The Spectrum of Creation

A framework for authorship, rights and distribution in the age of AI music

Whitepaper by Soundverse AI · July 2026 · www.soundverse.ai
Co-authored by Sourabh Pateriya, Riley Williams and Linnea Sundberg

Abstract

The public conversation about AI music is stuck on a false binary: “real” music versus “AI” music. In practice, creation runs along a continuous spectrum, from fully machine-generated to fully human-performed, and almost everything interesting happens in between. A finished track, moreover, is rarely the product of a single action: it is typically the outcome of many steps (generate, edit, layer, swap, extend) performed through a graphical studio, a conversational AI agent, or both, and its rights follow from that entire chain.

This paper introduces the **Spectrum of Creation**: a shared vocabulary of creation modes, a four-axis model of musical provenance, a three-tier clearance system, and a single composition rule, the weakest-link rule, that together answer, for any track made with AI, three questions the industry has struggled to answer consistently: who made this, can it be released, and where does the value belong? The framework is implemented in production across Soundverse, an AI-first music creation platform serving over three million creators, and is published here as a proposed common standard for platforms, distributors, rights-holders and creators.

About Authors

1. **Sourabh Pateriya**: Sourabh Pateriya is the Founder & CEO of Soundverse Inc., previously leading product and design at Spotify, Tobii Tech, and Samsung Research. He is a music-tech innovator with 10+ patents and a member of the Forbes Tech Council, co-designer of Samsung Bixby. At Spotify, he built music creation experiences and co-invented Basic Pitch, an open-source ML model for audio-to-MIDI transcription. As an electronic music producer (BlueNucleus), he blends creativity with deep technical insight. Sourabh spearheads Soundverse’s mission to build ethical and controllable AI music infrastructure, focusing on data attribution, artist-centric monetization, and tiered licensing models.
2. **Riley Williams**: Riley Williams is the Co-Founder and CTO of Soundverse, with prior experience as a research engineer at Meta Reality Labs and applied scientist at Amazon Photos. He has over a decade of expertise in computer vision and multimodal machine learning, paired with his background as an electronic music producer. Riley designs and scales Soundverse’s AI generation and attribution systems, ensuring controllability and reliability. His work centers on embedding licensing and rights management directly into the platform’s technical core.
3. **Linnea Sundberg**: Linnea Sundberg is a Soundverse Board Member and VP of Corporate Development at UnitedMasters, with previously senior leadership roles at Spotify and Splice. At Spotify, she helped shape global artist compensation and creator-facing business models, aligning incentives between the platform and rights holders. She brings 15 years of experience in music, tech, and creator-economy business models. Linnea specializes in rights management, IP strategy, and creator monetization. At Soundverse, she crafts fair, transparent licensing systems that balance innovation with artist and label interests.

Summary

Claims defended in the pages that follow:

- **AI music is a spectrum, not a binary.** Between “fully AI” and “fully human” lie dozens of distinct creation modes: a writer whose lyrics are composed by AI, a producer who hums a melody into a full arrangement, an artist who trains a model on their own voice. Each deserves its own name, and this paper gives each one a code and a definition.
- **“How much AI” and “can it be released” are different questions.** Position on the spectrum measures authorship. Clearance measures whether any third party's rights are entangled in the track. A fully human song built on an uncleared sample is unreleasable; a fully AI song made from clean inputs is not. Conflating the two axes is the most common error in the AI-music debate.
- **Provenance can be modelled with four axes.** Every track's origin story reduces to: where the lyrics came from, where the composition came from, whose voice is on it, and where the source audio came from. Uncleared third-party material on any axis, known as a “taint”, is what makes a track unreleasable.
- **Overall rights are represented by least permissive components.** Real tracks are made in many steps. The finished work's clearance is the minimum across every step in its lineage: one uncleared input taints the whole chain, and no later step can cure an earlier one. This single rule makes multi-step, multi-tool, multi-collaborator creation tractable.
- **Value follows human contribution.** Purely machine-generated output may not be copyrightable at all; human expression is. The more a creator writes, performs, and shapes, the more ownable their work becomes. A well-designed platform economy should reward creators for moving up the spectrum, and Soundverse's does.

