

Classical vs neural, the four named approaches, what to pick, and the voice-vs-noise trade-off.

## Two families

<b>Classical (subtract)</b>	Measure noise in the gaps, subtract it	Steady noise only; leaves musical noise
<b>Neural (keep voice)</b>	Pre-trained model keeps learned voice	Removes sudden noise; costs CPU/GPU + model

## The four named approaches

<b>WebRTC built-in NS</b>	Classical, free, in every browser	Fan / AC; misses barks and clatter
<b>RNNoise</b>	Hybrid GRU, ~85 kB, any CPU	~60x real time x86; ~10 ms look-ahead; BSD
<b>Krisp (SDK)</b>	On-device neural, Small + Big models	No audio upload; ~25 ms latency; near-field
<b>NVIDIA noise removal</b>	GPU model (Maxine / Broadcast)	Bigger model off the CPU; needs NVIDIA GPU

## Choose by use case

- General meeting, mixed devices
- Contact centre (open-plan)
- Telemedicine / finance (privacy)
- Podcast / recorded, no time limit
- Streamer on a GPU workstation
- Embedded / low-power device
- Music / instrument capture

- WebRTC NS, or RNNoise**
- On-device neural CPU model**
- On-device neural only**
- Largest neural model, offline**
- NVIDIA noise removal**
- Classical, or RNNoise if it fits**
- Turn suppression OFF**

## The voice-vs-noise trade-off

SNR = how far the voice (dB) sits above the noise (dB). Calls sit nearer 20 dB than 0 dB.  
More suppression removes more noise AND more voice - it is never free.  
Example: voice at 10 dB SNR, pull noise down 18 dB -> effective SNR ~28 dB (comfortable).  
Push too hard and consonants (t, k, s) get cut: lower noise, muffled voice - a bad trade on a call.  
Judge it with ITU-T P.835: SIG (voice damage), BAK (residual noise), OVRL (overall).

## Remember

- Classical suppression handles steady noise for free; it cannot remove a bark, a door, or other voices.
- Neural suppression has learned what voice looks like, so it removes sudden non-stationary noise.
- RNNoise is a small ~85 kB hybrid model that runs many times faster than real time on a plain CPU.
- Krisp runs on-device (no audio upload) with Small/Big models; NVIDIA runs a bigger model on the GPU.
- Match the model size to the device: a Big model stutters on a thin laptop or phone.
- Test with the P.835 SIG/BAK/OVRL split so a good noise number is not hiding voice damage.