2026 · LOW-RESOURCE AUDIO CODEC CHALLENGE

Call for Participation

LRAC 2.0 builds on the first edition with **two transparency-focused tracks** — and three new themes pushing low-resource neural audio coding toward transparency, across languages.

Overview

The growing demand for speech interfaces on resource-constrained devices calls for codecs efficient in **compute, memory, bitrate, and latency** — in real-world conditions without sacrificing quality. Recent neural methods optimize these constraints individually; meeting them *together*, at deployment quality, remains unresolved. LRAC 2.0 invites you to help close that gap.

New in LRAC 2.0

1. **Bitrate scaling toward transparency**
   Beyond low and ultralow rates: can low-resource neural codecs keep improving with bitrate and reach transparency? The first common benchmark to find out.

2. **A dedicated memory-efficiency track**
   Compute counts alone don't predict real-time behavior. New **Track 2** adds a hard parameter cap and a theoretical memory-bandwidth penalty, so leaner solutions rank higher at comparable quality and bitrate.

3. **Initial multilingual coverage**
   Training data expands in size and language diversity; evaluation extends beyond **English (ENG)** to **Spanish (SPA)** and **Mandarin Chinese (CMN)**.

The Challenge — Two Tracks

Both tracks ask you to develop **low-complexity, low-latency** speech codecs that achieve perceptual transparency of input speech as bitrate increases — including under mild noise and reverberation, and across multiple languages. They share the same constraints, data, and evaluation battery, and differ only in how model efficiency enters the ranking.

TRACK 1 Transparency Codecs

No parameter cap and no memory-bandwidth penalty. Systems are ranked purely on perceptual quality and intelligibility — minimizing degradation while meeting the complexity, latency, and bitrate constraints, and revealing which designs scale with bitrate toward transparency.

TRACK 2 Memory-Efficient Transparency Codecs

Everything in Track 1, with efficiency built into the ranking — rewarding leaner solutions at comparable quality and bitrate:

- **Hard cap of ≤1.5M parameters** — a firm eligibility constraint.
- **Theoretical memory-bandwidth penalty** applied to the ranking score.

FOCUS AREAS (BOTH TRACKS)

- Quality scaling across the bitrate caps to transparency in clean speech
- Robustness to mild real-world distortions (background noise, reverberation)
- Behavior across ENG, SPA, and CMN

Bring new designs. We explicitly invite diversity beyond the Conv-Encoder + RVQ + Conv/GAN-Decoder pattern that dominated LRAC 1.0 — e.g., hybrid parametric + neural designs, classical-DSP front ends with neural quantizers/decoders, diffusion/flow decoders, scalar/lattice/product quantization, multi-rate architectures, and neural-augmented classical codecs.

Data

TRAINING

A mandated, fixed public dataset — expanded to **~1,200 h** after curation (~2× LRAC 1.0), spanning **12 languages** with a roughly balanced ENG / non-ENG split. Sampling rate **24 kHz**. Training, validation, and tuning strictly on the designated challenge data.

BLIND TEST

A newly collected, withheld test set. Evaluation matches LRAC 1.0 Track 1 across five bitrate caps **{1, 6, 12, 24, 32} kbps**. Multilingual: clean-speech quality for SPA & CMN at higher rates; intelligibility for ENG, SPA, CMN. All evaluation languages appear in training.

Rules at a Glance*

≤700 MFLOPS TOTAL COMPLEXITY	**≤300** MFLOPS RECEIVE-SIDE
≤30 ms TOTAL LATENCY	**24** kHz INPUT / OUTPUT
CBR CONSTANT BITRATE ONLY	**1·6·12·24·32** KBPS BITRATE CAPS · ONE SYSTEM
TRACK 2 ONLY **≤1.5M** params PARAMETER HARD CAP	TRACK 2 ONLY **Memory-bandwidth penalty** APPLIED TO RANKING

- Systems must support all five bitrate caps within a single system (and a single decoder) — no variable or adaptive bitrate or entropy-coding optimization.
- Training, validation & tuning strictly on the designated challenge datasets.
- Each entry: processed test stimuli + a detailed system report + a completed complexity, latency, parameter count & memory-bandwidth spreadsheet (templates provided).
- Entry-submission instructions follow registration.

* See the Challenge website for the complete rules and conditions: lrac.short.gy/rules

Open Source & Awards

Open Source
Open-sourcing is encouraged but not required.

Innovation Award
For the submission judged most innovative by the organizers.

