

The 6 dB-per-bit rule, dynamic range, dither, and which depth to use per task.

The one rule: ~6 dB per bit

Each added bit lowers the quantization noise floor ~6 dB, widening dynamic range.

$$\text{SNR} = 6.02 \times N + 1.76 \text{ dB} \quad (\text{N} = \text{number of bits})$$

Dynamic range by depth

16-bit	~96 dB	Enough for human hearing once dithered. Deliver here.
24-bit	~144 dB	Headroom for recording, not better delivery. Real DACs ~123 dB.
32-bit float	~1528 dB	Can't digitally clip; capture first, set level later.

Human hearing ~120 dB threshold of hearing to threshold of pain
Dithered 16-bit reaches ~120 dB perceived range. More bits add no audible detail.

Dither — the final-reduction rule

- Reducing to 16-bit? Add TPDF dither BEFORE the cut, never truncate (chop bits).
- Truncation = grainy distortion in quiet fades; dither = smooth, inaudible hiss.
- Dither once, on the final reduction only. 24-bit / 32-bit float need no dither.

Which depth by task

Your task	Use	Why
Capture you can't re-record (field, live)	32-bit float	Can't clip; no level-setting
Studio recording & mixing in a DAW	24-bit	Headroom in; 32-bit float to process
Delivery to listeners (stream, download)	16-bit	Enough for ears; half the size
Archival master of source	24-bit	Future-proof headroom
Real-time voice (WebRTC, conferencing)	16-bit	Opus is 16-bit; saves CPU
Telephony / legacy narrowband	8-bit	Set by codec (G.711)