

The Auditory Pathway as a General-Purpose Data Channel

Capability scope, physical limits, and a dual-use analysis published alongside

Matthew "Bear" Schulz

Independent researcher

Belmont, California, United States

ORCID: [0009-0000-0956-5743](https://orcid.org/0009-0000-0956-5743)

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0)

Preprint · 27 August 2026

Keywords: sensory substitution · sensory expansion · bone conduction · implanted transceiver · dual-use · device governance · single-sided deafness

Companion papers: *Closing the Loop Behind the Cochlea* and *Ultrasonic Cochlear Excitation — Device Scoping*

0. What this document argues

Everything built for hearing assistance treats the auditory pathway as a *repair* target. But the pathway is a high-bandwidth sensory input that the brain is exceptionally good at parsing, it is already reached by implanted transducers in routine clinical use, and those transducers already accept arbitrary digital audio.

The version at issue here is an implanted transceiver in the mastoid carrying not only telephony but any stream that can be rendered to sound: environmental sensing, spatial navigation, proximity to fields a person cannot otherwise detect, or the output of a wearable instrument.

The step less often taken is to treat it as a general-purpose input channel for people whose hearing is fine. Note what this does *not* require: the implanted mastoid transceiver already exists as a regulated commercial device (\$1), so the question is not whether such hardware can be built but what it is permitted to carry and for whom. That framing is **not unoccupied**, and it would be wrong to claim otherwise: sensory *expansion* for normally-sensed people has been demonstrated, including a haptic wristband driven by infrared and ultraviolet sensors that let a wearer detect infrared cameras in darkness, and spinal cord stimulation has been trialled as a route for sound perception. The auditory, implanted version of that idea is where the literature is thinner — and thinner is a claim this document can support, where *unoccupied* is not.

This document argues that the capability is closer than it looks, that its binding constraint is not technical, and that its misuse profile is serious enough that it should be published with the capability rather than discovered afterward.

1. The components already exist

Nothing in the capability requires an invention. The parts are listed here so that no novelty is claimed for any of them.

- **An implanted mastoid transceiver is not a proposal — it is a shipping product.** The Osia active transcutaneous system places a piezoelectric actuator subcutaneously behind the ear, coupled to an osseointegrated titanium screw in the mastoid, with the external processor held magnetically over an implanted receiver coil across intact skin. Multicentre investigation, systematic review and meta-analysis, and paediatric and skull-base series are all published. That is precisely the architecture this document examines, already regulated, implanted and reimbursed — which removes any question of whether the hardware path exists and moves the entire discussion to what is sent through it.
- Percutaneous and transcutaneous bone conduction implants have three decades of use, with FDA-approved passive transcutaneous systems since 2013.
- **Direct digital audio streaming into implants** is commercial. Manufacturers ship streaming accessories and bimodal streaming partnerships that put arbitrary audio into an implanted device from a phone.
- **Sensory substitution is an established field.** Visual-to-auditory navigation aids for blind users have been built and trialled, and one such study explicitly recommends bone conduction *precisely because it leaves the*

ear canals unblocked so real-world sound is still heard — which is the same argument that makes an implanted channel attractive.

So the capability is an integration, not a discovery. That is the honest claim and it is a weaker one than it first appears.

2. What the auditory channel is actually good for

The auditory system is a strong choice for a synthetic data channel and the reasons are specific rather than rhetorical. It has high temporal resolution, so it distinguishes events separated by microseconds where vision cannot. It is omnidirectional and does not require the head to be pointed at anything. It operates in darkness, in smoke, and behind the wearer. And it is already the sensory system humans use for continuous background monitoring while attention is elsewhere.

Those properties suit exactly what Section 5 proposes to put through it.

3. The binding constraint is learning, not bandwidth

The channel is wide. The bottleneck is elsewhere, and decades of sensory substitution work have located it precisely.

