Contents

Sound Through the Night

The Threshold-Loop-Threshold Architecture of LullaWave's Nightscapes

A design paper grounded in evidence and tradition.

Preface

A note from the founder.

I'm a musician. I built LullaWave because, after years of making music, I noticed something I couldn't unsee: the audio that's marketed for sleep — the white-noise generators, the "spa music," the looping rain tracks — isn't really music. It's sound design. It can mask the world, but it can't *hold* you. And the audio that *is* music — the playlists labeled "sleep," the meditation soundtracks — is mostly made for streaming services, optimized for skips and saves rather than for what happens to a body across an actual night. There's a gap between what we know about how sound shapes the nervous system and what's actually being made for people who can't sleep.

I started building LullaWave's soundscapes by following the technical brief I'd written. They were correct on paper. They were lifeless to listen to. The drones held; the wind moved; the bells punctuated. None of it felt like anything. So I went looking — not for production tips, but for what the actual evidence and the actual cultural record say about how humans across history have used sound at night.

What I found surprised me. The clinical research is more conclusive than I expected on some questions (slow tempo, soft attack, instrumental, no lyrics — sleep music has a remarkably consistent acoustic profile across hundreds of studies) and less conclusive than I expected on others (continuous audio across a whole night is genuinely under-studied). And the cultural record — Compline, lullaby, *suantraí*, the Bedtime Shema, evening *sandhya*, the camp at fire — is doing something the clinical literature mostly hasn't named yet: humans across separate civilizations have, for thousands of years, used a *threshold-loop-threshold* structure. Music to mark the entry into sleep. Ambient sound through the night. Music to mark the exit. We didn't invent this pattern. We're translating it.

This paper is what we found, why it shaped the design we built, and where we're honest about what we don't know. It's longer than it needs to be for marketing copy and shorter than it needs to be for a textbook. We wrote it for the person who reads carefully — a sleep scientist, a music therapist, a sound designer, an insomniac who wants to understand what they're listening to before they trust it. We wanted it to stand up to that reader.

The body of the paper is in scholarly first-person plural — the voice of the LullaWave team. I'll return briefly at the end.

— *T*

Scope and Limits

Before we begin, the limits.

This paper describes a designed audio experience grounded in empirical research and cross-cultural sound traditions. It is **not** a clinical trial of LullaWave's nightscapes (none has been conducted), and it makes no claims about LullaWave functioning as a treatment for insomnia, sleep disorders, anxiety, depression, grief, or chronic pain. The clinical literature we cite establishes principles about *how* music and ambient sound affect sleep generally; we apply those principles to design choices. We do not extrapolate from the literature to claims about LullaWave's specific efficacy for any condition.

Where the evidence is strong, we say so. Where it is mixed, we say so. Where it is thin or contested, we name what's missing. We have tried throughout to distinguish what is well-documented from what is plausible-but-unproven, and to distinguish authentic cultural traditions from modern Western inventions that present themselves as ancient. The latter point matters: a meaningful fraction of what is sold as "ancient sound healing" — Tibetan singing bowls used meditatively, crystal bowls, gong baths, chakra-tuned tuning forks — is a Western construction from the late twentieth century, not the lineage it claims (Helffer 1994; Godwin 1995). LullaWave borrows neither the framing nor the marketing of those traditions, even where surface acoustic similarities exist.

We also want to be clear about audience. This paper is for the reader who wants to know what's underneath LullaWave's design. It is not a substitute for working with a sleep specialist, a therapist, a physician, or any other professional whose support a reader may need.

---

1. Introduction: A Modern Question with an Ancient Answer

How does a human being fall asleep? How does a human being stay asleep?

These look like one question, but they aren't. The first is about the threshold — the few minutes during which a person crosses from waking consciousness into sleep, a transition that under normal conditions takes between five and twenty minutes (Carskadon & Dement 2017; Roehrs et al. 2017). The second is about the rest of the night — the seven or eight hours during which the sleeping body cycles through stages of deep sleep, light sleep, and REM, and during which auditory input continues to be processed by a still-active auditory cortex even as conscious awareness has receded (Hayat et al. 2022; Strauss et al. 2015).

Most clinical sleep-music research has answered the first question. Music interventions reduce sleep-onset latency by 13–27 minutes on average across multiple meta-analyses (Jespersen et al. 2022; Wang et al. 2014). The acoustic features of effective sleep music are remarkably consistent across studies and cultures: slow tempo (60–80 BPM), soft attacks, instrumental, harmonically simple, lacking lyrics or sudden dynamic events (Dickson & Schubert 2019; Trahan et al. 2018; Kirk & Timmers 2025). We know what helps a person fall asleep.

The second question is largely unanswered by the clinical literature. The studies that exist on continuous all-night audio are mostly hospital-noise masking studies (Stanchina et al. 2005; Farokhnezhad Afshar et al. 2016) — concerned with whether ambient sound prevents arousal, not whether it actively supports sleep maintenance. Studies on whether music played throughout the night helps or harms sleep architecture are few and the findings are mixed, with one important cautionary signal: instrumental music played at bedtime increased the incidence of nighttime earworms (involuntary musical imagery) and worsened polysomnography-measured sleep in a 2021 study (Scullin et al. 2021). The all-night question is genuinely open.

LullaWave's product premise — that audio plays continuously through the night — therefore sits in a space the clinical research has not fully mapped. To design for that space responsibly, we need a second source of evidence beyond the clinical literature. We argue in this paper that the cross-cultural and historical record of how humans have actually slept provides that second source — and that it converges with the clinical evidence on a clear architectural principle: *sound at the thresholds of sleep, ambient through the middle.*

We call this the **threshold-loop-threshold** architecture. It is the structure of every nightscape in LullaWave's catalog. We did not invent it; we observed it.

The rest of this paper is organized as follows. Section 2 presents the threshold-loop-threshold architecture as it appears across separate human civilizations and across deep evolutionary time, drawing on ethnomusicology, anthropology, comparative religious studies, and soundscape ecology. Section 3 reviews the clinical sleep-music research and identifies the acoustic features the evidence consistently supports. Section 4 details how each of LullaWave's four nightscape lanes — Emotional, Somatic, Deep Rest, and Mental — implements the architecture for a specific user state, with the design rationale grounded in the literature. Section 5 addresses the production constraints and the specific decisions that follow from them. Section 6 names what remains uncertain. Section 7 is a brief closing reflection from the founder.

We have written for the careful reader. Citations are inline in the academic style, with the bibliography at the end.

---

2. The Threshold-Loop-Threshold Architecture as Cross-Cultural Pattern

2.1 What we mean by threshold

A threshold, in the sense we use it here, is a structured transition between two states of consciousness. The transition into sleep is one such threshold; the transition out of sleep into morning waking is another. What is striking, when one looks across cultures, is how often these thresholds are *sung.*

Consider Compline — the seventh and final canonical hour in the Christian monastic Liturgy of the Hours, codified in the Rule of Saint Benedict around 540 CE and stable across the Western liturgical tradition for nearly fifteen hundred years (Hiley 1993; Crocker 2000). Compline is sung after the evening meal, immediately before the *magnum silentium*, the great silence in which monks proceed directly from choir to dormitory. The texts are remarkably consistent across the medieval period: Psalms 4, 90, and 133, the hymn *Te lucis ante terminum* ("Before the ending of the light"), the *Nunc Dimittis* (Canticle of Simeon — *Lord, now lettest thou thy servant depart in peace*), and a Marian antiphon (typically *Salve Regina*). The musical character is similarly stable: plainchant in Mode I (Dorian) or Mode VIII (Hypomixolydian), narrow melodic ambitus often less than an octave, stepwise motion, strong gravitational pull to the final, breath-paced delivery (Hiley 1993; Boynton 2006). The agitated Phrygian materials of Mode III are explicitly avoided — the tradition recognized that certain modal characters were not appropriate at the threshold of sleep.