> *“The question is not whether a track is AI or human. The question is what, precisely, the human contributed, and whether anything in the track belongs to someone else.”*

1 · Why the binary fails

In the four years since generative music models became broadly capable, the industry has organised its thinking around a single question: is this track AI or not? Streaming platforms debate AI labels. Rights bodies debate AI registrability. Distributors debate AI catalogues. The framing assumes a clean line exists.

It does not. Consider five creators.

- One types a sentence and receives a finished song: lyrics, composition, and voice all machine-made.
- Another writes every word of her lyrics, then asks a model to compose and sing them.
- A third hums a melody he has carried for years, and AI arranges it into a produced record.
- A fourth spends a decade recording her own voice, trains a model on it, and releases songs sung by her own AI likeness.
- A fifth with deep technical skill, who chops, manipulates, or otherwise modifies an AI-generated stem.

All four “used AI.” To treat them identically (for credit, for rights, for revenue) is to erase distinctions that matter enormously to creators, to rights-holders, and to listeners.

Soundverse operates the full lifecycle from idea to audience (generation, editing, collaboration, attribution and distribution), which means we cannot leave these distinctions vague: our systems must decide, for every track, what it is and whether it can ship. This paper publishes the framework those systems run on. We offer it not as a legal treatise but as working infrastructure: a vocabulary precise enough to build products, contracts and policy on.

2 · The spectrum, and the orthogonality principle

Picture authorship as a continuous line. At the left end, the machine authors everything and the human contributes only intent. At the right end, the human authors everything and the machine is a passive utility: a recorder, an equaliser, a file converter. Every real track occupies a position on this line, determined by what the human actually contributed: words, melody, performance, arrangement, curation.

Independent of that position, every track also has a **clearance status**: is any third party's work or likeness entangled in it without permission? These two properties are **orthogonal**, and their independence is the foundation of everything that follows.

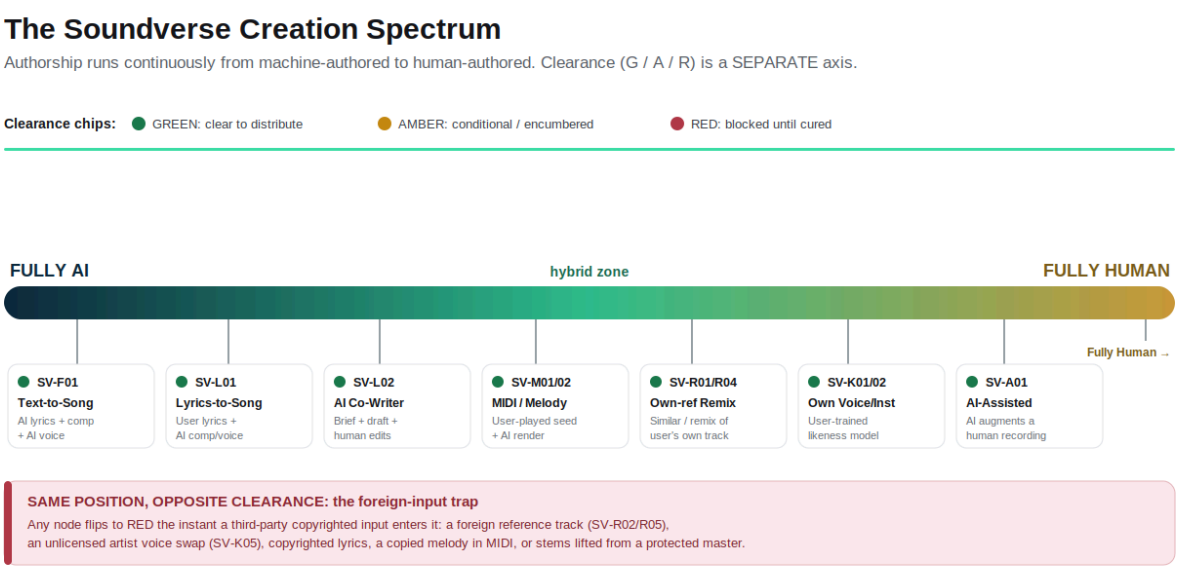


Figure 1. The Soundverse Creation Spectrum. Position measures authorship; colour measures clearance. The two are independent.

