

TELESCOPE

ANALYSIS

FREE · AGPLV3

VST3 · AU

TELESCOPE puts **thirteen lenses** over one analysis engine: loudness to the broadcast standard, dynamics, spectrum and sonograms, notes and key, stereo and pan, and your tonal balance against a reference. The thirteenth, **VERDICT**, turns what was measured into sentences — each one with the number that produced it and the id of the rule that fired. It is **not a mix assistant**: it measures, it doesn't judge. **Measurement, not taste**. And it never touches your audio: **0 samples of latency, bit-exact pass-through**.

// ANATOMY



- 1 **Bypass** — the power icon: on or off, the audio passes bit-exact

3 **A / B** — hold two settings and flip between them to compare

5 **Lens strip** — thirteen lenses, one click each; only what the open lens needs is computed

7 **Lens settings** — each lens has its own row; here LINES · TILT · HISTORY · RANGE · PALETTE
- 2 **Presets** — ◀ name ▶ and SAVE; the palette and the language travel with the state

4 **S · M · L** — three window sizes, each with its own metrics

6 **The open lens** — WATERFALL here: level as height, time receding, a 2.5D projection drawn in software

8 **Theme · Language** — THEME switches dark and light; LANGUAGE, the whole plug-in between six languages

// HOW IT WORKS

Measured, with the test next to it.

Every measurement quoted here comes from a test in the plug-in's repository, run against a synthetic signal — and what is not measured is not listed.



// CONTROLS — COMPLETE REFERENCE

LENSES

LOUDNESS — Integrated, short-term and momentary LUFS, LRA and true peak — also per channel. Against ITU-R BS.1770 gating, EBU Tech 3342 for LRA and the literal FIR table of BS.1770 Annex 2 for true peak.

DYNAMICS — How much headroom is left: PSR over the last 3 s, PLR since reset (AES TD1004), a 61-bin short-term histogram and clip events on a 10-minute timeline. Silence stays out of the histogram: only short-term values above -70 LUFS count.

SPECTRUM — FFT from 1 024 to 32 768 points with Hann, Blackman-Harris or Kaiser, and free, ⅓-octave (ISO 266) or Bark bands — every bin in dBFS referred to a full-scale sine, with an exact coherent-gain correction.

SPECTROGRAM — Level over time: 512 log rows from 20 Hz to 20 kHz, 10 / 30 / 60 s of history and level as colour — the instantaneous power, never an average.

WATERFALL — The same data in depth: level as height, time receding from now. A 2.5D oblique projection drawn in software — no GPU — with an occlusion proven identical to the painter's algorithm.

CQT — The spectrum by note: constant-Q at 24 bins per octave from A0, a chromagram and a key estimated against the Krumhansl & Kessler profiles — never shown without its confidence and its share of time.

SPIRAL — The same constant-Q coiled, one turn per octave: equal notes line up on one radius, so the octave structure of what is playing reads at a glance.

SCOPE — Correlation, width, balance and mono loss over 100 / 300 / 1000 ms, as Lissajous, POLAR SAMPLE or POLAR LEVEL: a half-circle with one ray per degree, where what is out of phase folds onto the base and prints its percentage.

BAND CORRELATION — SCOPE's figures per ⅓-octave band, so mono bass and open highs stop averaging into one number. Its band-wide figure, computed from the STFT bins, matches the time-domain meter to 0.00076 in correlation.

STEREO SPECTROGRAM — The sonogram with phase as colour — red out of phase, green wide, white mono — and level as brightness. A cell with no energy is black, not dim red.

FIELD — Energy by pan direction × frequency on a 64 × 96 grid, with decay and a trail, in a 2.5D view. The label under its axis cannot be switched off: **ENERGY BY PAN DIRECTION · NOT LOCALISATION**.

TONAL BALANCE — Your mix against a reference track you load, each normalised by its own integrated LUFS: tilt at equal loudness, so the delta sums to zero — checked band by band to 0.006 dB. Drag on the reference's waveform to compare against just one section of it; the section is saved with the session.

VERDICT — 21 deterministic rules over a per-second history: within range, how it will feel, where it translates, and what to check and where — every line with its number and its rule. Until the window has had audio above -70 LUFS it says “waiting for audio”, and nothing else. It reflows at 125 % and 150 % in all six languages, with its accents and a MODE button that names the whole mode. **MEASUREMENT, NOT TASTE.**

GLOBAL & I/O

LENS STRIP — Thirteen lenses on the left, one click each. Only the modules the open lens needs are computed — loudness always, so its integrated value never has a hole.

SIZE · S/M/L — Three window sizes, each with its own metrics — not the same drawing scaled.