Compline is sung. Then the night is silent — for the monks. Then, hours later, the night office of Matins or Vigils begins, and finally Lauds at dawn — sung, melodic, brighter. *Sound at the threshold of sleep, silence (for the sleeper) through the night, sound at the threshold of waking.*

This pattern is not unique to Christianity. It recurs, with cultural and theological variation but with structurally identical shape, across many of the world's religious and contemplative traditions:

- In Jewish tradition, the *Kriat Shema al ha-Mitah* (Bedtime Shema) is recited before sleep, with the *Hashkiveinu* prayer ("Cause us, O Lord, to lie down in peace") sung in the evening *Ma'ariv* service in the *Magen Avot* mode, distinguishable from the brighter modes of morning prayer (Slobin 1989; Schleifer 2008). Morning prayer, sung again at waking, uses different *nusach*.
- In Sufi practice across the Islamic world, *wird* — personal litanies after the *isha* (night) prayer — is the documented evening practice, distinct from the ecstatic public *sama'* ceremonies of the Mevleviyya (Schimmel 1975; Ernst 1997). At dawn, *fajr* is again sung.
- In Theravada Buddhist practice, bedtime *parittas* (protective recitations including the *Karaniya Metta Sutta* and *Mangala Sutta*) are chanted across Sri Lanka, Thailand, and Burma. Morning chant follows pre-dawn.
- In Hindu daily practice, *sandhya-vandana* is performed at the three junctions (dawn, noon, dusk), with night ragas (Yaman, Bageshri, Malkauns, Darbari Kanada) elaborated in evening contexts and dawn ragas (Bhairav, Bhairavi) at waking. The night ragas use *komal* (flat) swaras at slower tempos in lower tessitura — a documented temporal-modal system codified in Bhatkhande's *Hindustani Sangeet Paddhati* (1909–1932) and earlier in the thirteenth-century *Sangita-Ratnakara* (Wade 1979; Bor 1999).

Outside formal religious structure, the same pattern repeats. The lullaby is the most universally distributed musical form on Earth — Mehr et al. (2019), working with the Natural History of Song corpus of 118 songs from 86 small-scale societies, demonstrated that lullabies are reliably distinguishable from other songs by naïve listeners across cultures, sharing a tight cluster of acoustic features (slow tempo, narrow melodic range, descending phrase contours, repetitive structure, breathy phonation, vocables over lexical content). Bainbridge et al. (2021), in *Nature Human Behaviour*, demonstrated that infants from one culture relax in response to lullabies sung in unfamiliar foreign languages — suggesting infants respond to *universal acoustic features* rather than learned cultural cues. Adults sing lullabies; lullabies end; sleep follows.

In the medieval Irish harp tradition, the three classical strains — *geantraí* (joy), *goltraí* (lament), and *suantraí* (sleep-strain) — were named genres recognized in textual sources including the *Acallam na Senórach* (Buckley 2005). *Suantraí* is one of the only documented adult-targeted sleep-music traditions in European sources; the medieval Irish recognized that adults, not just infants, benefit from music at the threshold of sleep, and they had a name for it.

What is striking about all of these is not that they are similar — they are not, in their specific musical content. Compline plainchant sounds nothing like a Yiddish *vigndlid* lullaby, which sounds nothing like a Hindustani Yaman *alap*, which sounds nothing like an Anmatyerr *wapirra*. They are similar in their *position* — at the threshold — and in *acoustic features that survive across the differences*: slow tempo, low tessitura, breath-paced phrasing, narrow range, drone or sustained pitch ground, descending contours, smooth phonation. These features are documented in the cross-cultural musical-feature literature (Mehr et al. 2018, 2019; Savage et al. 2015) and converge with the clinical sleep-music acoustic profile to a degree that is, frankly, hard to dismiss as coincidence.

2.2 The night between thresholds

If the thresholds are universally sung, what about the night between them?

The strongest evidence on this question comes not from musicology but from anthropology. David Samson and colleagues, working with the Hadza foragers of Tanzania, monitored 33 Hadza adults with actigraphy across 220 nights and reported that for only 18 minutes of the entire study period were all participants simultaneously asleep (Samson et al. 2017, *Proceedings of the Royal Society B*). The "sentinel hypothesis" follows: in groups that sleep communally, vigilance is distributed across staggered chronotypes such that the camp is never wholly unaware. The implication for soundscape is direct: ancestral human sleep happened *amid continuous low-level ambient sound* — fire crackling, low conversation, sleeping bodies stirring, distant insect chorus, the occasional infant nursing. It was not silent. It was, in the language of acoustic ecology, *high-fidelity* — sounds were spatially distinguishable, sources were identifiable, the soundscape carried information about safety and the state of the surrounding environment (Schafer 1977; Krause 2012).

Carol Worthman and Melissa Melby, in "Toward a comparative developmental ecology of human sleep" (Cambridge UP, 2002), document that traditional sleep environments worldwide are characterized by what they call "open, ambient, social, polyphasic" conditions — fire, low conversation, infant nursing, dogs, livestock, intermittent vocalization. Solitary, silent sleep is the recent WEIRD anomaly (Henrich et al. 2010), not the human norm. Roger Ekirch's *At Day's Close: Night in Times Past* (2005) documents the same pattern for pre-industrial Europe through the early modern period: bedrooms were communal, fires were tended, watchmen called the hours, livestock was nearby, church bells marked offices throughout the night. Thomas Wehr's NIMH photoperiod experiments (Wehr 1992, 1995) showed that humans deprived of artificial light spontaneously developed bimodal sleep with a quiet-wakeful interval near midnight — physiological corroboration that the segmented-sleep pattern Ekirch documented historically reflects an underlying biology.

The strongest empirical findings on what *kinds* of sounds the human nervous system reads as soothing across the night converge on a consistent profile. Gould van Praag et al. (2017), in *Scientific Reports*, demonstrated using fMRI and HRV that natural soundscapes (versus matched artificial sounds) shift autonomic balance toward parasympathetic dominance, particularly in subjects starting in a sympathetically activated state. Alvarsson, Wiens, and Nilsson (2010) showed water sounds accelerate stress recovery. Ratcliffe and colleagues (2013, 2016) documented birdsong as the most consistently associated natural sound with perceived stress recovery, with nuance — soft, melodic, non-territorial calls are soothing; harsh corvid alarms are not. Lynn (2014), in *Evolutionary Psychology*, showed campfire stimuli lowered blood pressure with effect strengthening over exposure duration, supporting a hearth-as-safety hypothesis.

The asymmetry — why these sounds and not others — maps cleanly onto signal theory. Sounds that are read as soothing are *continuous, broadband, low-modulation, non-transient*; sounds that are read as alarming are *sudden, high-amplitude, transient, or biologically threat-coded.* A predator approaching a savanna campsite produced not noise but the *cessation* of biophony — birds and insects fall silent in alarm. Continuous biophonic and geophonic sound therefore signaled, to a nervous system shaped by ancestral environments, *no immediate threat.* Stephen Porges's polyvagal theory (Porges 2011) provides a mechanism: prosodic, low-frequency, rhythmic sound engages the ventral vagal complex via the social engagement system, supporting parasympathetic dominance.

This is the soundscape humans evolved within. It is not silent. It is not melodic. It is *ambient, continuous, biophonic, low-modulation, broadband, and informative-of-safety.*

2.3 The wake threshold

The wake threshold has its own architecture, and again it is sung across cultures. Matins and Lauds in Christian monastic tradition. *Fajr* in Islam. The morning Shema in Jewish practice. Dawn ragas in Hindustani classical. The morning section of the *Sandhya.* And, biologically: the dawn chorus — the universal pre-dawn vocalization of birds across nearly every climate zone (Krause 2012), peaking just before sunrise as ambient temperature shifts and species coordinate territorial calls. Humans wake, across cultures, into a sonically richer environment than the one they slept through.

The clinical literature on waking corroborates this pattern. McFarlane and colleagues (2020), in *PLOS One*, analyzed reported waking sounds and found that *melodic* sounds were associated with reduced sleep inertia, while *neutral* sounds (rated as neither melodic nor unmelodic) were associated with *increased* sleep inertia. Hilditch and colleagues (2020) reviewed auditory countermeasures for sleep inertia and identified gradual-onset, melodic, mid-frequency-emphasis (500–1500 Hz range) audio as most effective for cognitive transition out of sleep. A 2024 NIH review reported that rising-tone alarms reduce subjective grogginess by up to 40% versus abrupt buzzers.

The cultural pattern and the clinical evidence agree: wake is supported by *gradual, melodic, mid-frequency, ascending or arch-shaped* audio. Not by silence, and not by a sudden tone.

2.4 What the architecture is

Across these traditions and across the empirical literature, then, the same three-part shape recurs:

1. **A threshold of sleep,** marked by sung/structured music with specific acoustic features (slow tempo, narrow range, descending contours, breath-paced delivery, drone or sustained pitch, modal restraint, smooth phonation).
2. **The night between thresholds,** characterized by ambient, continuous, biophonic, low-modulation sound — informative of safety, never melodic, never demanding, but never silent.
3. **A threshold of waking,** marked by gradual, melodic, mid-frequency-emphasis sound that supports cognitive transition out of sleep without alarming the nervous system.

