

VapourWiki — Audion VS Engine preset reference (EN)

Full English reference for all 22 presets across 4 palettes. Starts with a **decision tree** ("which preset for which task"). Followed by detailed per-preset description: what it does, on what footage, key parameters, recommended encoder.

Per-preset docstrings live inside each `.vpy` file in `system_core/presets/<palette>/`.
This document is the navigation-friendly companion.

Decision tree — which preset for which task

Read top to bottom: the first matching case is your preset.

Footage symptom	Preset	Palette
Interlaced legacy (DV, HDV, VHS digitization, broadcast TS)	<code>qtgmc_deinterlace</code>	restoration
NTSC 29.97 fps with 3:2 pulldown (telecine, film → video)	<code>tivtc_ivtc</code>	restoration
High-ISO noise, low-light digital, want to keep detail	<code>mvtools_mcdegrain</code>	restoration
Rainbow / dot crawl on VHS-rip / composite capture	<code>derainbow_decross</code>	restoration
Over-compressed H.264/MPEG (YouTube-rip, WhatsApp, SD broadcast)	<code>deblock_h264_artefacts</code>	restoration
Haze / low contrast / want to "lift" the image	<code>dehaze_local_contrast</code>	restoration
Digital noise only in shadows ("low-light" digital footage)	<code>shadow_denoise_sota</code> ★	precision

Footage symptom	Preset	Palette
Mild uniform noise, hardware-agnostic (no CUDA/OpenCL)	<code>mild_denoise</code>	precision
Chroma noise only (dirty blue channel)	<code>chroma_cleanup</code>	precision
Banding in gradients (sky, walls)	<code>deband_safe</code> or <code>deband_fine_grain</code>	precision
Want one preset "everything" for filmic material	<code>filmic_rebuild</code> ★	precision
Pre-grade prep for DaVinci	<code>pregrade_prep</code>	precision
Clean archival master, no grain	<code>archive_clean</code>	precision
Subtle filmic, safe default	<code>cinematic</code>	film_looks
Specific Kodak 35mm stock (250D / 500T / 50D)	<code>film_35mm</code>	film_looks
16mm organic, lifted blacks	<code>film_16mm</code>	film_looks
Super 8, faded 70s, heavy grain	<code>super8</code>	film_looks
High contrast desaturated (Se7en, Saving Private Ryan)	<code>bleach_bypass</code>	film_looks
VHS / CRT aesthetic	<code>vhs_crt</code>	retro
90s amateur camcorder	<code>camcorder_90s</code>	retro

Pipeline scenarios (multiple presets chained via `apply-profile-batch` or manually):

Scenario	Steps
Night timelapse → film	<code>shadow_denoise_sota</code> → <code>filmic_rebuild</code> → <code>cinematic</code>
VHS archive → restored master	<code>qtgmc_deinterlace</code> → <code>derainbow_decross</code> → <code>mvtools_mcdegrain</code> → <code>archive_clean</code>
Old YouTube-rip → project-usable	<code>deblock_h264_artefacts</code> → <code>dehaze_local_contrast</code> → <code>pregrade_prep</code>
Telecined NTSC DVD → 24p master	<code>tivtc_ivtc</code> → <code>archive_clean</code>

Scenario	Steps
Modern digital → film look	<code>mild_denoise</code> → <code>film_35mm</code>

PRECISION palette — technical pipeline (8 presets)

Menu order = pipeline order: Stage 1 (denoise) → Stage 2 (deband) → Stage 3 (compositions).

Stage 1 — denoise

`mild_denoise` — gentle universal denoise

DFTTest spectral denoiser. Pure CPU, hardware-agnostic — works the same on any processor.

Recommended as a **first pick** when you don't know much about the footage: light noise, small touch-up, low risk.

- Parameters: `strength` = light (sigma=4.0) / medium (8.0) / strong (14.0)
- Encoder: `h264_crf17` or `h264_crf14` for archival quality
- Adobe/DaVinci equivalent: only Temporal NR / Noise Reduction, much coarser result

`shadow_denoise_sota` ★ — flagship shadow denoiser

BM3D with float32 pipeline + smooth luma-mask "shadows only". Noise is suppressed **only** in zones below `shadow_threshold`, the rest of the image is untouched. Chroma planes are always denoised (chroma noise is ugly everywhere). Optional micro-grain return via `grain_back` so shadows don't look "plastic".