2.1 Four axes of provenance

Every track's origin reduces to four questions, each an independent axis:

Axis	Question	Values (cleanest → most encumbered)
L · Lyrics	Where did the words come from?	None · AI-written · creator-written · co-written · third-party copyrighted (LX)
C · Composition	Where did the music come from?	AI from prompt · creator MIDI · creator melody · AI on creator's own reference · AI on a copyrighted reference (CX)
V · Voice	Whose voice (or likeness) is on it?	Instrumental · AI-generic · creator's own trained voice · licensed artist model · unlicensed real-person likeness (VX)
S · Source	Where did any source audio come from?	Generated in-platform · creator's own original upload · third-party copyrighted upload (SX)

The values marked **X** share one property: an uncleared third-party right has entered the track. We call this a **taint**. Taints, not AI use, are what make tracks unreleasable.

2.2 Three clearance tiers

Tier	Meaning	What it means for the creator
GREEN	No third-party taint on any axis: the track is built from AI output, the creator's own original input, and in-platform generation.	Eligible for one-click distribution and full monetisation.
AMBER	Releasable, but with a condition attached: typically a licensed artist model (which carries an artist royalty, tracked by Soundverse TRACE), a documented consent, or an attestation.	Distributable once the condition is met; the encumbrance travels with the release.
RED	An uncleared third-party right is present: a copyrighted reference, sample, lyric, melody or voice.	Not eligible for distribution. Soundverse shows exactly which input caused it and how to cure it; the track can remain private.

What RED is, and is not

RED is not a punishment and not a judgment of quality. It is a statement of fact: something in this track belongs to someone who has not agreed to its use. Releasing it would expose the creator to takedowns and claims, and would degrade the standing of every other creator sharing the same distribution catalogue. The cure is always specific: remove the tainted input, replace it, or license it.

3 · How creative tools move provenance

A modern AI music platform is a toolbox, and each tool interacts with provenance in one of three ways: it originates provenance (creates a fresh axis value), passes it through (transforms without adding rights), or imports it (can carry a third-party right into the track). Understanding which is which turns rights reasoning from case-by-case argument into mechanical lookup.

Tool family	Behaviour	Provenance effect
Song / music generation from text	Originates	Fresh AI lyrics, composition, voice. Clean origin.
Lyric writing & co-writing	Originates	AI or co-written words; clean unless copyrighted text is pasted in.
Lyrics-to-song	Passes through	Honours the creator's words; composes around them.
MIDI / melody / motif to song	Imports	Clean if the seed is the creator's own; tainted if the seed encodes someone else's melody.
Reference-conditioned generation (“similar song”)	Imports	The reference is the risk: own reference clean; copyrighted reference tainted.
Style extraction → fresh generation	Passes through	Only abstract, unprotectable style travels (genre, mood, tempo feel).
Remixing	Imports	A remix is a derivative of its source: own source clean, copyrighted source tainted.
Stem separation & stem generation	Imports / originates	Separated stems inherit their source's status; newly generated stems are clean.
Voice swap & voice/instrument model training	Imports / originates	Own voice clean; licensed artist model conditional; unlicensed real-person likeness tainted at the model itself.
Editing (inpainting, region edits, extend, effects, mastering)	Passes through	Inherits the track's status. Editing preserves; it never cures.
Lyric translation with melody-preserving re-sing	Imports	Own song clean; a copyrighted song doubly tainted (translated lyric + preserved melody).
Copyright & watermark checking	Gate	Diagnostic instruments that support clearance decisions.

Two structural facts worth memorising

Editing never launders. No amount of inpainting, effects, extension or mastering converts a tainted track into a clean one. Transformation of an uncleared work is still a derivative of it.

The reference is the risk. Across every importing tool, liability enters through the uploaded object: the reference track, the source file, the target voice, the seed melody. Prompt text is cheap; uploads carry rights.

4 · A shared vocabulary: the creation-mode catalogue

The industry lacks names for the modes between the poles, so debates talk past each other. We propose the following taxonomy. Each mode has a stable code, a name, and a default clearance outcome. Platforms, distributors and rights-holders are welcome to adopt and extend it.