LullaWave's nightscapes are this architecture, translated for a single listener in a modern bedroom who lacks both the village campfire and the monastic choir. The Arrival is the sung threshold of sleep — Compline, lullaby, *suantraí.* The Loop Bed is the night between thresholds — the camp at fire, the inhabited stillness. The Wake Sequence is the dawn chorus, the morning office, the sung threshold of waking.

The product premise — that audio plays continuously through the night — is not a novel claim about what humans need. It is a *restoration* of an ancestral acoustic condition that solitary, silent, technologically-mediated modern bedrooms have removed.

This framing also clarifies what we are *not* doing. We are not adding music to sleep — the night between thresholds is intentionally non-melodic in our design, consistent with the cultural and evolutionary pattern. We are not playing a continuous song. We are providing sounds at the right places with the right acoustic features, and ambient continuity in between.

---

3. The Clinical Sleep-Music Evidence

We have argued that the threshold-loop-threshold architecture is grounded in cross-cultural and evolutionary pattern. We turn now to the clinical literature, which speaks directly to the *acoustic features* of sound used in proximity to sleep — what the human nervous system, measured under controlled conditions, actually responds to.

3.1 Effect sizes

The clinical case for music as a sleep-onset intervention is reasonably well-established. Multiple meta-analyses converge on moderate-to-large effect sizes for music interventions on subjective sleep quality. Wang et al. (2014), in a meta-analysis of 10 randomized controlled trials, reported significant improvement in sleep quality (Hedges' g of approximately 0.74). Jespersen et al. (2022), updating their earlier Cochrane review with 13 RCTs and 1,007 participants, reported moderate-certainty evidence that music groups experienced better sleep quality compared to no intervention or treatment as usual. A meta-analysis of music interventions in adults with mental health problems (Gao et al. 2024) reported a moderate effect size on sleep quality (g = -0.66). Across these reviews, music interventions reduced sleep-onset latency by 13 to 27 minutes on average and improved global sleep-quality measures (typically the Pittsburgh Sleep Quality Index).

The objective polysomnographic evidence is less consistent. Loewy et al. (2013) showed effects on objective sleep parameters in NICU populations, but adult RCTs more often show subjective improvement without robust polysomnographic confirmation (Trahan et al. 2018; Cordi et al. 2019). The likely reading: music supports the *subjective experience* of sleep (perceived quality, time-to-sleep, restfulness on waking) more reliably than it shifts objective sleep architecture. For a product like LullaWave whose users are evaluating their own experience, this is a meaningful effect; for the literature, it is appropriately modest.

3.2 The acoustic profile of sleep music

What features distinguish music that supports sleep from music that does not?

Across studies and across cultures, the convergence is striking. Dickson and Schubert (2019), in "Musical features that aid sleep," analyzed acoustic features across multiple sleep-music studies and identified a consistent profile: slow tempo (typically 60–80 BPM), soft and smooth melodic contours, instrumental, simple harmonic structure, low brightness/timbral energy, and minor or modally restrained tonalities. Trahan et al. (2018), in a survey of 651 participants on what music they used for sleep and why, found enormous diversity in *specific tracks* but consistency in *acoustic features* — slowness, softness, smoothness, instrumental character, simplicity. Kirk and Timmers (2025), analyzing Spotify sleep playlists, identified loudness, energy, acousticness, and instrumentalness as the strongest acoustic discriminators between sleep-labeled and non-sleep-labeled music.

The Frontiers narrative review on elements of music that improve sleep (Wang & Liu 2025) summarized the convergent profile this way: "Music that was slow in tempo (60–80 bpm), soft and smooth in melodies, instrumental, and simple in structure, often classical or new age, was most effective." We note for transparency that one notable study (Cordi et al. 2019) found *major* mode rather than minor most effective, against the broader literature; Dickson and Schubert's review concluded that mode is contested, with individual differences likely larger than absolute mode effects.

These features are mechanistically explained. Slow tempo and breath-paced phrasing activate the parasympathetic nervous system through baroreflex resonance; Bernardi et al. (2006), in *Heart*, showed that slow-tempo music with silence between phrases (including rosary recitation and yoga mantras) produces breathing at approximately 6 breaths per minute, the resonance frequency for heart rate variability maximization. Soft attacks avoid the sudden transients that constitute the strongest correlate of auditory arousal during sleep. Lyrics activate language-processing networks that compete with the cognitive disengagement sleep requires (Scullin et al. 2021). Familiarity and distinctive melodic content increase the risk of involuntary musical imagery (earworms), which a 2021 polysomnographic study found to worsen sleep quality across the night (Scullin et al. 2021; Liikkanen & Jakubowski 2020).

LullaWave's nightscapes implement this acoustic profile in the Arrival and Wake sections (where music is present) and avoid all of its risk factors in the Loop Bed (where music is intentionally absent). The Arrival uses slow ametric phrasing, soft instrumental attacks, instrumental palette only, modally restrained voicings, and lullaby-typical descending phrase contours. The Loop Bed uses sustained drone and ambient texture without melodic content, eliminating earworm risk. The Wake Sequence resolves with a soft, brief, ascending or arch-shaped melodic phrase in the 500–1500 Hz range, consistent with the sleep-inertia evidence.

3.3 The 17.7 dB arousal threshold

A specific finding shapes our envelope decisions. Buxton et al. (2012), in a polysomnographic study of hospital noise and sleep arousal, found that the threshold for auditory arousal during sleep is not absolute volume but *change from baseline* — approximately 17.7 dB of increase. A continuous sound at moderate level is far less disruptive than a brief sound that exceeds the baseline by even a few dB more than this threshold.

This finding is what makes the Loop Bed's structure possible. A continuous drone and ambient pad — even at a level that would be conspicuous in silence — does not arouse the sleeping listener because there is no *change* large enough to cross the threshold. Conversely, any event introduced into the Loop Bed (the rare felt-piano notes, the slow brightening of an upper layer) must be calibrated so that its level rise above the immediate baseline is well under 17.7 dB. The catalog spec (Section 3.1 of the master spec) limits volume changes to 3 dB across any 30-second window in the Loop Bed — an order of magnitude below the arousal threshold.

This is also why we designed against the original spec's bell shimmer events distributed across the base layer: each bell event was below the arousal threshold individually, but the *cognitive registration* of recurring events — even quiet ones — risks pulling the listener toward conscious attention. The cross-cultural research on bells reinforces this: bells across traditions mark transitions, not continuous time. The base layer in v2 retains at most two bell-like cues, both in the first two minutes (the threshold opens), with no further events through the remaining six minutes.

3.4 The pink noise question

A particular care needs taking with broadband noise. The clinical literature on white and pink noise as continuous sleep aids is genuinely mixed. The Sleep Foundation review (Riemann et al., periodic systematic review) characterized the evidence on continuous noise machines as "inconclusive." A 2026 study at Penn Medicine reported that pink noise played at 50 dB across the night was associated with a reduction in REM sleep of approximately 19 minutes — a non-trivial effect on a stage of sleep critical to memory consolidation and emotional regulation. The earlier finding that pink noise *enhances* slow-wave sleep applied to *closed-loop auditory stimulation* — brief pink-noise pulses timed to ongoing slow oscillations via real-time EEG monitoring (Ngo et al. 2013; Papalambros et al. 2017) — not to continuous broadband playback.

The reading: pink noise as a *therapeutic intervention timed to brain state* (which requires EEG monitoring not feasible in a consumer product) appears to support deep sleep; pink noise as *continuous background* across the night has weaker evidence and at moderate level may carry some risk to REM. We therefore reduced the role of pink noise across the LullaWave catalog. The base layer's pink-noise wind bed, which was the loudest element in the original specification, is reduced by 6–8 dB and demoted to a supporting textural role rather than the lead. The pad — a slow-evolving, harmonically restrained suspended voicing — takes the lead. This shifts the catalog from being closer to a noise machine with musical seasoning to being closer to a designed musical environment with ambient seasoning, consistent with the stronger evidence base for music over noise.