- Parameters: `sigma=2.5` (BM3D luma), `use_cuda=0/1`, `grain_back=0.6`, `shadow_threshold=0.20`, `transition=0.10`
- On NVIDIA: `--use-cuda 1` gives 25–35% wall-time gain on 4K
- Encoder: `h264_crf17` / `prores_lt`
- Not in NLE: zone-targeted denoise with smooth blend, no cheap "threshold mask"

`chroma_cleanup` — chroma-only cleanup

DFTTest on planes=[1,2]. Luma untouched. Use when "blue channel is dirty" (typical issue of old compact cameras, low-bitrate AVCHD).

- Parameters: `strength` = light / medium / strong
- Encoder: `h264_crf17`

- Not in NLE: pure chroma-only DFTTest with no luma side-effects

Stage 2 — deband

deband_safe — safe deband

neo_f3kdb with light settings, no grain. Removes banding in gradients (sky, walls, night lighting) without losing sharpness.

- Parameters: `range=15`, `y=64`, `cb=64`, `cr=64`
- Encoder: `h264_crf17` / `h265_crf17`

deband_fine_grain — deband + fine grain restore

Same neo_f3kdb + AddGrain. After debanding a thin monotonic grain is added — prevents the "too clean" look (signature of cheap compression).

- Parameters: `range=15`, `grain_var=1.5`
- Encoder: `prores_lt` (for downstream grading)

Stage 3 — compositions

filmic_rebuild ★ — flagship composition

Full chain: BM3D denoise → neo_f3kdb deband → 3-zone luma grain (shadows / midtones / highlights). Grain is zoned — heavier in shadows (like real film), lighter in highlights. Recommended as a "one preset for everything" on filmic material.

- Parameters: `sigma=2.0`, `use_cuda=0/1`, `deband_range=14`, `grain_shadow=1.5`, `grain_mid=0.9`, `grain_high=0.4`, `shadow_threshold=0.30`, `high_threshold=0.65`
- Encoder: `prores_lt` / `h265_crf17`

archive_clean — neutral archival master

DFTTest + neo_f3kdb. **No grain.** Goal — maximally clean "flat" picture for archival. Don't use if grading is planned (grading works better on grainy material).

- Parameters: `denoise=medium`, `range=15`
- Encoder: `prores_422hq` / `prores_422hq_mxf`

pregrade_prep — DaVinci handoff

Minimum touch: only light DFTTest. No deband, no grain. Goal — give Resolve the cleanest source so the colorist doesn't grade on top of compression artefacts.

- Parameters: **strength=light**
 - Encoder: **prores_lt_mxf** (for Adobe / Avid round-trip)
-

FILM_LOOKS palette — film emulations (5 presets)

Each look uses the shared **audion_lib** library (MTF softening, halation bloom, zone grain, gamma curve, black-lift, desaturation).

cinematic — universal subtle filmic

Safe default. Light MTF softening + halation + micro-grain. Not tied to any specific stock. Works on any modern digital footage.

- Parameters: **intensity=1.0** (range 0..2)
- Encoder: **prores_lt** / **h264_crf17**

film_35mm — Kodak 35mm with stock variants

Emulates real Kodak stocks: 250D (daylight), 500T (tungsten), 50D (fine grain). Each stock has its own gamma curve, balance, grain density.

- Parameters: **stock=250D|500T|50D**, **intensity=1.0**
- Encoder: **prores_lt** / **prores_lt_mxf**

film_16mm — organic, more grainy

Lifted blacks, more pronounced grain, softer than 35mm. Suits documentary / arthouse aesthetic.

- Parameters: **intensity=1.0**
- Encoder: **prores_lt**

super8 — strongest look

Heaviest grain, softest optics, faded 70s. For music video / stylized inserts.

- Parameters: **intensity=1.0**

- Encoder: `h264_crf17` (grain already baked in, aggressive compression OK)

`bleach_bypass` — high contrast desaturated

"Se7en" / "Saving Private Ryan" look. High contrast, silver greys, hot highlights. Strong stylistic choice.

- Parameters: `intensity=1.0`
 - Encoder: `prores_lt`
-

RETRO palette — analog character (2 presets)

`vhs_crt` — VHS / CRT

Chroma bleed (horizontal color spread), soft optics, analog noise. Emulates VHS tape playback on a CRT TV.

- Parameters: `intensity=1.2`
- Encoder: `h264_crf17`

`camcorder_90s` — amateur camcorder

Softer than VHS, slight overexposure, minimal chroma bleed. 90s home video aesthetic.

- Parameters: `intensity=1.0`
 - Encoder: `h264_crf17`
-

RESTORATION palette — heavy guns (7 presets, Phase 18.B + NNEDI3 2x)

The set that is **missing or poorly implemented** in Adobe Premiere / DaVinci Resolve.