Cluster F: Fully AI

Code	Name	Definition	Default
SV-F01	**Genesis**	Text-to-song: a prompt yields AI lyrics, composition, and voice. The human contributes intent alone.	GREEN
SV-F02	**Aria Score**	Text-to-music: prompt to fully AI instrumental.	GREEN
SV-F03	**Scene Score**	Media-to-music: AI scores a creator-supplied image or video. Clean only if the media itself is the creator's or cleared.	GREEN

> **Clear to release, but perhaps not to own**
>
> Fully AI tracks carry no third-party taint, so nothing blocks their release. But most major copyright regimes require human authorship for protection: a purely machine-generated work may belong, in the copyright sense, to no one. “I can release it” and “I can stop others copying it” are different powers, a distinction we return to in Section 8, because it is where the economics of the spectrum come from.

Cluster L: Lyric-led

Code	Name	Definition	Default
SV-L01	**Penned**	The creator writes original lyrics; AI composes and performs them. The words are human-authored expression.	GREEN
SV-L02	**Co-Pen**	Creator and AI co-write: brief, draft, regeneration, line-by-line human editing.	GREEN
SV-L03	**Spoken**	Narration or a cappella singing generated from the creator's own text.	GREEN

The trap: pasting lyrics the creator did not write (a verse of a chart song, a poem in copyright) taints the lyric axis and makes the track RED no matter how original the music is. There is no safe number of borrowed lines.

Cluster M: Melodic and structural human input

Code	Name	Definition	Default
SV-M01	**Programmed**	The creator programs MIDI; AI renders it into a full arrangement.	GREEN
SV-M02a	**Hummed, own words**	The creator hums or plays a melody and writes the lyrics; AI produces the record.	GREEN
SV-M02b	**Hummed, AI words**	The creator's melody, AI-written lyrics.	GREEN
SV-M03	**Motif**	A short original motif seeds the generation.	GREEN

The trap: performing an input does not make it yours. MIDI that transcribes a protected melody, or a hum that reproduces a recognisable hook, taints the composition axis. The act of playing describes the performance, not the authorship.

Cluster R: Reference-conditioned

Code	Name	Definition	Default
SV-R01	**Echo (Own)**	“Similar song” generation from the creator's own earlier track: borrowing from oneself.	GREEN
SV-R02	**Echo (Foreign)**	Similar-song generation conditioned on a copyrighted reference.	AMBER/RED
SV-R03	**Restyle**	A reference is reduced to abstract style descriptors (genre, mood, tempo feel), and a fresh work is generated from text. Only unprotectable style travels.	GREEN/AMBER
SV-R04	**Remix (Own)**	AI remix of the creator's own track.	GREEN
SV-R05	**Remix (Foreign)**	AI remix of an uploaded copyrighted track: a derivative of both composition and recording.	RED

> **The most misunderstood case in AI music**
>
> A creator uploads a copyrighted track as the musical reference and writes entirely original lyrics on top. Intuition says the original words should count for something. The law of derivative works says otherwise: lyrics and composition are **separate rights**, and originality in one does not cure borrowing in the other. The track inherits the status of its weakest axis; here, RED. The productive path is Restyle (SV-R03): extract the abstract style you admire, and let the model compose fresh expression inside it. Style is free; expression is owned.

Cluster K: Likeness and trained models

Code	Name	Definition	Default
SV-K01	**Mirror Voice**	The creator trains a model on their own voice and releases songs sung by it. Self-owned likeness.	GREEN
SV-K02	**Mirror Instrument**	A model trained on the creator's own instrument, played back as a virtual instrument.	GREEN
SV-K03	**Artist DNA, licensed**	Generation through a consent-trained, revenue-sharing artist model. The artist is paid on every use, tracked by TRACE.	AMBER
SV-K04	**Voice swap, licensed**	A vocal re-voiced into a licensed artist model.	AMBER
SV-K05	**Voice swap, unlicensed**	A vocal re-voiced into a real artist's likeness without permission.	RED
SV-K06	**Clone, consented**	A voice clone of another real person with documented consent.	AMBER
SV-K07	**Clone, non-consented**	A clone of a real person without consent. The model itself is tainted; every output inherits it.	RED

> **Voice is a right of its own**

A voice can be infringed without a note being copied. Personality and publicity rights protect a person's voice, name and likeness independently of copyright; courts, including in India, have enforced them against unauthorised AI voice imitation. This is why the taint in K05/K07 attaches to the **model**, not merely to individual songs: a model trained on an unconsenting person poisons everything it produces. Consent-first artist models, where the artist opts in and earns from every generation, are the only sustainable architecture, and they are the architecture Soundverse builds.