3.5 Cognitive shuffling and the racing mind

A specific intervention — distinct from music — bears on the design of the Mental lane. Cognitive shuffling, also called serial diverse imagining (SDI), was developed by the cognitive scientist Luc Beaudoin in the 2010s based on what he calls the "somnolent information processing theory" (Beaudoin 2014; Beaudoin et al. 2016). The technique asks the user to visualize random, emotionally-disconnected words and images at intervals of approximately 5 to 10 seconds. The theoretical basis: rumination and worry drive cognitive arousal; the mind transitions toward sleep when its information-processing load is *non-narrative and non-emotionally-engaging*. By imagining unrelated, non-charged stimuli at slow pace, the user mimics the fragmented, nonlinear cognition characteristic of the sleep-onset hypnagogic state. Research evidence on the technique is preliminary but supportive: a Beaudoin et al. study reported sleep-onset latency improvements and reduced pre-sleep arousal in subjects using a digital implementation of SDI versus a no-treatment control (Beaudoin et al. 2024 abstract; the technique has been deployed commercially in MySleepButton and adapted in Calm's "Cognitive Shuffling" content for nearly a decade).

This finding directly shaped the Mental lane's design. The original specification used an ordered four-note figure (E-G-B-D) as a "soft focus object" for racing minds. After reviewing the cognitive shuffling literature, we recognized that an ordered figure is closer to what racing minds latch onto — patterns the mind can anticipate, melodic logic the mind can follow — than to what the cognitive shuffling research validates. We replaced the four-note figure with a *randomized note-event system*: notes drawn from the E Dorian palette, sequenced randomly, timed at random intervals between 5 and 15 seconds, placed randomly in the stereo field. The mind has nothing to track because there is nothing to track; the random stream provides the auditory equivalent of the cognitive shuffling cadence.

We acknowledge directly: this design decision has no cultural lineage. It is a modern construction informed by recent cognitive science. We name the technique, cite the research, and let the lane stand on the evidence rather than on tradition.

3.6 What we do not claim from the clinical literature

We want to be explicit about what we do not extrapolate.

The clinical sleep-music literature establishes effects on sleep onset, subjective sleep quality, and parasympathetic activation. It does *not* establish that music intervention treats clinical insomnia, anxiety disorders, depression, grief, or chronic pain — though there is associated literature suggesting modest effects on each. We do not market LullaWave as a treatment for any of these conditions, and the master specification's Part 7 makes this scope explicit. We make design decisions on the basis of the literature; we do not make therapeutic claims on the basis of the literature.

The literature on continuous all-night audio is, again, thin. We design with the architecture-level evidence (the threshold pattern is universal; the night ambient is biologically expected) and the acoustic-level evidence (specific acoustic features distinguish sleep-supporting from sleep-disrupting audio). We do not have direct evidence that LullaWave's specific catalog improves sleep across an actual night for an actual user; that would require a clinical trial we have not conducted. We design responsibly within what is known and acknowledge what isn't.

---

4. The Four Lanes: Design Rationale

The catalog is organized as four "lanes," each tuned to a specific user state at the moment of arrival. The user, in choosing a lane, is not selecting a category of music but identifying a starting condition: emotionally heavy, physically restless, exhausted, or mentally racing. Each lane begins where the user begins and provides a particular kind of support across the night.

4.1 Lane 1 — Emotional (Suantraí + Goltraí)

The Emotional lane addresses the user who arrives with overwhelm, grief, or anxiety. The intent is to meet — not lift, not console, not fix — the listener's emotional state. The acoustic vocabulary is drawn directly from the cross-cultural lullaby and lament literature.

The lane is voiced in F Aeolian, with a tonic drone at 87 Hz (F2 — chosen as the actual tonic of the mode, replacing the original specification's 174 Hz, which had been justified by the (folkloric) "solfeggio pain relief" claim). The harmonic motion is reduced to two centers (Fm and Db), each held for four to six minutes, drifting between them via shifts in drone partials rather than chord changes. The lullaby-typical descending phrase contour governs the felt-piano material in the Arrival.

A defining feature of this lane is the **breath layer**: filtered pink noise shaped to rise and fall on a six-second cycle, matching coherent breathing at approximately ten breaths per minute. This corresponds to the baroreflex resonance frequency of approximately 0.1 Hz documented by Bernardi et al. (2006) and by the broader heart-rate-variability biofeedback literature (Lehrer et al. 2014). Slow-paced breathing at this frequency produces phase coherence between cardiac oscillations and respiration, increases vagally-mediated heart rate variability, and supports parasympathetic dominance. Of all the elements in the LullaWave catalog, the breath layer is among the most directly evidence-grounded; it is one of the lane's most important, and the six-second period must not be approximated.

We name the lane *suantraí + goltraí* in recognition of the medieval Irish harp tradition's named genres for sleep-strain and lament. The Emotional lane sits in this lineage explicitly: the recognition that grief and sleep often inhabit the same emotional register, and that the right musical response to emotional heaviness at night is not consolation but companionship.

4.2 Lane 2 — Somatic (Trouble Settling, Pain)

The Somatic lane addresses the user whose body cannot settle — physical restlessness, chronic pain, or the inability to release physical tension. This lane does the most active *physiological* work of the four, and its acoustic vocabulary is correspondingly more concerned with rhythm and low-frequency content than with melody.

The lane is voiced in C Dorian, with a three-layer root drone stack at 64, 128, and 256 Hz. The stack creates a perceived "floor" that carries through the body via bone conduction (the 64 Hz layer), through standard playback systems (128 Hz), and through phone speakers in worst-case playback conditions (256 Hz). This redundancy is deliberate: the lane works regardless of playback hardware.

The lane's defining feature is its **slow sub-bass pulse** — a filtered, soft-attack pulse at 50 BPM, with the fundamental at 50 Hz heavily low-passed at 80 Hz. The pulse is calibrated to be *felt* rather than heard at typical listening levels. Its function is rhythmic entrainment: at 50 BPM, the pulse is below typical resting heart rate for most adults, providing an external rhythm that the body's autonomic regulation can entrain toward. The literature on slow rhythmic auditory entrainment (Hove et al. 2016; Hurless et al. 2013) establishes that prolonged exposure to slow stable rhythms produces measurable EEG and heart-rate responses, with the strongest effects in the theta band associated with relaxed and pre-sleep states.

A subtle structural feature of this lane is the *pulse drift*: at minute 12 of the 25-minute Loop Bed cycle, the pulse slows from 50 BPM to 48 BPM over 90 seconds, then holds. Across multiple loop cycles in a long session, the listener's body is repeatedly drawn toward 48 BPM. The cumulative effect is a small but real downward nudge to resting cardiac rhythm — consistent with the slow-tempo entrainment literature, though we have not measured this effect specifically in LullaWave users.

We are honest about a limitation of this lane: the literature on music and chronic pain shows that effects on pain-related sleep quality typically emerge over weeks of repeated use rather than within a single night (Garza-Villarreal et al. 2014; Wang et al. 2024). The Somatic lane is not designed to "fix tonight"; it is designed to support the body to settle, with cumulative effects across repeated exposure. We frame this honestly in user-facing copy.

This lane has the least direct cultural lineage of the four. While slow rhythmic auditory practices appear in some shamanic, drum-meditation, and contemplative traditions, the specific implementation here — a sub-audible pulse calibrated for entrainment under solo listening conditions — is a modern construction grounded in clinical entrainment research. We name this directly. The lane works because the entrainment evidence is solid, not because the design is ancient.

4.3 Lane 3 — Deep Rest

The Deep Rest lane addresses the user who is already tired, already somewhat settled, and seeking depth rather than emotional or physical settlement. It is the most minimal of the four lanes, and the closest to authentic contemplative drone tradition.

The lane is voiced in A Aeolian, but the lane treats A more as a *drone key* than a modal one — there is essentially no harmonic motion. The drone foundation is at 55 Hz (A1), with a 110 Hz octave layer joining at the two-minute mark of the Arrival.

A defining feature is the **Shepard tone**: a psychoacoustic illusion of perpetually descending pitch, implemented via three sine-wave layers at octave intervals with envelopes that fade in at the top of each layer's range and fade out at the bottom. The Shepard tone produces a sensation of *falling* — appropriate at the threshold of sleep onset — that has been used in psychoacoustic research and in some ambient and avant-garde compositional contexts (Shepard 1964; Risset 1986).