Devices that translate one sense into another *work* — participants navigate obstacle courses and identify objects. The reported limitations are perceptual rather than technical: encodings that are not intuitive require extended training, and magnitudes mapped onto loudness are hard to interpret because loudness carries no intrinsic scale.

A claim about abandonment needs correcting rather than repeating. Non-use of implanted auditory devices is not uniformly high, and the pattern is informative. A thirty-year single-institution series of 356 recipients found cochlear implant non-use at **2.8%**, while middle-ear implant non-use ran at **27.3%** — with device failure and surgical complication cited for the latter group. Separately, up to **16%** of adult cochlear implant recipients are reported to have poor long-term outcomes, defined as under 30% word or sentence recognition at twelve months.

Read together: people keep a device that clearly works and abandon one that fails or disappoints, which is unsurprising but sets a bar. A channel whose benefit depends on the user learning a new encoding sits closer to the second category than the first until the encoding is demonstrated.

The design consequence is that encoding choice dominates hardware. A channel that delivers ten times the data through an encoding the user cannot internalise delivers less than a narrow channel through one they can. Any programme here should spend its first effort on encoding and training studies with non-implanted bone conduction hardware, which is cheap, reversible and available today. **If an encoding does not work through a headset, implanting it will not rescue it.**

4. The constraint that actually governs: surgical risk in healthy people

This is the argument that decides whether the vision is responsible, and it is not a technical one.

Implanting a device in a person with normal hearing means accepting the risks of temporal bone surgery — facial nerve injury, taste disturbance, infection, cerebrospinal fluid leak, device failure and revision surgery, and the permanent presence of hardware — in exchange for a benefit that is **convenience rather than restoration**.

For a deaf person, that trade is straightforward: the alternative is not hearing. For a hearing person, the alternative is a headset that costs nothing to remove. **A surgical risk profile that is clearly acceptable against deafness is a very different proposition against the inconvenience of earbuds**, and any paper proposing this must lead with that asymmetry rather than bury it.

But the asymmetry can run the other way, and that argument deserves equal weight. A large elective market is how a surgical safety envelope gets characterised at all. Volume refines technique, drives complication rates down, funds long-horizon outcome data, and amortises the regulatory and manufacturing cost of a device class — and it does so on a population that can consent freely and can wait. Small populations cannot fund that discovery themselves; they inherit it. The history of implantable devices, imaging and semiconductor sensing is substantially a history of small clinical populations receiving mature technology that a larger market paid to make safe.

On that reading, an elective consumer market is not a competitor to the orphan indication but its likeliest route to a well-characterised device — and a population living with progressive loss may weigh the same risks very differently from a person weighing them against earbuds. That judgement is theirs.

Two cautions attach to the inversion rather than defeating it. A safety envelope characterised on healthy young adults does not automatically transfer to altered anatomy and a growing lesion, so the transfer must be shown rather than assumed. And a model in which one population bears the risk of learning while another receives the benefit is an ethical structure that deserves to be named out loud rather than treated as a happy accident.

This document states the asymmetry; it does not resolve it. Whether the trade is worth making is a judgement for patients, clinicians, regulators and the people who would live with the device, and it is not one an engineering scope is entitled to make on their behalf. What a scope owes them is the asymmetry stated clearly, in both directions, with the numbers where numbers exist.

4A. Does this help retrocochlear patients?

Partly, and not in the way the rest of this document might suggest.

Not on the affected side. Any auditory channel — acoustic, ultrasonic, electrical — depends on the nerve it is delivering into. A compressed nerve degrades all of them equally, and nothing in this capability bypasses it.

But there is a working ear, and that changes what is available. A person with one intact side has a full-bandwidth auditory channel that current single-sided solutions use only to relay sound. Routing devices move the dead side's audio to the good ear and cannot restore spatial hearing, because the classical binaural cues need two cochleas.