Cluster A: AI-assisted, human base

Code	Name	Definition	Default
SV-A01	**Augment**	AI adds new instrumentation (a bass line, a pad, a drum layer) to the creator's own recording.	GREEN
SV-A02	**Surgery**	AI repairs or reimagines a selected region of the creator's own track.	GREEN
SV-A03	**Polish**	EQ, reverb, compression, mastering applied by AI.	GREEN
SV-A04	**Re-timbre**	The creator's own hummed or sung line becomes a realistic instrument (a sitar, a bansuri, a guitar) and can be layered into the arrangement.	GREEN
SV-A05	**Extend**	AI continues the creator's own track with new sections.	GREEN
SV-A0X	**Assist, foreign base**	Any of the above applied to a copyrighted upload; assistance inherits the base's taint.	RED

This cluster is where AI behaves most like every studio tool before it. A compressor never claimed authorship of the songs it compressed; an AI equaliser is no different. A human work stays fully the creator's through any amount of clean assistance. The only unrescuable case is a base that was never clean.

Cluster T: Localisation

Code	Name	Definition	Default
SV-T01	**Localize (Own)**	The creator's own song, translated and re-sung in another language with the melody preserved: one song, many markets.	GREEN
SV-T02	**Localize (Foreign)**	The same operation on a copyrighted song: the translation is a derivative of the lyric, and preserving the melody reproduces the composition.	RED

5 · Multi-step creation and the weakest-link rule

Real tracks are rarely one shot. A creator generates, inpaints a chorus, adds a stem, swaps a voice, extends an outro, masters, sometimes conversationally with an AI agent, sometimes hands-on in a multi-track studio. Every operation appends to the track's **provenance ledger**, and the finished work's clearance follows from the whole ledger by a single rule:

> *Final clearance equals the minimum clearance of every step in the lineage. One tainted step taints the chain, and no later step can restore what an earlier step lost.*

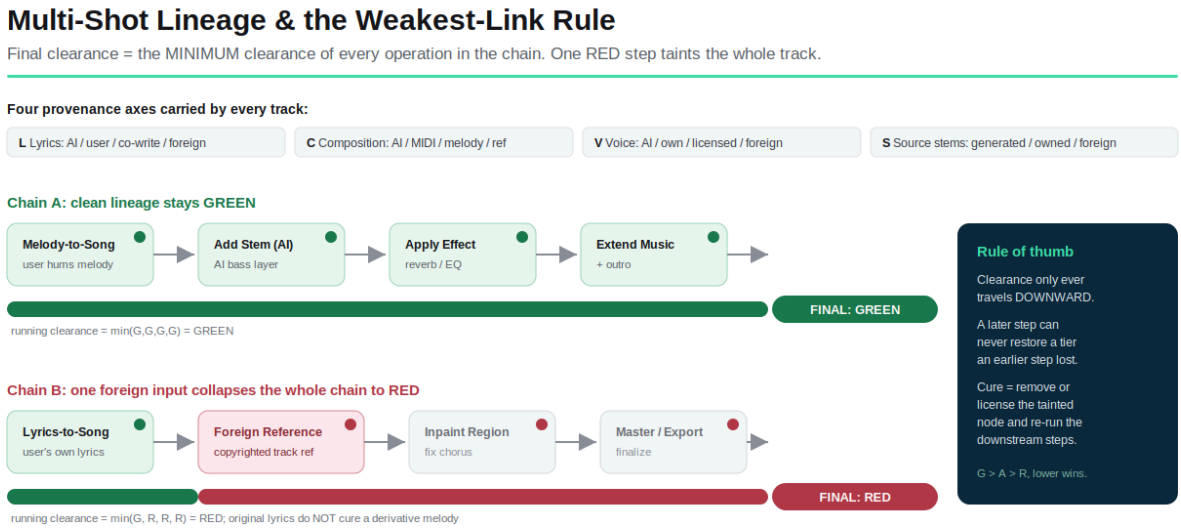


Figure 2. Two lineages. A clean chain stays GREEN end to end; a single foreign reference collapses an otherwise-original chain to RED.