The original specification deployed the Shepard tone continuously through the Arrival and Loop Bed. We reduced this in the v2 specification. Continuous psychoacoustic illusions edge toward the "overly mystical" register that LullaWave's brand framework explicitly avoids, and a perpetual descending tone risks reading as a production decision rather than as an ambient phenomenon. The v2 specification has the Shepard tone present through the Arrival, tapering through the first eight minutes of the Loop Bed, and returning briefly every fifteen to twenty minutes thereafter. The tone is present approximately 30% of the time across a Loop Bed cycle. The promoted suspended pad — A1 + E2 + A2 + B3, stacked fifth and ninth, no third — carries the continuous warm element instead, addressing the original specification's tendency toward sterility.

This lane is the most vulnerable to sounding "boring" in A/B comparison. We resisted, in design, the urge to make it more interesting. Its function is to be the background against which sleep happens, not to be the sleep music itself. The cultural record across contemplative traditions converges on minimal, drone-focused composition for the most inward states (Indian classical *anahata nada* tradition, Tibetan *dbyangs* chant fundamentals, Byzantine *ison*, the Western organum tradition). Drone alone is what most contemplative traditions converge on. We follow.

4.4 Lane 4 — Mental (Racing Mind, Rumination)

The Mental lane addresses the user with a racing mind — rumination, intrusive thoughts, the inability to disengage from cognitive activity. This is the lane where we made the largest single design change between v1 and v2, on the basis of the cognitive shuffling literature reviewed above.

The lane is voiced in E Dorian, with a tonic drone at 82.5 Hz (E2). The harmonic motion across the Loop Bed drifts very slowly: E for fourteen minutes, A for four, D for four, back to E for three, implied via shifts in drone partials.

The defining feature is the **randomized note-event stream**, which replaces v1's ordered four-note focus figure. Notes are drawn from the E Dorian palette across three octaves (E3 to E6), sequenced randomly with no melodic logic between consecutive notes, timed at random intervals between 5 and 15 seconds, placed randomly in the stereo field. The stream provides the racing mind with isolated, non-narrative, non-melodically-related stimuli at the cadence of cognitive shuffling — a slow stream of "random imaginings" that the mind can lightly attend to without being able to engage with as music. Across a 25-minute Loop Bed cycle, there are approximately 165 events at a mean inter-event gap of 9 seconds.

The technical implementation deserves note. We provide three production approaches in the master specification (Section 4 of that document): a Logic Pro Scripter JavaScript that generates the stream procedurally during playback; a Step Sequencer with low per-step probability values; or manual composition using a random-number generator. The recommended approach is the Scripter script, which produces output that can be captured to MIDI and edited for any unintended clusters.

The Wake Sequence in the Mental lane is the catalog's most "musical" wake — a clear ascending phrase appears, in deliberate contrast to the Loop Bed's randomness, signaling that the mind can now engage with ordered material as a positive cue. The contrast carries the wake.

We name directly that this lane has no cultural lineage. The cognitive shuffling technique was developed in the 2010s and has been deployed commercially for nearly a decade in MySleepButton and adapted by Calm. Beaudoin's *somnolent information processing theory* is a cognitive-science framework, not an ancient tradition. The Mental lane works because the research evidence supports it, and we place no false historical mantle on the design.

4.5 Variants and Repertoire

Each of the four lanes has three variants (A, B, C) — three musical takes on the same lane identity. A listener encounters one variant per session, randomly or rotationally selected by the app. The variant system is a structural feature of the catalog, not an aesthetic embellishment, and its rationale is worth naming explicitly.

The clinical literature on involuntary musical imagery (earworms) demonstrates that repeated exposure to the same audio strengthens involuntary recall rather than building habituation tolerance (Liikkanen & Jakubowski 2020; Scullin et al. 2021). A user who listens nightly to a single fixed track in a given lane will, over weeks and months, accumulate hundreds of identical exposures — the precise condition under which earworms most reliably form. With three variants per lane, the same user encounters each variant approximately every seven to ten nights, a frequency below the threshold where earworm-style memory consolidation strongly takes hold.

The variant system also addresses a compositional problem. Single-variant design forces tight melodic constraint within the Arrival, because any phrase will be heard nightly. Multi-variant design distributes that constraint across three musical takes, allowing each variant to be more developed musically — more present felt-piano writing, more developed wake melodies, more articulated harmonic moments — without any single variant carrying every-night repetition risk. The Arrival in particular benefits: the brand promise of "musician-built soundscapes" is most expressed in this section, and three variants give the section room to actually sound like music rather than a restrained gesture toward it.

The cross-cultural pattern argues for the variant system as well. Lullabies historically were not single fixed songs but repertoires; a parent sang different lullabies on different nights. Compline rotates through Marian antiphons by liturgical season — *Salve Regina* in ordinary time, *Alma Redemptoris Mater* in Advent, *Regina Caeli* in Easter — so monks did not sing the same closing antiphon every night across the year. Indian raga-time theory has multiple "night ragas" (Yaman, Bageshri, Malkauns, Darbari Kanada), and a listener over a season will hear several. The single-track-per-lane design is the anomaly against the cultural pattern; the three-variant design moves us toward how the architecture has historically worked.

The catalog therefore comprises 48 nightscape masters (4 lanes × 4 durations × 3 variants) plus the ambient base layer, totaling 49 files. Within each lane, all three variants share the lane's modal palette, drone foundation, instrumentation, reverb space, and architectural shape; they differ in musical content — what the felt piano does, where the harmonic weight settles, how the wake melody resolves. A listener encountering different variants of the same lane on different nights should feel they're in *the same lane* — same emotional address, same physiological work — but with a different musical face on it.

4.6 The 30-Minute Arrival

The Arrival in the v3 specification is 30 minutes in length, an extension from the 10-minute Arrival of the original spec. The change was prompted by user feedback during early production: users wanted "more time with a piece that actually sounds like music" before the audio transitioned into the Loop Bed's drone-rich ambience. The clinical literature supports the extension: sedative music interventions for sleep onset typically use 30 to 45 minutes of content (Wang & Liu 2025; Jespersen et al. 2022), with some interventions running up to 60 minutes for users with mental health concerns. The original 10-minute Arrival fell well below this range — an under-implementation that predated our review of the bedtime-music duration evidence.

The 30-minute Arrival is structured in four sections within each lane: an Entry (minutes 0–4) that introduces the lane's drone and harmonic foundation, a Music-Rich section (minutes 4–14) that contains the Arrival's most developed musical material, a Settling section (minutes 14–22) that reduces note density and melodic motion, and a Release/Handover section (minutes 22–30) that transitions to the Loop Bed's ambient state. A user falling asleep within the typical 5–20 minute sleep onset window encounters music for most or all of their pre-sleep period, consistent with the clinical-research design of bedtime-music interventions.

A consequence of the 30-minute Arrival is that the 30-minute duration master in the catalog now functions as a standalone bedtime music piece: the Arrival in its entirety, with a graceful fade-out across the final minute, no Loop Bed or Wake Sequence. This makes the 30-minute duration a distinct use case from the 6/8/10-hour durations and gives the catalog a pedagogical entry point — a user not yet ready for the all-night experience can listen to a complete 30-minute piece in their chosen lane.

4.7 The Mental Lane's Three-Phase Arrival

The Mental lane required a structural variation. The Mental lane's defining feature is the cognitive-shuffling random-event system (Section 3.5), which is by design non-musical. A 30-minute Arrival composed entirely of random non-melodic events would not meet the user-feedback request for "more time with music"; conversely, a fully melodic Mental Arrival would lose the cognitive-shuffling identity that makes the lane work for ruminating users. The v3 specification implements a three-phase compromise: minutes 0–12 are melodic and developed, minutes 12–22 transition gradually from melodic phrases through scattered events, and minutes 22–30 are fully randomized in the cognitive-shuffling style.

This structure preserves both functions. In the music-rich first 12 minutes, the racing mind is given something to track — descending modal phrases on felt piano, sustained viola counterpoint, the slow opening of a Pharlight pad voiced in E Dorian. The cognitive shuffling work begins later, when the user has had some preliminary calming and is closer to sleep onset. The clinical research on cognitive shuffling (Beaudoin et al. 2016; Beaudoin 2014) supports the technique most strongly for the *transition into sleep* itself rather than for early pre-sleep arousal reduction; the Mental Arrival's sequencing aligns with this framing.

4.8 The Ambient Base Layer

The single non-variant deliverable in the catalog — the eight-minute ambient base layer — plays whenever the LullaWave app is open and a nightscape has not yet been selected, and serves as the underscore for the voice ritual that precedes lane selection. Its function is liminal: it makes the app feel inhabited while a user browses, signals to the nervous system that the wind-down has begun, supports the voice ritual without competing with it, and provides harmonic neutrality that crossfades into all four lanes without clash.