Sensory substitution offers a different answer to the same problem, and it is the one the evidence supports. Instead of trying to reconstruct binaural cues that cannot form, deliver an explicitly *learned* spatial code through the working ear — the same approach that lets blind users navigate obstacle courses with a visual-to-auditory device. The code does not have to imitate anything natural. It has to be learnable and it has to leave the ear canal open, which is why that literature recommends bone conduction in the first place.

That reframes single-sided deafness from a restoration problem into an encoding problem, and the encoding problem has a research community, a method, and cheap non-implanted hardware to test on. It is also the one route in this whole family of documents that does not require the compressed nerve to do anything.

4B. Encoding a spatial code into a single working ear

If single-sided deafness is an encoding problem rather than a restoration problem, the design question becomes what code to send. One structure is worth setting out because it is quantitative and testable.

Interleave. Assign alternating frequency bands to two purposes: ordinary speech and environmental sound in one set, an explicit spatial code in the other. The question is whether there is spectral room, and there is. A normal cochlea resolves roughly 24 to 30 critical bands, while noise-band vocoder work established that speech remains highly intelligible on a handful of spectral channels. Splitting alternate bands leaves about 14 for speech — comfortably above what speech needs — and about 14 for the code.

Fourteen channels is a great deal for spatial information. It would not reproduce the roughly one-to-two degree azimuth resolution of normal binaural hearing, but a coarse code of eight to sixteen zones would exceed what a single-sided listener has now, which is close to nothing.

Spend the gaps the brain already fills. Phonemic restoration is a documented effect: interrupted or masked speech is perceptually reconstructed from context. The spectral cost of interleaving is therefore lower than a naive channel count suggests, because some removed bands are supplied by prediction rather than lost. That is the auto-complete argument, and it is the reason to prefer a comb-like interleaving over simply compressing speech into half the spectrum.

Space the code logarithmically. A constant-ratio comb has constant spacing on the cochlear place axis, so an interleaved code sits evenly across the partition; a constant-delay comb does not, and crowds the code into one cochlear region while leaving another untouched. Over 500 Hz to 8 kHz the place spacing of a linear comb varies by a factor of 6.7, against 1.23 for a constant-ratio one.

One further cost, specific to music. Cochlear implant users' difficulty with music is well documented and is generally attributed to coarse spectral resolution and weak pitch coding rather than to any deficit of the listener. Speech survives on few channels because it is redundant and predictable; music is neither, and it depends on

precisely the fine spectral structure an interleaved code would consume. A scheme that costs a listener little speech intelligibility may still cost them music, and that trade should be measured rather than assumed away.

What would kill it, and all three are testable through headphones. Interleaving may degrade speech *in noise* even where it does not in quiet, because noise removes the redundancy restoration depends on — which would impose the cost precisely where single-sided listeners are already worst. The code may prove unlearnable within the training time people will actually give. And listeners may reject the altered timbre of their only good ear, which is a real risk given what non-use rates show about devices that alter hearing a person already has.

None of this requires an implant to test. It requires a bone conduction headset, normal-hearing listeners, and a training protocol.

5. Applications worth naming

Assistive, where the case is strongest. A blind user's spatial-mapping instrument rendering obstacle geometry to sound while the ear canals stay open for ordinary hearing; a deafblind user for whom the auditory channel is the remaining high-bandwidth route; a person whose hands and eyes are committed and whose alternative is nothing at all. In each, the benefit is capability rather than convenience, and the risk calculus changes accordingly.

Sensory extension. Rendering to sound what a person cannot otherwise perceive: proximity to energised conductors, gas concentration, structural vibration, radiation dose rate. There is a real precedent in industrial practice — audible alarms already do this crudely, and a continuous encoded stream is a better instrument than a threshold beep.

Navigation without light. Echo-ranging tied to head direction is the obvious case and it has a natural model, since human echolocation by trained blind users is a documented capability rather than a curiosity.

Communication. The telephony case is the one that would drive volume, and it is the one with the weakest benefit-to-risk ratio and the strongest misuse profile. Those facts should be stated together.