Open Codec Awards
Open weights, open inference code, and a research-permissive license are the minimum bar; the top-ranked eligible system wins. If no entry qualifies, no award will be made. Training code release encouraged, but not required. One award per track.

Important Dates

Date	Milestone
Jul 27, 2026	SP Grand Challenge proposals deadline
Aug 7, 2026	SP Grand Challenge acceptance notifications *(accepted)*
Aug 25, 2026	Registration opens
Sep 1, 2026	Challenge runs — start
Nov 17, 2026	Registration closes
Nov 30, 2026	Challenge runs — end; withheld test stimuli released
Dec 1, 2026	Processed test stimuli, system reports and spreadsheets due
Dec 21, 2026	Notification of final results
Jan 7, 2027	2-page invited papers due
Jan 21, 2027	2-page paper acceptance notification
Jan 28, 2027	Camera-ready 2-page paper deadline
May 2027	ICASSP 2027, Toronto — SP Grand Challenge session^

Dates are tentative and subject to change. ^Dissemination through the SP Grand Challenge at ICASSP 2027 in Toronto, including up to five invited two-page papers eligible for extension into peer-reviewed SP-OJ publications (details on the official ICASSP website: https://2027.ieeeicassp.org).

Organizers

CISCO SYSTEMS

Kamil Wojcicki · Australia
Yusuf Isik · Türkiye
Mansur Yesilbursa · Poland
Ivana Balic · Switzerland
Erfan Loweimi · UK
Anna Bartlett · UK
Rafał Łaganowski · Poland
Nezih Topaloglu · Türkiye

CO-ORGANIZERS & ACADEMIC PARTNERS

Shinji Watanabe
Carnegie Mellon University, USA

Minje Kim
University of Illinois Urbana-Champaign, USA

Yossi Adi
Hebrew University of Jerusalem, Israel

Evaluation

Participants process open and withheld test sets and submit the stimuli. Both objective metrics and crowdsourced subjective tests are used; objective metrics are informational only. **Final rankings rest on subjective scores** — accounting for speech quality scaling toward transparency and speech intelligibility. **In Track 2**, the score is additionally adjusted by a theoretical memory-bandwidth penalty favoring leaner solutions at comparable quality and bitrate. The same battery below applies to both tracks.

	Condition	Languages	# Files	Bitrates (kbps)	Goal	Test	Test Range	Weight
a.	**Clean speech** (ENG)	ENG	60	1, 6, 12, 24, 32	Quality scaling toward transparency	MUSHRA-1S	[0, 100]	25%
b.	**Clean speech** (multilingual)	SPA, CMN	30, 30	12, 32	Multilingual clean-speech quality	MUSHRA-1S	[0, 100]	10%, 10%
c.	**Real-world** light noise & reverb	ENG	90	1, 6, 12, 24, 32	Robustness to light noise & reverb	DCR	[1, 5]	30%
d.	**Simultaneous talkers**	ENG	20	12, 32	Preservation of overlapping speakers	DCR	[1, 5]	10%
e.	**Speech intelligibility** (clean)	ENG, SPA, CMN	TBA	1	Multilingual intelligibility	DRT	[−100, 100]	5%, 5%, 5%

Subjective evaluation battery — applies to both tracks. Each listed bitrate is a cap: an upper limit a system's constant rate may meet but not exceed (for example, 16 kbps satisfies the 32 kbps cap; 48 kbps does not). Clean-speech MUSHRA-1S assesses transparency against the hidden clean reference, with results for Opus at 32 kbps reported as a strong conventional-codec baseline. The framework accommodates full-battery testing of at least 7 submissions per track. If more are received, an objective-metric-based qualification round, detailed in the Official Rules, determines which submissions advance to the full evaluation. Rankings are based on a weighted battery of the subjective evaluation results. Track 1 is ranked on these subjective scores alone; Track 2 adds a memory-bandwidth penalty, favoring solutions with lower memory-bandwidth requirements at comparable quality and bitrate. The highest-ranked submission in each track is declared the winner. No financial prizes are awarded.

GET INVOLVED

Join the LRAC 2.0 Challenge

Building practical speech codecs, researching novel neural architectures, or pushing low-resource coding toward transparency across languages? Register and help shape the next generation of low-resource communication systems.

Register → lrac.short.gy/lrac2.0

Questions & Slack invite
lrac-challenge@cisco.com

Scan to register
lrac.short.gy/lrac2.0