The base layer's harmonic foundation is built on A — the one note common to all four lane scales (F Aeolian, C Dorian, A Aeolian, E Dorian). The held chord voicing is stacked fifth-and-ninth on A (A2 + E3 + B3), suspended, harmonically ambiguous, neither major nor minor. The drone is at 55 Hz (A1).

The original specification used a low-passed pink noise wind bed as the loudest element of this layer. The v2 specification replaced that with a vocal-granular pad (Pharlight) that promoted the harmonic content over noise. The current specification (v3) revised this further: the lead element is now Noire's felt piano with the Particle engine engaged, holding the A2 + E3 + B3 chord in sustain throughout the eight-minute duration. The Particle engine generates the evolving atmospheric overtones that earlier versions sourced from synthetic pads or vocal granular instruments.

This evolution reflects a learning specific to the base layer's role under the voice ritual. Vocal-granular content underneath actual human voice produces what we came to call a "voice under voice" doubling — two vocal-source signals that the auditory system parses as competing rather than complementary. Piano-DNA atmospheric content does not produce this conflict; it sits underneath voice cleanly. The vocal-warmth element that earlier versions sourced from Pharlight in the base layer is now present in the catalog only in the four lane tracks, where the voice ritual has already concluded and where the cross-cultural vocal warmth signal can do its work without competing with actual speech. The base layer's musician-built character is now unambiguously delivered by piano — the literal instrument the listener might most readily associate with bedtime music — while still satisfying all of the layer's functional constraints (voice support, harmonic neutrality across lanes, low information density, parasympathetic priming).

The base layer also implements one of the cross-cultural findings we noted above: bells mark transitions, not time. At most two felt-piano notes appear in the base layer, both in the first two minutes — *the threshold opens* — and the remaining six minutes are pad-and-drone only. The original specification distributed bell events across all eight minutes; we corrected this in v2.

---

5. Production Constraints

A design grounded in evidence imposes specific constraints on production. We name several here that may be of interest to the careful reader, particularly if any portion of the catalog is reproduced or adapted by other producers.

5.1 The earworm-prevention constraint

Repeated nightly listening to the same audio strengthens earworm risk rather than habituating away from it (Liikkanen & Jakubowski 2020; Scullin et al. 2021). For a product whose users will play the same track over many nights, this is a sleep-safety concern that must shape composition. We applied four constraints across the catalog:

- **Embedded melodic content, never foregrounded.** Any melodic material sits inside the ambient bed at lower volume than the bed itself.

- **No phrase repeated more than twice within a 25-minute Loop Bed cycle.** This includes pitch-identical repetitions of single notes.
- **No melodic leaps wider than a fifth.** Stepwise motion is the strong default. Fourths and fifths are exception, not norm.
- **Modal rather than melodic distinctiveness.** The lane should be recognizable by its *mode* (a listener should feel "this is the Dorian lane") but not by its *tune* (no listener should be able to hum a fragment back).

5.2 Volume calibration

The catalog targets -20 LUFS integrated for nightscape masters and -28 LUFS integrated for the ambient base layer, with true peak ceilings of -1 dB and -3 dB respectively. These targets are approximately 10 LUFS quieter than commercial popular-music masters. The reasoning is twofold: (a) sleep audio plays at user-set low volume, so headroom matters, and (b) lower integrated loudness allows greater dynamic range without crossing the 17.7 dB arousal threshold within the audio itself.

We also designed for the *playback condition*, which for most users is a phone speaker at low volume rather than studio monitors. The catalog is tested through phone speakers at simulated 25–30 dB SPL at the ear, through laptop speakers, through over-ear headphones, and through a single speaker with a pillow over it (simulating side-sleeping). Mono compatibility is verified for every master.

5.3 The shared reverb space

A single reverb instance — ValhallaRoom (preferred) or Logic's Space Designer (fallback), large hall, ~5 second decay, aggressive high-frequency damping above 4 kHz, low cut at 200 Hz on the return — is deployed across all 17 catalog files. Every pitched element across all lanes goes to this reverb at varying send levels. The space is constant; the wetness of any individual element varies. This constraint produces the catalog's recognizable acoustic identity — the sense that all 17 nightscapes inhabit the same room — without requiring melodic or thematic similarity.

5.4 The non-third constraint

Across the catalog, no pad uses a chord-voiced major or minor third. Pads are voiced as stacked fifths, ninths, sus2, or sus4. This is the unifying compositional grammar. It serves three simultaneous purposes: (a) it gives the catalog a unified compositional fingerprint, (b) it preserves harmonic ambiguity (a brand and emotional choice — *warm but not performative*, *held without explanation*), and (c) it makes crossfades between the base layer and any lane work without retuning.

5.5 Tuning

The catalog tunes to A = 444 Hz. The original specification cited a (folkloric) solfeggio justification that we have removed; the practical justification for 444 Hz is registral warmth and catalog consistency. The 4 Hz difference from concert standard 440 Hz is small and audible only on direct A/B comparison; we maintain it because all 17 files in the catalog must be in tuning agreement, and 444 Hz was the original choice. Producers reproducing or adapting the catalog should match.

6. What We Do Not Know

We have argued, in the foregoing sections, that LullaWave's nightscape architecture is grounded in convergent evidence from three sources: cross-cultural ethnomusicology, evolutionary and historical anthropology, and clinical sleep-music research. We have explained the design rationale for each lane. It would be a poorer paper if we did not also name what remains uncertain.

The all-night question is genuinely open. The clinical literature on continuous audio across the full sleep period — as opposed to bedtime music for sleep onset — is thin. We have made the case (Section 2) that ancestral human sleep was always *ambient*, never silent, and that LullaWave's continuous Loop Bed is a translation of that ancestral acoustic condition. The cross-cultural and evolutionary case is, we believe, strong. The clinical case is weaker, because the relevant studies have not been done. We have designed the Loop Bed to minimize known risk factors (no melodic content, no familiar songs, no lyrics, no events that cross the arousal threshold) and to align with the cultural-evolutionary evidence, but we cannot point to a randomized controlled trial demonstrating that LullaWave's specific Loop Bed improves sleep maintenance versus silence or versus an alternative audio condition. This is the most significant uncertainty in our design.

Effects on objective sleep architecture (versus subjective sleep quality) remain unclear. The clinical sleep-music literature shows reasonably robust effects on subjective measures (PSQI, ISI, sleep diary reports of quality and time-to-sleep) but more variable effects on polysomnographic measures of sleep stages, arousal frequency, sleep efficiency, and REM duration. We do not claim to alter sleep architecture in any specific way. We claim to design audio that the literature suggests will be experienced as supportive at the threshold of sleep, and that follows the cross-cultural pattern for ambient audio across the night.

Individual variation is substantial. Trahan et al. (2018) documented enormous individual diversity in what music users select for sleep. Cordi et al. (2019) found major-mode music more effective in their sample, against the broader literature's slight preference for minor. Personal history, taste, learned associations, and individual nervous-system reactivity all shape what audio supports a given listener's sleep. LullaWave's four-lane design acknowledges this — each lane addresses a different starting state — but it does not eliminate the underlying individual variation. Some users will find lanes that fit them well. Some will not. We do not yet know what proportion will fall in each category.

The cumulative effects question is also open. Some intervention literatures (chronic pain, fibromyalgia, anxiety-driven insomnia) suggest that music-based interventions show effects that emerge over weeks or months of repeated use, not within a single session (Garza-Villarreal et al. 2014; Wang et al. 2024). LullaWave is designed for nightly use; cumulative effects may matter more than acute ones. We have no longitudinal data on LullaWave specifically.

The Mental lane's randomized note system has not been independently validated as an audio intervention. The cognitive shuffling literature validates the *cognitive* technique — the user mentally generating random images. The Mental lane translates the cadence and randomness into auditory form, but no published study has tested whether listening to randomized note streams produces the same sleep-onset benefits as the active cognitive technique. We hypothesize, on the basis of the cognitive-shuffling theoretical framework and the broader auditory-attention literature, that exposure to non-narrative non-melodic random stimuli at the cognitive-shuffling cadence will support similar sleep-onset effects. We name this as a hypothesis, not a finding.

The Shepard tone's use as a sleep aid is undocumented in clinical literature. Shepard tones are well-established psychoacoustically (Shepard 1964; Risset 1986) and have been used in ambient and avant-garde compositional contexts, but no controlled study has tested whether the Shepard illusion supports sleep onset. We use it in the Deep Rest lane on the basis of psychoacoustic plausibility and on the absence of any countervailing evidence; if the technique is not effective for sleep specifically, the Deep Rest lane's drone foundation and promoted suspended pad will carry the lane regardless.