5A. Prior work this builds on, credited

Ultrasonically powered millimetre-scale implants. Work on neural dust and its stimulating successor established that a cubic-millimetre implant can be both powered and addressed over a single ultrasonic link, with data returned by backscatter. That work was encountered here during a reference check for a different paper rather than as a starting point, and the systems differ substantially — but two things from it are adopted with credit. The first is a scale reference: it establishes how small active in-body hardware can actually be built, which bounds several arguments in this document. The second is the **active-probing pattern** — ping and read, rather than transmit and hope — which is a useful metaphor for any implanted element operating on an external carrier, and a direction worth pushing further than this document does.

A note on terminology, offered as a distinction rather than a criticism. Those units are *notes*: they carry delegated intelligence on board, which is a legitimate and well-earned design point for the environments they were built for. Units intended for a bloodstream sit at the opposite, near-zero-agency end of the same axis, because a bloodstream tolerates almost no autonomy. The difference is one of mission scope.

Citation to be completed at deposit: the primary references should be pulled and verified rather than carried forward from a secondary mention.

6. Dual-use analysis

Published with the capability, not after it. The purpose is to let regulation be written on the front end rather than in response to an incident.

A note on scope. What follows is a taxonomy of risk categories and their governance implications. It deliberately does not include implementation detail for any misuse pathway. A risk analysis and a design guide are different documents, and only one of them helps a regulator — naming a hazard clearly is what enables a rule; specifying how to build it is not.

6.1 The property that generates most of the risk

It cannot be taken off. Every governance mechanism that works for a phone — leave it at home, switch it off, put it in a drawer, hand it to someone else, throw it away — depends on the device being separable from the person. An implant is not. Consent given at implantation is consent that persists, and withdrawing it requires surgery.

Nearly every category below is a consequence of that single property.

6.2 Categories

Coercive and captive delivery. A channel the user cannot mute or remove is a captive audience in the strict sense. Advertising, ideological content, or harassment delivered through it has no user-side remedy. In an abusive relationship, an institutional setting, or a custodial one, the asymmetry is severe.

Surveillance. A device with a microphone that is surgically attached to a person is a recording device that cannot be left behind. Even absent recording, the metadata of a body-worn transceiver is a location and association record.

Coerced or pressured adoption. If the capability confers advantage — in employment, in the military, in competitive settings — pressure to adopt follows, and consent given under economic pressure is a weak form of consent for an irreversible surgical procedure.

Security. Wireless implanted medical devices have a documented history of vulnerability, and the consequences here include not only data compromise but direct sensory access. An attack surface that terminates in a person's perception is a different class of problem from one that terminates in their data.

Incapacitation. A channel into the auditory system with controllable output is a channel that can be used to disable, disorient, or deprive of sleep. This must be addressed by hard output limits in hardware rather than by software policy.

Perceptual authenticity and evidentiary integrity. If synthetic content is delivered through the same pathway as real sound, a wearer's testimony about what they heard becomes contestable, and their own confidence in it may not be reliable. This has consequences for consent, for legal process, and for mental health that are not addressed by any existing device framework.

Equity and stratification. A surgical augmentation available by purchase creates a capability divide that is not merely economic but embodied.

Irreversibility and the exit problem. A person who wants out faces surgery, cost, and risk. Any framework must treat explantation as a right with a funded pathway, not as an unexamined possibility.

6.3 Governance implications, stated as requirements

Written as things a regulator could require, because a risk named without a remedy is only an alarm.

- **Hardware-enforced output ceiling**, below any incapacitating level, not defeatable in software or firmware.
- **A physical, user-controlled mute** that interrupts the signal path independently of software — the implant equivalent of a lens cap.
- **Mandatory, funded explantation pathway** as a condition of approval.
- **A prohibition on non-consensual content classes** — advertising, political messaging, third-party content the user has not requested — enforced at the device authorisation layer rather than by policy.
- **Security requirements at the level applied to implanted cardiac devices**, including authenticated command channels and a documented disclosure process.
- **Employment and custodial protections** making adoption non-coercible, drafted before a market exists rather than after.
- **Distinct regulatory classes for restorative and elective use.** The evidence bar, consent process and post-market surveillance appropriate to restoring hearing are not appropriate to elective augmentation, and merging them would let the elective case inherit the restorative case's precedents.