THEME — At the foot of the strip: dark (the default) or light, remembered per user. In the light theme the four data lenses keep their dark screen and palette.

LANGUAGE — Under THEME: switches the whole plug-in between six languages, each shown by its own name.

HOVER READOUT — Every plot gives a readout from the data, never from the pixel colour; where there is no data, there is no readout.

PRESETS · A/B — ◀ name ▶ and SAVE for the bank; A/B holds two settings to compare. The palette and the language travel with the state.

BYPASS — The power icon. On or off, the audio passes bit-exact: TELESCOPE never processes it.

// LANGUAGES

CODE	LANGUAGE	STATUS	NOTE NAMES
EN	English	reviewed (default)	letters (C D E)
ES	Español	reviewed	sofège (Do Re Mi)
PT	Português	translated · pending native review	sofège
FR	Français	translated · pending native review	sofège
DE	Deutsch	translated · pending native review	German (B = B♭, H = B natural)
IT	Italiano	translated · pending native review	sofège

150 of 150 interface keys and 96 of 96 VERDICT phrases in every language, with a per-key fallback to English that no key needs (0 fall back); all 13 lenses read from it.

// PARAMETERS · EXACT RANGES

		RANGE	DEFAULT			RANGE	DEFAULT
FFT · SPECTRUM		1 024 – 32 768	–	CHROMA		0.5 · 2 · 5 s	–
WINDOW FN	Hann · B-Harris · Kaiser β=9		–	SCOPE WINDOW		100 · 300 · 1000 ms	300 ms
CHANNEL	L · R · M · S · L+R		–	POLAR LEVEL DECAY		12 · 24 · 48 dB/s	–
BANDS	free · ⅓-oct (30) · Bark (24)		–	BAND WINDOW		0.3 · 1 · 3 s	–
SLOPE · DISPLAY		dB / oct	3	FIELD DECAY		0.5 · 1 · 2 s	1 s
AVERAGE	none · τ 0.1–10 s · infinite		–	PALETTE	ovni · inferno · viridis · spectrum		–
HOLD		0–60 dB/s	12	THEME		dark · light	dark
SMOOTHING · DISPLAY	off · 1/24 · 1/12 · 1/6 oct		1/12	LANGUAGE		en · es · pt · fr · de · it	en
CLIP THRESHOLD		–3 ... 0 dBTP	–1.0	INGAIN · OUTPUT (HOST)		not in the UI	0 dB
HISTORY		10 · 30 · 60 s	–	MONOSAFE (HOST)		not in the UI	Off
RANGE		60 · 90 · 120 dB	–	BYPASS		Off · On	Off
WATERFALL LINES		60 · 90 · 120	–				

// WHAT EACH LENS MEASURES – AND ITS TEST

LENS	QUANTITY · STANDARD OR DEFINITION	TEST
LOUDNESS	Integrated / short-term / momentary LUFS · LRA · true peak — ITU-R BS.1770 double gate (–70 LUFS, –10 LU) · EBU Tech 3342 · BS.1770-5 Annex 2 §3 literal FIR	[ebu] Tech 3341 1-5 · 3342 1-4 · [tp] 3341 15-19
DYNAMICS	PSR (true peak of the last 3 s – short-term) · PLR (AES TD1004) · per-channel true peak · clip events	[dyn] 10 bursts → 10 events
SPECTRUM	dBFS per bin, coherent-gain corrected · ⅓-octave ISO 266 (30) · Bark (24)	[spectrum] –20.0000 dB, 3 windows
SPECTROGRAM	Instantaneous power · 512 log rows, 20 Hz – 20 kHz	[spectrogram] p 0.99998
WATERFALL	The same data in depth · oblique software projection, horizon occlusion	[waterfall][horizon] = painter's algorithm
CQT	Constant-Q, B = 24, f_min 27.5 Hz, Q 34.127 · key vs the 24 Krumhansl & Kessler profiles	[cqt] –19.998 / –19.997 / –20.000 dB
SPIRAL	Note → (angle, radius) · constant exactly 0.75	[settings] positionFor to 10 ^{–9}
SCOPE	corr · width · balance · mono loss	[stereo]
POLAR LEVEL	181 one-degree rays · average (τ 0.3 s) + peak-hold (12/24/48 dB/s)	[hemis] mono 1 of 181 · noise 181 of 181
BAND CORRELATION	The 5 sums per ⅓-octave band, over the STFT bins	[bands] Δcorr 0.00076 · Δmono 0.0033 dB
STEREO SPECTROGRAM	Per-bin coherence (colour) + level (brightness)	[sgram-st] L=R 255 · L=–R 0 · silence 128
FIELD	Energy panning × frequency · pan = –cos 2θ (constant power) · 64 × 96 grid	[pan] worst error 4.9·10 ^{–8} · [field]
TONAL BALANCE	Tilt delta at equal loudness · band power by fractional overlap	[tonal] 0.006 dB · 30 of 30 bands
VERDICT	21 deterministic rules · per-second history at 1 Hz	[verdict] 6 languages · file == live · [history] 60/60