Crossfades, loop seams, and long-session continuity are tested via composer ear and beta listener feedback, not by polysomnography. We have designed and tested for inaudibility of seams in conscious listening, but we cannot confirm that loop transitions across a 6-, 8-, or 10-hour session do not produce micro-arousals during sleep stages. The 60-second crossfades and the harmonic-neutrality choice for loop points are meant to address this, but the empirical confirmation would require a polysomnographic study we have not conducted.

We are honest about all of the above. The paper cannot end with claims it cannot defend. What we can defend: that the architecture is grounded in convergent evidence, that the acoustic features of each section align with the relevant literature, that the production-level design follows the constraints the evidence and the brand framework imply, and that we have built the catalog to support the listener as well as the available evidence allows.

---

7. Closing Reflection

A return to the founder's voice.

I started this work because the audio I was making for sleep wasn't doing what I wanted it to do — and I couldn't, at first, articulate why. The technical brief was correct; the production was clean; the mixes met the targets. None of it felt like anything.

What I learned, going through the literature and the cultural record, was that the missing thing wasn't a *production* missing thing. It was an *architectural* missing thing. We were treating sleep audio like a single artifact — eight minutes, eight hours — that should be uniform across its duration. Of course it wasn't doing the work; that's not how humans use sound around sleep, and never has been. We were making continuous *something* and calling it a sleep aid, when what humans across thousands of years and dozens of separate civilizations have done is something quite specific: *sound at the entry, ambient through the middle, sound at the exit.* We had been building one thing where we needed three.

When I see the threshold-loop-threshold structure now, I see it everywhere. The mother who hums while she puts the child down, then quiets while the child sleeps, then says good morning. The Compline office and the great silence and Lauds. The lullaby and the hearth and the dawn. None of these are inventions of a sleep app. They're inventions of human beings who needed to sleep and who, working without empirical research and without product designers, found a shape that worked.

What I take from that is humility about the work. We didn't design the architecture; we recognized it. What we *did* design — the specific musical content, the harmonic grammar, the cognitive-shuffling Mental lane, the no-third pad voicings, the breath layer at the resonance frequency, the felt piano as the human signature — those choices are ours, and they live or die on whether they serve the architecture well. The architecture itself is older than any of us.

There are things in the work I still don't know. Whether the all-night Loop Bed will hold up to polysomnographic study. Whether the random Mental lane will help ruminators at the rates the cognitive-shuffling research suggests. Whether the choice to demote pink noise was correct in every case or only in most cases. I am honest with myself that the work is provisional and that listeners will tell us, over months and years, what works and what doesn't.

What I am confident about: I would rather build something grounded in evidence and tradition that I can defend than something marketable that I cannot. And I would rather make a sleep tool that *holds* than one that performs. The brand brief I wrote at the beginning of this project asked LullaWave to feel "like something already there before the user arrived — not something that was built for them, but something they have found." If we have done this well, that is what these nightscapes will feel like. Not a clever app. A shape humans have always known.

— *T*

Bibliography

Clinical sleep-music research and meta-analyses

Bainbridge, C. M., Bertolo, M., Youngers, J., Atwood, S., Yurdum, L., Simson, J., Lopez, K., Xing, F., Martin, A., & Mehr, S. A. (2021). Infants relax in response to unfamiliar foreign lullabies. *Nature Human Behaviour*, 5(2), 256–264.

Beaudoin, L. P. (2014). *A design-based approach to sleep-onset and insomnia: super-somnolent mentation, the cognitive shuffle and serial diverse imagining*. Conference paper, ACS 2014.

Beaudoin, L. P., Digdon, N., O'Neill, K., & Racour, G. (2016). Serial Diverse Imagining Task: A new remedy for bedtime complaints of worrying and other sleep-disruptive mental activity. *Sleep*, 39(Suppl A410).

Bernardi, L., Porta, C., & Sleight, P. (2006). Cardiovascular, cerebrovascular, and respiratory changes induced by different types of music in musicians and non-musicians: the importance of silence. *Heart*, 92(4), 445–452.

Buxton, O. M., Ellenbogen, J. M., Wang, W., Carballeira, A., O'Connor, S., Cooper, D., Gordhandas, A. J., McKinney, S. M., & Solet, J. M. (2012). Sleep disruption due to hospital noises: a prospective evaluation. *Annals of Internal Medicine*, 157(3), 170–179.

Carskadon, M. A., & Dement, W. C. (2017). Normal Human Sleep: An Overview. In *Principles and Practice of Sleep Medicine* (6th ed.). Elsevier.

Cordi, M. J., Ackermann, S., & Rasch, B. (2019). Effects of Relaxing Music on Healthy Sleep. *Scientific Reports*, 9, 9079.

Dickson, G. T., & Schubert, E. (2019). Musical features that aid sleep. *Musicae Scientiae*, 1029864920972161.

Gao, Y., et al. (2024). A systematic review and meta-analysis of music interventions to improve sleep in adults with mental health problems. *European Psychiatry*.

Garza-Villarreal, E. A., Wilson, A. D., Vase, L., Brattico, E., Barrios, F. A., Jensen, T. S., Romero-Romo, J. I., & Vuust, P. (2014). Music reduces pain and increases functional mobility in fibromyalgia. *Frontiers in Psychology*, 5, 90.

Hilditch, C. J., et al. (2020). Auditory Countermeasures for Sleep Inertia: Exploring the Effect of Melody and Rhythm in an Ecological Context. *Clocks & Sleep*, 2(4).

Hove, M. J., Stelzer, J., Nierhaus, T., Thiel, S. D., Gundlach, C., Margulies, D. S., Van Dijk, K. R., Turner, R., Keller, P. E., & Merker, B. (2016). Brain Network Reconfiguration and Perceptual Decoupling During an Absorptive State of Consciousness. *Cerebral Cortex*, 26(7), 3116–3124.

Jespersen, K. V., et al. (2022). Listening to music for insomnia in adults. *Cochrane Database of Systematic Reviews*.

Kirk, R., & Timmers, R. (2025). Characterizing music for sleep: A comparison of Spotify playlists. *Musicae Scientiae*.

Lehrer, P. M., & Gevirtz, R. (2014). Heart rate variability biofeedback: how and why does it work? *Frontiers in Psychology*, 5, 756.

Liikkanen, L. A., & Jakubowski, K. (2020). Involuntary musical imagery as a component of ordinary music cognition: A review of empirical evidence. *Psychonomic Bulletin & Review*, 27(6), 1195–1217.

Loewy, J., Stewart, K., Dassler, A. M., Telsey, A., & Homel, P. (2013). The Effects of Music Therapy on Vital Signs, Feeding, and Sleep in Premature Infants. *Pediatrics*, 131(5), 902–918.

McFarlane, S. J., Garcia, J. E., Verhagen, D. S., & Dyer, A. G. (2020). Alarm tones, music and their elements: Analysis of reported waking sounds to counteract sleep inertia. *PLOS ONE*, 15(1), e0215788.

Mehr, S. A., Singh, M., Spelke, E. S., Glowacki, L., & Krasnow, M. M. (2018). Form and function in human song. *Current Biology*, 28(3), 356–368.

Mehr, S. A., Singh, M., Knox, D., Ketter, D. M., Pickens-Jones, D., Atwood, S., et al. (2019). Universality and diversity in human song. *Science*, 366(6468), eaax0868.

Ngo, H. V., Martinetz, T., Born, J., & Mölle, M. (2013). Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. *Neuron*, 78(3), 545–553.

Papalambros, N. A., Santostasi, G., Malkani, R. G., Braun, R., Weintraub, S., Paller, K. A., & Zee, P. C. (2017). Acoustic Enhancement of Sleep Slow Oscillations and Concomitant Memory Improvement in Older Adults. *Frontiers in Human Neuroscience*, 11, 109.

Roehrs, T., Carskadon, M., Dement, W. C., & Roth, T. (2017). Daytime Sleepiness and Alertness. In *Principles and Practice of Sleep Medicine* (6th ed.). Elsevier.

Savage, P. E., Brown, S., Sakai, E., & Currie, T. E. (2015). Statistical universals reveal the structures and functions of human music. *PNAS*, 112(29), 8987–8992.