7. Honest position

The capability is largely assembled from parts that already exist, so the engineering barrier is lower than the discussion usually assumes. The perceptual barrier — encoding and training — is real and is where a serious programme would start, using non-implanted hardware.

The barrier that should actually govern near-term work is the risk asymmetry of Section 4: surgical risk for convenience is a bad trade, and the applications where it becomes a good trade are assistive ones for people with a genuine capability gap. **That is where this belongs, and a paper proposing it should say so plainly rather than lead with the consumer vision.**

The dual-use profile is serious, it follows almost entirely from a single property — that the device cannot be removed — and it is best addressed by requirements written before a market exists. Publishing the misuse analysis alongside the capability is the point: a hazard that regulators learn about from an incident gets regulated badly, and a hazard they learn about from the people proposing the technology can be regulated well.

References

Retrieved and verified during preparation. Entries marked **[partial]** were identified by title and source but their complete bibliographic record was not obtained; they are listed with enough detail to be found. No author list, volume or year appearing below without that marker was inferred.

Implanted bone conduction devices

1. Mylanus, E.A.M.; Hua, H.; Wigren, S.; et al. Multicenter clinical investigation of a new active osseointegrated steady-state implant system. *Otology & Neurotology* **2020**, 41(9), 1249–1257. doi:10.1097/MAO.0000000000002794
2. Goycoolea, M.; Ribalta, G.; Tocornal, F.; et al. Clinical performance of the Osia system, a new active osseointegrated implant system: results from a prospective clinical investigation. *Acta Oto-Laryngologica* **2020**, 140(3), 212–219. doi:10.1080/00016489.2019.1691744
3. Willenborg, K.; Avallone, E.; Maier, H.; Lenarz, T.; Busch, S. A new active osseointegrated implant system in patients with single-sided deafness. *Audiology & Neurotology* **2022**, 27(1), 83–92. doi:10.1159/000515489
4. Nowlen, et al. Osia implantation in patients with prior craniotomy for skull base tumors. *Otolaryngology-Head and Neck Surgery* **2026**. doi:10.1002/ohn.70283
5. Systematic review and meta-analysis of a new active transcutaneous bone conduction implant, 14 studies, 15 cohorts, n = 314. **[partial]**
6. Placke, L.; Appelbaum, E.; Patel, A.; Sweeney, A. Bone conduction implants for hearing rehabilitation in skull base tumor patients. *Journal of Neurological Surgery Part B: Skull Base* **2019**, 80(2), 139–148. doi:10.1055/s-0039-1677690

Ultrasonically powered implants

7. Seo, D.; Neely, R.M.; Shen, K.; Singhal, U.; Alon, E.; Rabaey, J.M.; Carmena, J.M.; Maharbiz, M.M. Wireless recording in the peripheral nervous system with ultrasonic neural dust. *Neuron* **2016**, 91(3), 529–539. doi:10.1016/j.neuron.2016.06.034
8. Piech, D.K.; et al. A wireless millimetre-scale implantable neural stimulator with ultrasonically powered bidirectional communication. *Nature Biomedical Engineering* **2020**, 4(2), 207–222. doi:10.1038/s41551-020-0518-9

Psychoacoustics of the interleaving proposal

9. Shannon, R.V.; Zeng, F.-G.; Kamath, V.; Wygonski, J.; Ekelid, M. Speech recognition with primarily temporal cues. *Science* **1995**, 270(5234), 303–304. doi:10.1126/science.270.5234.303 — *high speech recognition was obtained with only three bands of modulated noise; this is the source of the spectral-headroom argument in §4B.*
10. Phonemic restoration: perceptual reconstruction of interrupted or masked speech. **[partial — classic literature identified by description only]**