// LEVEL-MAP PALETTES

RAMP	WHAT IT IS	MONOTONIC IN LUMINANCE
OVNI	the label's own, computed from the brand colours (not tabulated)	yes
INFERNO	black → purple → orange → yellow · the exact 256 entries of matplotlib	yes
VIRIDIS	dark blue → teal → yellow · legible with red-green colour blindness	yes
SPECTRUM	blue → cyan → green → yellow → red → white	no, on purpose

One ramp shared by SPECTROGRAM, STEREO SPECTROGRAM (its level axis only), WATERFALL and FIELD, stored in the state like the language. In the light theme these four lenses keep their dark screen and the same ramps, so the data reads the same in both themes (D-109). spectrum is exempt by name: there the hue orders the level. inferno and viridis are by Smith and van der Walt (2015), CC0.

// UNDER THE HOOD

LOCK-FREE BUS → ANALYSIS THREAD

processBlock reads L and R, pushes them to a lock-free bus and writes nothing back: 0 samples of latency, 0 tail, bit-exact. An analysis thread drains the bus and computes only the modules the open lens needs — loudness always, so the integrated value, LRA and clip count never have a hole. The audio never waits for a measurement.

STEREO FROM THE SAME SUMS

corr, width, balance and mono loss from ΣLL, ΣRR, ΣLR, ΣMM and ΣSS, broadband and per 1/2-octave band. POLAR LEVEL draws 181 one-degree rays — average in time (τ 0.3 s) and peak-hold — and folds what is out of phase onto the base with its percentage. FIELD bins energy panning on a 64 × 96 grid.

LOUDNESS TO THE LETTER

ITU-R BS.1770 double gate (absolute −70 LUFS, relative −10 LU), LRA per EBU Tech 3342 (P95 − P10), K-weighting recalculated per sample rate from the analogue prototype, and true peak with the literal 4 × 12 polyphase FIR of BS.1770-5 Annex 2 §3 — 4× below 96 kHz, 2× up to 176.4 kHz, 1× from 192 kHz.

VERDICT: A TABLE, NOT A MODEL

21 deterministic rules read a per-second history — 1 Hz, 600 rows, 30 bands plus loudness, stereo, DC and key. Every sentence carries its number and its rule id, the same signal gives the same report word for word, and a file analysis yields the same rows as live, to the bit.

CONSTANT-Q BY NOTE

Brown & Puckette (1992) kernels, B = 24 from A0 (27.5 Hz), Q = 34.127, pruned at 0.0054-max. The key is a Pearson correlation against the 24 Krumhansl & Kessler (1982) profiles, always with its confidence and share of time; silence reports no key.

DRAWN AT PHYSICAL SCALE

The six lenses that rasterise by hand bake their pixel caches at the panel's physical scale. Against the same render at 1× stretched, the gradient energy rises 2.96–4.25× (criterion 2.75), and the test checks every cache was baked at 2.0 in the 2× render.

// MEASURED PERFORMANCE

Pass-through · 13 lenses, bypass on/off	bit-exact · 0 mismatches	Latency · tail	0 samples · 0
Full-scale peak through	1.000000 · identity	EBU Tech 3341 · loudness	tests 1-5 · 15-19
EBU Tech 3342 · LRA	tests 1-4	K-weighting vs the 48 kHz table	max error 8.9·10 ⁻¹⁶
True peak · test 3341-17 at 4×	-6.3160 dBTP vs -6.0	Sine on an exact bin · 3 windows	-20.0000 dB
Spectrogram · monotonic sweep	p 0.99998	CQT · same dB over 3 octaves	-19.998 / -19.997 / -20.000
Band correlation vs time meter	Δcorr 0.00076	FIELD pan · constant power	worst error 4.9·10 ⁻⁸
TONAL BALANCE · zero sum	max 0.006 dB · 21 bands	File analysis vs live	equal to the bit
File analysis vs 0.1.0 · 13 goldens	equal to the bit	60 s of audio analysed in	0.63-1.82 s loaded · test < 6 s
VERDICT · rules · languages	21 · 6		

Conditions — every value above is the output of a test in the plug-in's repository, run against a synthetic signal ([telescope][null], [ebu], [kw], [tp], [spectrum], [spectrogram], [cqt], [bands], [pan], [tonal], [file]). Most measurement runs at 48 kHz; the EBU vectors and the K-weighting also at 44.1 and 96 kHz. File speed was measured again on 0.2 (06dcb69), on the house M4 under load 9.4 to 12.9: 0.63-1.82 s in ten runs. EBU vectors still pending: Tech 3341 7-14 and 20-23, Tech 3342 5-6 — listed as pending, not as passing.