Scullin, M. K., Gao, C., & Fillmore, P. (2021). Bedtime Music, Involuntary Musical Imagery, and Sleep. *Psychological Science*, 32(7), 985–997.

Shepard, R. N. (1964). Circularity in Judgments of Relative Pitch. *Journal of the Acoustical Society of America*, 36(12), 2346–2353.

Trahan, T., Durrant, S. J., Müllensiefen, D., & Williamson, V. J. (2018). The music that helps people sleep and the reasons they believe it works: A mixed methods analysis of online survey reports. *PLOS ONE*, 13(11), e0206531.

Wang, C. F., Sun, Y. L., & Zang, H. X. (2014). Music therapy improves sleep quality in acute and chronic sleep disorders: a meta-analysis. *International Journal of Nursing Studies*, 51(1), 51–62.

Wang, et al. (2024). Effects of Music Intervention on Pain, Mood, Sleep, and Heart Rate Variability in Patients with Chronic Pain. *Journal of Pain Research*.

Cross-cultural ethnomusicology and historical sources

Bhatkhande, V. N. (1909–1932). *Hindustani Sangeet Paddhati*. Sangeet Karyalaya.

Bor, J. (1999). *The Raga Guide: A Survey of 74 Hindustani Ragas*. Nimbus Records.

Boynton, S. (2006). *Shaping a Monastic Identity: Liturgy and History at the Imperial Abbey of Farfa*, 1000–1125. Cornell University Press.

Buckley, A. (2005). Music in Ireland to c. 1500. In *A New History of Ireland*, Vol. 1. Oxford University Press.

Crocker, R. L. (2000). *An Introduction to Gregorian Chant*. Yale University Press.

Ernst, C. W. (1997). *The Shambhala Guide to Sufism*. Shambhala.

Helffer, M. (1994). *Mchod-rol: Les instruments de la musique tibétaine*. CNRS Éditions.

Hiley, D. (1993). *Western Plainchant: A Handbook*. Oxford University Press.

Levin, T. C., & Süzükei, V. (2006). *Where Rivers and Mountains Sing: Sound, Music, and Nomadism in Tuva and Beyond*. Indiana University Press.

Lortat-Jacob, B. (1995). *Sardinian Chronicles*. University of Chicago Press.

Schimmel, A. (1975). *Mystical Dimensions of Islam*. University of North Carolina Press.

Schleifer, E. (2008). *Studies in Jewish Music* (collected essays).

Slobin, M. (1989). *Chosen Voices: The Story of the American Cantorate*. University of Illinois Press.

Wade, B. C. (1979). *Music in India: The Classical Traditions*. Prentice-Hall.

Anthropology, evolutionary biology, soundscape ecology

Ekirch, A. R. (2001). Sleep We Have Lost: Pre-industrial Slumber in the British Isles. *American Historical Review*, 106(2), 343–386.

Ekirch, A. R. (2005). *At Day's Close: Night in Times Past*. W. W. Norton.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2-3), 61–83.

Krause, B. (2012). *The Great Animal Orchestra: Finding the Origins of Music in the World's Wild Places*. Little, Brown.

Porges, S. W. (2011). *The Polyvagal Theory: Neurophysiological Foundations of Emotions, Attachment, Communication, and Self-Regulation*. W. W. Norton.

Samson, D. R., Crittenden, A. N., Mabulla, I. A., Mabulla, A. Z. P., & Nunn, C. L. (2017). Chronotype variation drives night-time sentinel-like behaviour in hunter-gatherers. *Proceedings of the Royal Society B*, 284(1858), 20170967.

Schafer, R. M. (1977/1994). *The Soundscape: Our Sonic Environment and the Tuning of the World*. Destiny Books.

Wehr, T. A. (1992). In short photoperiods, human sleep is biphasic. *Journal of Sleep Research*, 1(2), 103–107.

Wehr, T. A. (1995). The durations of human melatonin secretion and sleep respond to changes in daylength (photoperiod). *Journal of Clinical Endocrinology and Metabolism*, 73(6), 1276–1280.

Wilson, E. O. (1984). *Biophilia*. Harvard University Press.

Worthman, C. M., & Melby, M. K. (2002). Toward a comparative developmental ecology of human sleep. In M. A. Carskadon (Ed.), *Adolescent Sleep Patterns: Biological, Social, and Psychological Influences*. Cambridge University Press.

Specific empirical studies on natural sound

Alvarsson, J. J., Wiens, S., & Nilsson, M. E. (2010). Stress recovery during exposure to nature sound and environmental noise. *International Journal of Environmental Research and Public Health*, 7(3), 1036–1046.

Gould van Praag, C. D., Garfinkel, S. N., Sparasci, O., Mees, A., Philippides, A. O., Ware, M., Ottaviani, C., & Critchley, H. D. (2017). Mind-wandering and alterations to default mode network connectivity when listening to naturalistic versus artificial sounds. *Scientific Reports*, 7, 45273.

Lynn, C. D. (2014). Hearth and campfire influences on arterial blood pressure: defraying the costs of the social brain through fireside relaxation. *Evolutionary Psychology*, 12(5), 983–1003.

Ratcliffe, E., Gatersleben, B., & Sowden, P. T. (2013). Bird sounds and their contributions to perceived attention restoration and stress recovery. *Journal of Environmental Psychology*, 36, 221–228.

Psychoacoustic foundations

Risset, J.-C. (1986). Pitch and Rhythm Paradoxes: Comments on "Auditory Paradox Based on Fractal Waveform". *Journal of the Acoustical Society of America*, 80(3), 961–962.

Vaschillo, E. G., & Vaschillo, B. (2017). Heart rate variability and the resonance principle. In *Self-regulation and resonance breathing*. Springer.

Erratum and framework refinement note (May 8, 2026)

Between this paper's writing and its first wide reading, the master specification it described moved from version 3 to version 4. The framework's architecture — threshold-loop-threshold; the four lanes; the cross-cultural threshold-tradition convergence; the 30-minute Arrival; the cognitive-shuffling principle in the Mental lane; the honest-scope claims about what this specification does and does not promise — remains current. The specific compositional and production prescriptions instantiating that framework changed materially.

Three refinements warrant noting:

Modal palette. v3 specified Aeolian and Dorian only across the catalog, with major modes restricted to wake-sequence resolution moments. The Phase 1 research deep-dive that informed v4 found this premise not defensible at the strength v3 implied: clinical sleep-music research (Dickson & Schubert 2022 in particular) finds successful sleep-aiding music more often in major mode than minor; the cross-cultural threshold-and-night-music tradition is mode-mixed when actually examined (Compline's Salve Regina exists in both Mode I Dorian and Simple Tone V Lydian; Indian raga evening practice prominently uses Yaman, a Lydian-equivalent); the universal lullaby record's universals are timbral and dynamic, not modal; and Western tonal-music affective-character research does not support "minor = calming" for Western listeners. v4 commits all four lanes to major modes (F major, C major, A major, E major), with the base layer's 55 Hz A drone serving as the catalog-wide pivot pitch — A is in F major, C major, A major, E major at the same scale degrees as in v3's parallel minor/Dorian assignments. Catalog-consistency engineering preserved across the mode shift.

The breath-paced amplitude-modulated layer in Lane 1. v3 specified the LFO at a 6-second period, citing Bernardi 2006's baroreflex resonance frequency. The 6-second period actually corresponds to 10 breaths per minute, not 6 — the Bernardi/Vaschillo/Lehrer literature establishes the resonance frequency at 0.1 Hz, period 10 seconds, six breaths per minute. v4 corrects the LFO period to 10.0 seconds and deploys the layer across all four lanes at sub-perceptual levels as a catalog-wide unifying signature element grounded in autonomic physiology.

The R&D-vs-composer line. v3 prescribed specific notes, voicings, melodic phrases, and dB levels at strengths the research could not actually support. v4 strips those prescriptions and replaces them with principle-level constraints plus per-lane affective briefs. The composer holds compositional and ear-level production authority within the framework's evidence-bound constraints. The "musician-built" brand promise is strengthened by v4 — the composer is composing, not following AI's prescriptions.

The v4 master specification lives at `Soundscape Dev/nightscape_master_spec_v4.md`. v3 is archived at `Info for old builds/nightscape_master_spec_v3.md`. The full audit and research deep-dive that drove the v3 → v4 refinement live in `Projects/07_Soundscape_Spec_v4_Rework/`.

End of paper.

LullaWave Nightscape Paper, version 1.0 (with May 8, 2026 erratum).