Sensory substitution and expansion

11. Pilot study of a visual-to-auditory navigation aid in blind participants (n = 7), recommending bone conduction so that the ear canals remain unblocked. **[partial]**

12. Haptic sensory substitution and expansion, including infrared and ultraviolet sensing through a wrist-worn device; and spinal cord stimulation trialled as a route for sound perception (2024). **[partial]**

Device use and non-use

13. Thirty-year single-institution series of 356 implant recipients reporting cochlear implant non-use at 2.8% and middle-ear implant non-use at 27.3%; and poor long-term outcomes in up to 16% of adult cochlear implant recipients. **[partial — the second figure was cited within a secondary source]**
-

Declaration

This is a preprint. It has not been peer reviewed. The literature pass is limited and incomplete: searches were targeted and natural-language against general web indices, and did **not** include structured database searching, systematic patent searching by classification code, clinical trial registries, or non-English literature. Where a claim is described as not found, the correct reading is *not found by these searches*, not *absent from the art*.

Prepared by a part-time, unaffiliated research effort, published free to the world. An AI language model (Anthropic Claude) assisted with structural drafting, literature recall and targeted search. No automated tool generated any observation from a patient, device or instrument. The author is responsible for the correctness of all content and citations.

Deliberate omission. Section 6 states misuse categories and their governance remedies without implementation detail for any misuse pathway. That omission is intentional and is described in §6 itself.

Table 1. Provenance table.

quantity or claim	tag	source / note
active transcutaneous mastoid implant (piezoelectric actuator on osseointegrated screw, magnetic external processor) is a shipping regulated device	SOURCED ×4	multicentre clinical investigation; systematic review and meta-analysis (14 studies, n=314); single-sided-deafness series; skull-base and paediatric series
CI users' music perception is poor relative to speech	SOURCED ×1	widely reported; primary literature not opened here
transcutaneous bone conduction implants FDA-approved 2013	SOURCED ×1	clinical outcomes review
direct digital audio streaming into implants is commercial	SOURCED ×2	manufacturer product documentation; paediatric streaming-training study
visual-to-auditory navigation aids trialled in blind participants	SOURCED ×1	pilot study, n=7, obstacle course and object recognition
bone conduction recommended for sensory substitution to keep canals unblocked	SOURCED ×1	same pilot study, stated in its discussion
sensory substitution limits are perceptual — training burden, non-intuitive magnitude coding	SOURCED ×1	pilot study's reported limitations
CI non-use 2.8%, middle-ear implant non-use 27.3% (356 recipients, 30 years)	SOURCED ×1	single-institution retrospective, La Paz
up to 16% of adult CI recipients have poor long-term outcomes	SOURCED ×1	cited within a sensory-substitution paper; primary not opened
sensory expansion demonstrated in normally-sensed people (IR/UV haptic wristband)	SOURCED ×1	haptic sensory substitution/expansion review
spinal cord stimulation trialled for sound perception	SOURCED ×1	2024 sensory-substitution study
ultrasonically powered cubic-mm implant, backscatter readout	UNVERIFIED at primary source	described from a secondary mention; primary references to be pulled and checked before deposit
~24–30 critical bands; speech intelligible on a handful of vocoder channels	UNVERIFIED	standard psychoacoustic results, not re-checked here
phonemic restoration	UNVERIFIED	well-documented effect; primary not opened
comb place-spacing figures ($6.7\times$ vs $1.23\times$)	COMPUTED	Greenwood map, same calculation as the companion paper
temporal bone surgical risk categories	UNVERIFIED	standard otologic risk categories, rates not checked and deliberately not quoted
human echolocation as a trained capability	UNVERIFIED	widely reported; primary literature not checked here
implanted medical device security vulnerability history	UNVERIFIED	documented in the security literature; not checked for this document