The rule sounds strict, and it is, deliberately. Its strictness is what makes it usable: any participant in the chain, human or machine, can compute a track's status locally, without global knowledge, by tracking one number that only ever goes down. It also implies the correct cure procedure: find the lowest node, remove or license it, and regenerate everything downstream of it. Patching later steps achieves nothing.

5.1 Worked lineages

Chain	Ledger	Final
Hum → own lyrics → AI stem → master	G → G → G → G	GREEN
AI song → inpaint → extend → swap to own voice	G → G → G → G	GREEN
AI song → swap to licensed artist voice	G → A (artist royalty attaches)	AMBER
Own lyrics → copyrighted reference → inpaint → master	G → R → R → R	RED
Own track → stem lifted from a hit → extend	G → R → R	RED

6 · The hard cases

Fourteen situations that account for most genuine confusion. Each pairs the intuition with the correct analysis.

#	Situation	The analysis	Outcome
1	**Copyrighted reference, original lyrics**	Separate rights; original words don't cure a borrowed composition. Use style extraction instead.	RED
2	**MIDI transcribing a famous melody**	Performing an input is not authoring it.	RED
3	**Humming a copyrighted hook**	Same principle, different instrument.	RED
4	**Reusing stems separated from a hit**	Stems are derivatives of the master; your added parts don't cleanse the borrowed one.	RED
5	**Remixing an uploaded copyrighted track**	Derivative of composition and recording; “transformative” is not self-certified.	RED
6	**Voice-swapping to a real artist, no licence**	Publicity/personality rights are infringed even if no recording is copied.	RED
7	**A model trained on a non-consenting person**	The taint lives in the model; every output inherits it.	RED
8	**“Make it sound like [living artist]”**	A fresh composition steered at a named persona raises soundalike and publicity concerns.	AMBER
9	**Translating and re-singing a copyrighted song**	Derivative twice over: the translated lyric and the preserved melody.	RED
10	**A generic AI voice that happens to resemble a star**	Defensible if not trained on or steered toward that person; steering re-introduces the risk.	GREEN/AMBER
11	**A clean similarity scan on a foreign-reference track**	No scanner is complete; a clean scan supports, but does not manufacture, clearance.	AMBER
12	**A watermark detected on an upload**	The audio has a history, AI-origin elsewhere or protected, that must be resolved first.	AMBER/RED
13	**A “style” transfer that carries the hook through**	If recognisable expression survives, it was never just style.	AMBER/RED
14	**A track with co-writers or partners**	Multiple humans hold rights; distribution needs all of them to agree.	AMBER

7 · From clearance to distribution

Distribution is where the framework becomes consequential: the moment a track leaves the platform for streaming services, social platforms, or commercial licensees. Soundverse applies the tiers as follows.

Tier	Distribution	Condition	Monetisation
GREEN	One-click eligible.	Standard originality attestation.	Full.
AMBER	Eligible after a gate.	Licence or consent on file; encumbrance (e.g. artist royalty) accepted; attribution logged via TRACE.	Full, net of the encumbrance.
RED	Not eligible.	A specific cure path is shown: remove, replace, or license the tainted input. The track may remain private.	None until cured.

Two properties of this gate deserve emphasis. First, it is **transparent**: a blocked creator is never told merely "no," but exactly which input caused the block and what would cure it, because the provenance ledger records where every element came from. Second, it is **collective protection**: a distribution catalogue is a shared reputation. Streaming services increasingly police uncleared and mass-generated uploads at the catalogue level, so every infringing track that slips through degrades the standing, and ultimately the reach, of every clean creator alongside it. A strict gate is the most creator-friendly gate.