// LATENCY · PER SAMPLE RATE

44.1 KHZ	48 KHZ	96 KHZ	
REPORTED LATENCY	0 samples / 0.0 ms	0 samples / 0.0 ms	0 samples / 0.0 ms

Reported and real latency: 0 samples, and a tail of 0 — the output is the input, sample by sample, compared with == ([telescope][null]). Nothing to compensate.

// TRUE-PEAK OVERSAMPLING · BS.1770 ANNEX 2

SAMPLE RATE	OVERSAMPLING	RESULT
< 96 KHZ	4×	176.4 / 192 / 352.8 kHz
96 - 176.4 KHZ	2×	192 - 352.8 kHz
≥ 192 KHZ	1×	already at that resolution


Known 4× under-read, tabulated by the document itself: test 3341-17 (fs/6 at 60°) reads -6.3160 dBTP against a real -6.0 — inside the standard's tolerance (+0.2 / -0.4 dBTP), and a real under-read.

// THE WINDOW AND YOUR DAW · 0.2


WHAT	IN 0.2	TEST · SOURCE
EDITOR BACKGROUND	baked once — in 0.1.0 it was re-baked on every lens frame, 30 times a second: 12–80 ms per frame, and over a second of the main thread per second at size M on Retina	[editorbudget]
LENSES AT REST	a lens repaints only when something visible changes — in 0.1.0, 11 of the 13 redrew the same silence 30 times a second with the transport stopped	[editorbudget]
125 % AND 150 %	each lens lands on a whole physical pixel — it used to be resampled on every frame: 3–6 ms at M on 125 %, against 0.2–2.3 ms at the origin	[editorbudget] · [ancla]
FRAME RATE	30 fps per lens; VERDICT, which re-reads a report that changes once a second, 12 fps	Lens.h · VerdictLens.cpp

This is the fix that matters most in 0.2: an open TELESCOPE window could stall the whole DAW, and no longer does.

// LIGHT THEME · 0.2



The light theme — chosen at the foot of the strip (THEME) and remembered per user; dark is still the default. Here LOUDNESS: the frame, the labels and the controls go light.



...and the data keeps its screen — WATERFALL in the light theme: SPECTROGRAM, WATERFALL, STEREO SPECTROGRAM and FIELD keep their dark screen and palette, so the data reads the same in both themes.

// SHARPNESS AT PHYSICAL SCALE

LENS	1× STRETCHED	2× REAL	RATIO
SPECTROGRAM	8.53	27.61	3.24
STEREO SPECTROGRAM	4.13	12.92	3.12
WATERFALL	28.79	122.36	4.25
FIELD	2.42	8.70	3.60
SCOPE	178.58	528.17	2.96
SPECTRUM	15.71	64.47	4.11

VISUAL[hd] compares, on the same signal, the gradient energy of the scale-2 render against the scale-1 render stretched — only inside each lens's pixel cache. The criterion is 2.75 because this metric's floor is 2.0; the test also checks that every cache was baked at the paint's physical scale (1.00 in the 1× render, 2.00 in the 2× one). Measured on 0.2 (06dcb69).

// SAMPLE RATES TESTED

WHERE	SAMPLE RATES
MEASUREMENT (MOST TESTS)	48 kHz
EBU VECTORS AND K-WEIGHTING	44.1 and 96 kHz as well as 48 ([ebu], [kw])
TRUE PEAK (ALL THREE REGIMES)	44.1 · 48 · 88.2 · 96 · 176.4 · 192 kHz ([tp])
FILE ANALYSIS	44.1 and 48 kHz ([file])
INstantiate / PREPARE / Lifecycle	44.1 · 48 · 96 kHz × blocks of 64 / 512 / 2048 ([smoke], [lifecycle])
HOST RENDER	up to 192 kHz, verified by aual (11.025 · 22.05 · 44.1 · 48 · 96 · 192 kHz)

No measurements at 96 or 192 kHz are claimed beyond that: what is not in the table was not tested.

