboard, in time.

WORLDS — 50 built-in worlds swap physics, geometry and colour in one click; drop in an image, photo sequence or video to reshape the field.

VARIATION — Per-domain hotknobs, a global Variation and a Mutate button explore thousands of deterministic variations of the active world.

BYPASS — Freezes the reactivity; the audio always passes through bit-exact, bypassed or not.

// PARAMETERS • EXACT RANGES


		RANGE	DEFAULT
INTENSITY		0-100%	50%
CHAOS		0-100%	30%
SIZE	0-100% • 0.5-4.0 px		40%
GLOW • BLOOM		0-100%	50%
CUTOUT • SUBJECT		0-100%	0%
DENSITY		1-100%	100%
SCATTER • IMAGE↔CLOUD		0-100%	0%
GRAVITY	-100...+100% • ±0.60		0%
BLAST • EXPLOSION FORCE		0-100%	50%
SHIMMER • HIGHS		0-100%	50%
BREATHE • RMS		0-100%	50%
MOTION	Contours • Matter • Wave • Vortices • Radial	Contours	
SHAPE	Dot • Disc • Ring • Dash • Tri • Quad • Spark	Dot	
TRAILS		0-100%	0%
LINKS • PLEXUS		0-100%	0%
SPEED	×0.25-×4 • log	×1	50%
PUMP • KICK		0-100%	30%

		RANGE	DEFAULT
ROTATE		-100...+100% • ±30°/s	0°/s
DEPTH • 3D		0-100%	0%
ROT X / ROT Y		-180...+180°	0°
ORBIT		-100...+100% • ±45°/s	0°/s
KALEIDO	0ff • 2 • 4 • 6 • 8		0ff
SATURATION	0-100% • 50 = neutral		50%
HUE		-180...+180°	0°
HUE CYCLE		-100...+100% • ±60°/s	0°/s
FIGURE	Image • Sphere • Spiral • Rings • Grid • Helix	Image	
FORM • IMAGE↔FIGURE		0-100%	100%
PALETTE • LOOK RAMP		Original + 16 looks	Original
COLOR AMOUNT		0-100%	100%
BACKGROUND • PAPER		0-100%	0%
VARIATION • CURATED		0-100%	0%
EXPLODE • MIDI TRIGGER		0ff • 0n	0ff
BYPASS		0ff • 0n	0ff

// HOW IT WORKS

A real engine, not a filter.

SUPERNOVA is a GPU synthesizer for light — HDR compositing, filmic tonemapping and per-preset physics — wrapped around an audio path that stays untouched.



// UNDER THE HOOD

262,144 GPU PARTICLES

A 512×512 grid — a quarter-million particles — samples your image or video and is advected every frame on the GPU by an analysis of the sound (onsets, frequency bands, RMS). Each of the 50 worlds is a different force field — flow, vortices, gravity, radial blast — coupled to its own geometry, so a preset changes how the field MOVES, not just how it looks.

VISION CUTOUT + 3D SCULPTURE

Apple's Vision framework computes a saliency mask of your subject in real time; CUTOOUT dissolves the background so only the subject's particles remain, and the mask feeds a real 3D cutout — a rotatable sculpture of the subject (additive, order-independent, no z-sort) that you can tilt and orbit in depth.

BIT-EXACT AUDIO TAP

As a plug-in the audio is a READ-ONLY analysis tap: out equals in, sample for sample, at 0 samples latency. SUPERNOVA measures your sound to drive the light and never touches the signal — it is not a DSP effect, adds no EQ, dynamics, noise or colour, and is safe anywhere in the chain.

HDR LINEAR → ACES FILMIC

Every particle is accumulated in a linear, additive HDR buffer — order-independent, no z-sort — and only then resolved through an ACES filmic tonemap. Highlights bloom and roll off like film instead of clipping to a flat white, so a dense core stays luminous and coloured rather than blowing out to paper.

FOUR LFOs, ON EVERY FRAME

Four LFOs modulate any of 12 targets — hue, saturation, intensity, size, glow, chaos, rotate, orbit, scatter, trails, links, depth. RATE covers 21 beat divisions from a sixteenth to four bars, each with its triplet (-T) and dotted (-D), plus a FREE Hz mode (0.05–20 Hz) that runs off the clock instead of the tempo — so it keeps moving with the transport stopped or with no host. Six waveform shapes, and per LFO: BI/UNI polarity, PHASE 0–360° and RETRIG. At 100% depth an LFO sweeps exactly the target's range and is clamped to it, and the value is recalculated on every rendered frame — a fast LFO reads as a wave, not a staircase.

EXPORT WITH SOUND

Record the visuals to MP4 / H.264 in four aspect presets with the audio muxed IN SYNC — a 30 s real-time ring buffer of the incoming audio is written to an AAC track aligned to the frames — and any length from 1 to 600 s. Picture and sound come from that SAME 30 s window, so a longer clip repeats it with both locked together — the take loops, but it never drifts out of sync. Thirty seconds is long enough that a 30–45 s piece no longer brings the same bar back, and the loop point is crossfaded over 10 ms instead of cutting, so the seam does not tick.