Every release is stamped with its creation-mode code, its provenance ledger, its tier and its encumbrances, so that attribution, royalty runs, audits and any later dispute can trace a straight line from the stream back to the session.

8 · Where the value sits: economics of the spectrum

The spectrum is not only a rights map; it is a value map. The key legal fact from Section 4, that purely machine-generated work may be unprotectable while human expression is fully protectable, has a direct economic translation: **the more a creator contributes, the more ownable their work is, and the more of its value they should command.**

Soundverse's revenue model is built on that principle. Its exact terms live in our Terms of Service and partner agreements; the shape is as follows:

- **Streaming revenue flows down a waterfall.** The streaming service takes its platform share (industry-standard, roughly 30%); the distributor takes a competitive margin for carriage, attribution and catalogue protection; the remainder is the rights pool.
- **The creator's share of the rights pool rises with their contribution.** On the Soundverse model it rises steadily across the spectrum: from an even split on fully AI-generated work, where the durable value lies chiefly in the models and infrastructure, to a substantial creator majority (up to 90%) on fully human work, where the platform's role is carriage alone. Writing the words, humming the melody, performing, arranging: every step up the spectrum is a step up in the creator's share.
- **Consent-trained artists are paid on every use.** Where a track uses a licensed Artist DNA model, a royalty for the artist comes off the top of the rights pool, logged per-generation by TRACE. Likeness is compensated as a first-class right, not an afterthought.
- **Collection societies keep their role.** Human-authored elements (lyrics, melodies, arrangements) remain registrable works that performing-rights organisations can administer for their usual fee. The spectrum does not route around the existing rights economy; it tells it, precisely, which parts of a hybrid work it applies to.

Why the incentive points the right way

A revenue curve that rewards human contribution does three things at once: it compensates creators in proportion to what they can actually own; it steers the catalogue toward distinctive, protectable work rather than interchangeable output; and it aligns the platform's commercial interest with the creator's creative growth. The platform earns most where its models did the work, least where the human did, **which is exactly the incentive a creator should want their platform to have.**

9 · Releasable versus ownable

A final distinction ties the framework together. This paper has treated two questions separately, and the separation matters most at the far AI end of the spectrum:

Question	Governed by	Fully AI work	Human-contributed work
Can it be released?	Clearance (taints)	Yes, if built from clean inputs.	Yes, if built from clean inputs.
Can it be owned?	Human authorship	Often not: most regimes deny copyright to purely machine output.	Yes: the human expression (words, melody, performance, arrangement) is protectable.

For creators, the practical guidance is one sentence: **everything you add makes it more yours.** Write the hook, hum the line, sing the verse, shape the arrangement: each contribution converts unprotectable output into ownable work, raises your share of its value, and makes the track harder to imitate. The spectrum is not a scale of legitimacy. It is a ladder of ownership, and it is open at every rung.

10 · Conclusion: provenance as infrastructure

The music industry has resolved every previous technological disruption (recording, sampling, synthesis, streaming) the same way: not by prohibition, but by building the rights infrastructure that let the new practice pay its participants fairly. AI music is at that juncture now, and the missing piece is not better models. It is provenance: a shared, machine-readable answer to who made this, from what, with permission from whom.

The Spectrum of Creation is Soundverse's contribution to that infrastructure, published so that others can adopt, criticise and improve it. The framework runs in production today, in our provenance ledgers, our clearance gates, our TRACE attribution system, and our consent-first artist models, across a community of more than three million creators and fifteen million songs. We believe its core commitments travel well beyond one platform: **name the modes precisely; keep authorship and clearance orthogonal; compose rights by the weakest link; pay likeness on every use; and reward the human contribution.** An industry that adopts them will spend less time litigating the binary and more time doing what the technology finally makes possible: letting anyone with an idea carry it all the way to an audience.

About Soundverse AI

Soundverse is an AI-first music creation platform covering the full lifecycle from idea to audience: conversational creation with Agent One, a collaborative audio-visual studio, consent-trained Artist DNA models, TRACE attribution and fair-pay infrastructure, distribution, and a B2B enterprise API. Learn more at www.soundverse.ai.