// FILE ANALYSIS

QUANTITY	VALUE	TEST
IDENTITY WITH LIVE	to the bit: integrated, LRA, true peak, M/S maxima, the 30 measured bands and the history rows	[file] · [history]
IDENTITY WITH 0.1.0	to the bit: 13 golden dumps of the engine, generated with the 0.1.0 analysis code and compared byte by byte	[golden]
BLOCK-SIZE INDEPENDENCE	blocks of 1 / 7 / 64 / 512 / 4096 → same results to the bit	[file]
SPEED	60 s of audio in 0.63–1.82 s on the house M4 under load (1-minute load 9.4 to 12.9), ten runs on 0.2 with the 06dcb69 test binary; the test requires < 6 s (10× real time)	[file]
CANCELLATION	cancel() returns in < 200 ms (test criterion; measured 2–3 ms)	[file]
FORMATS	WAV · AIFF · FLAC · Ogg always; macOS: whatever CoreAudio reads (MP3, AAC, ALAC); Windows: WMA and MP3	—
SAMPLE RATE	the file's; nothing is resampled	[file]

// MEMORY

STRUCTURE	CAPACITY	BYTES
SECONDHISTORY · VERDICT	1 Hz · 600 rows = 10 minutes	252,008 B
LOUDNESSHISTORY · LOUDNESS	10 Hz · 6,000 points = 10 minutes	48,008 B
CLIPHISTORY · DYNAMICS	1 Hz · 600 points = 10 minutes	2,408 B
SPECTRUMFRAME · ORDER 15	16,385 bins × 2 spectra × 2 arrays · 3-slot TripleBuffer	≈ 790 KB
SPECTROGRAM RING	60 s × 60 columns/s × 512 rows of 1 byte, fixed	≈ 1.8 MB

Measured with sizeof over the real headers. The integrated value and LRA are computed over every block since the last RESET, not over a 0.1 dB histogram: exact, at ≈ 2.3 MB per hour of continuous measurement; RESET empties it. The spectrum buffers are sized for the worst case and do not grow with use.

// SPECIFICATIONS

SYSTEM · FORMATS		AUDIO · PROJECT	
FORMATS	VST3 · AU (macOS) · VST3 (Windows)	SAMPLE RATES	host render up to 192 kHz · see table
SYSTEMS	macOS 11.0+ · Windows 10+ (x64)	LATENCY	0 samples · bit-exact · tail 0
ARCHITECTURE	Universal · arm64 + x86_64	LOUDNESS	BS.1770 · EBU Tech 3341 / 3342
I/O	mono / stereo in → stereo out	TRUE PEAK	BS.1770 Annex 2 · 4× / 2× / 1×
PLUG-IN	VST3 Fx Analyzer · AU auxf TLsc Ovni	MEMORY	per-second history 246.1 KB · 10 min
LANGUAGES	en · es · pt · fr · de · it	LICENSE	AGPLv3 · open-source

// TRANSPARENCY & HONESTY

TELESCOPE does not process: its output is its input, bit-exact. It does not localise sources — FIELD and POLAR LEVEL measure energy by pan direction, with no HRTF, no ITD and no azimuth. Stereo only: no surround, no Ambisonics (EBU Tech 3341 test 6, 5.0 channels, is out of scope). No speech-intelligibility or dialogue metric. No model, no telemetry, and it sends nothing — the only way out is the card that introduces EYEPIECE, on the Mac and once, which opens [ovniaudio.com/eyepiece](#) in the browser only if you click it. VERDIOT is a deterministic threshold table — measurement, not taste — and its per-device checks are generic definitions of each class of system, not curves of real speakers, printed as a forecast. “No findings” is not “it’s finished”. The label’s chassis injects inGain, output and monoSafe, which the host lists and the UI does not show; at their defaults (0 dB, 0 dB, off) the audio is bit-exact.

// INSTALL (FREE & OPEN-SOURCE)

Free & open-source (AGPLv3), for **macOS and Windows**. On macOS, grab the `.pkg` installer from [ovniaudio.com](#) and double-click: it is **signed with a Developer ID and notarized by Apple**, so macOS’s own Installer puts the **VST3** and the **AU** where your DAW looks for them (`/Library/Audio/Plug-Ins`) and asks for your admin password itself — nothing lands in quarantine. The installer ships no standalone app, and not `telescope-measure` either: that one is in the source. **On Windows 10+** it is a VST3 x64 in an unsigned ZIP (SmartScreen may warn): unblock it and copy the `.vst3` folder to `C:\Program Files\Common Files\VST3` . If macOS ever warns you about this `.pkg` , do not bypass it: download it again and check its SHA-256 against `SHA256SUMS.txt` . **AGPLv3 · full source at [github.com/ovniaudio/ovni](#)**.