MULTI-SCALE BLOOM & OPTICS


The HDR frame drives a multi-scale bloom — several blurred octaves summed — for a halo that reads as real light, followed by a cinematic optics pass: film grain, ordered dither to kill banding, and a touch of chromatic aberration at the edges. It sits AFTER the tonemap, grading the finished image rather than decorating it.

DRIVER-FREE SYSTEM AUDIO

The standalone app captures your Mac's own output with zero drivers and no BlackHole routing — point it at Spotify, YouTube, Serato or a DAW and it just reacts. On macOS 14.2+ it does this with Core Audio process taps, so the permission macOS asks for is the small, exact one — System Audio Recording — and it is asked when you PICK the source, not when the app opens. That TCC gate is how the OS authorises system-audio capture: it is NOT the microphone, not a single pixel of your screen is read, and nothing is recorded. On macOS 13–14.1 those taps do not exist yet, so the app falls back to ScreenCaptureKit and macOS gates it behind the older Screen Recording switch — same capture, same privacy, simply the only route those versions offer. System Audio needs macOS 13+; on 11–12 you pick an input device instead.

// THE STANDALONE APP

The same engine outside the DAW: a window of its own that hears everything your Mac plays.



1 A window of its own — the title bar says it: this is an app, no DAW required

2 System Audio (this Mac) — it hears whatever is playing: Spotify, YouTube, Serato, a DAW

3 Syphon & fullscreen — feed OBS / Resolume / VDMX, or throw it on the projector

4 Mic / MIDI — pick an input device instead, and play the visuals from a keyboard

// THE WORLDS

Every world is a preset — its own physics, geometry and colour ramp. One click swaps the whole system. Ten of the 50 built in:



// MEASURED PERFORMANCE

Particle engine	262,144 • 512×512 grid	Frame rate	120 fps sustained • GPU 2–6 ms/frame
Audio → visual	8.9 ms p50 @512 smp / 60 Hz • 6.5 ms @120 Hz	GPU pipeline	Apple Metal • HDR linear
Compositing	linear additive → ACES filmic	Colour ramps	17 OKLab looks (Original + 16)
Audio latency	0 samples • zero-latency	Audio path	bit-exact • out = in (read-only tap)
System audio (app)	Core Audio taps 14.2+ • SCK 13–14.1	Offscreen render	deterministic • byte-exact goldens
Export	MP4 / H.264 • 30 s ring looped in sync		

Conditions — frame rate and GPU time measured on an Apple Silicon Mac with a ProMotion display, canvas 1332×750, all 262,144 particles live (04-09-2026); on a 60 Hz display the ceiling is 60 fps. Audio-to-visual latency measured with 512-sample blocks: median 8.9 ms at a 60 Hz refresh, 6.5 ms at 120 Hz. The particle count is fixed by design and the offscreen render is deterministic, verified byte-exact against golden frames. The audio path is a read-only pass-through: 0 samples latency, out == in — so THD, SNR and frequency response do not apply. System-audio capture (standalone app) needs macOS 13+: Core Audio process taps on 14.2+, ScreenCaptureKit on 13–14.1.

// LATENCY • PER SAMPLE RATE

	44.1 KHZ	48 KHZ	96 KHZ
REPORTED LATENCY	0 samples / 0.0 ms	0 samples / 0.0 ms	0 samples / 0.0 ms

// SPECIFICATIONS

SYSTEM • FORMATS		AUDIO • PROJECT	
FORMATS	VST3 • AU • Standalone app	I/O	stereo in → out • bit-exact tap
SYSTEMS	macOS 11+ (Metal) • System Audio: 13+	LATENCY	0 samples • zero-latency
ARCHITECTURE	Universal • Apple Silicon + Intel	COLOUR	17 OKLab look ramps • ACES
ENGINE	262,144 GPU particles • HDR linear	OUTPUT	window • fullscreen • Syphon
CAPTURE	System audio, driver-free • macOS 13+	EXPORT	MP4 / H.264 • 1–600 s • 30 s window looped
MIDI	Notes • CC • MIDI-learn	PRESETS	50 worlds (factory)
MODULATION	4 LF0s • 21 divisions • free Hz • 12 targets	LICENSE	AGPLv3 • open-source

// TRANSPARENCY & HONESTY

The audio is bit-exact pass-through — a read-only analysis tap, out equals in sample for sample, at 0 samples latency. SUPERNOVA is not a DSP effect: it adds no EQ, dynamics, noise, harmonics or colour, so THD+N, frequency response and SNR do not apply — the signal is unmodified. Everything measured here is the visual engine. The plug-ins and standalone app are macOS-only (Apple Metal); there is no Windows build.

// INSTALL (FREE & OPEN-SOURCE)

Free & open-source (AGPLv3), macOS only. Grab the `.pkg` installer from ovniaudio.com and double-click: it is **signed and notarized by Apple (Developer ID)**, so it just installs — the **VST3**, the **AU** and the **standalone app**. No Gatekeeper detour. The first time you pick **System Audio**, macOS asks for permission to listen to your Mac’s audio: turn SUPERNOVA on in **System Settings › Privacy & Security › System Audio Recording** and hit **REOPEN**. Nothing is recorded. **AGPLv3 • full source at github.com/ovniaudio/ovni**.