`qtgmc_deinterlace` — reference deinterlace

havsfunc.QTGMC (NNEDI3 + MVTools motion estimation). Reconstructs progressive frames from interlaced source via neural-network upscaling of each field + motion-compensated temporal smoothing. **Better than any NLE out-of-the-box.**

- Parameters: `field_order=tff|bff`, `qtgmc-`
`preset=Faster|Fast|Medium|Slow|Slower|Placebo`, `output_fps=single|double`

- When to pick: DV, HDV, VHS digitization, broadcast TS, S-VHS / U-matic transfers
- Encoder: `prores_lt` (for downstream work) or `h264_crf17` (final master)

`tivtc_ivtc` — inverse telecine

TIVTC.TFM (field matching) + TIVTC.TDecimate (drop duplicate). Reference-quality 3:2 pulldown removal: NTSC 29.97i → 23.976p. Restores the original 24p film master from a telecined source.

- Parameters: `pp=6` (TFM post-processor), `cycle=5`, `rdrop=1` (standard NTSC pattern)
- When to pick: NTSC DVD, broadcast prints, digitized film transfers
- Encoder: `prores_422` / `prores_422hq_mxf`

`mvtools_mcdegrain` — motion-compensated denoise

MVTools motion estimation → MDeGrain temporal averaging along motion vectors. **Removes noise WITHOUT detail loss** — what Topaz Video Enhance AI and Neat Video do internally. DaVinci Temporal NR is a coarse implementation of the same; this is reference-grade.

- Parameters: `radius=2` (5-frame window), `thsad=200`, `blksize=16` (HD) / `8` (4K detail)
- When to pick: high-ISO low-light footage, preserving skin texture / fabric detail
- Encoder: `prores_lt` / `h265_crf17`

`derainbow_decross` — NTSC composite cleanup

MVTools-based motion-compensated chroma-only smoothing. Removes rainbow / dot crawl / cross-color on composite-captured footage (VHS-rip, U-matic, BetaSP, S-Video → SDI). Luma untouched.

- Parameters: `strength=0.6`, `blksize=16`
- When to pick: digitized VHS, color "running" on thin lines, mosquito noise on B&W edges
- Encoder: `prores_lt` / `h264_crf17`

`deblock_h264_artefacts` — over-compressed rescue

havsfunc.Deblock_QED — edge-aware deblocker for H.264 / MPEG-2 / MPEG-4. Smooths 8x8 block boundaries, ringing around edges, mosquito noise. Edge-aware — does not touch real detail.

- Parameters: `quant1=24`, `quant2=26`, `aoffset=1`, `boffset=1`
- When to pick: old YouTube-rips, WhatsApp/Telegram re-encodes, low-bitrate SD broadcast TS
- Encoder: `h264_crf17` (after deblock the material is cleaner, can re-encode at higher CRF)

dehaze_local_contrast — clarity without halos

Luma-only local contrast via high-pass + zone-weighted MaskedMerge. Zone-mask (parabolic, peaks in midtones) prevents highlight burnout and crushed shadows. No color shift (luma-only). 16-bit math.

- Parameters: `strength=1.0`, `radius=8` (clarity) / `12-16` (haze removal)
- When to pick: hazy footage, flat low-contrast material, want to "wake up" the picture without destroying highlights
- Not in NLE: Resolve "Dehaze" burns highlights and shifts color; this one has no halos
- Encoder: `prores_lt` (handoff to colorist) or `h264_crf17` (final)

nnedi3_upscale_2x — deterministic non-ML 2x upscale

ZNEDI3/NNEDI3-based enlargement through `nnedi3_resample`. This is the Lite answer for clean, predictable upscaling without MLRT or ONNX models.

- Parameters: `quality=fast|balanced|best`, `chroma=spline36|nnedi3`
- When to pick: SD/HD archive material that needs a clean 2x resize without hallucinated ML texture
- Encoder: `prores_lt` for handoff or `h264_crf14` / `h265_crf14` for final

Smoke / bench status

`install\Bench-AllPresets.{cmd,ps1}` walks every `.vpy` preset through `vspipe -c y4m --end <Frames-1> | ffmpeg -f null` without writing output files. It checks preset parsing, plugin namespaces, frame delivery, and the CPU/CUDA BM3D path when requested.

Current local references for Lite:

- CPU fallback: **22/22 PASS** with `Frames=1`, `Cuda=off` (2026-05-16).
- RTX 5070 validation: **30/30 PASS** with `Frames=1`, `Cuda=sweep` (2026-05-27). The 30 rows are 22 presets plus the extra CPU/CUDA Precision sweep.

```
system_core\powershell\pwsh.exe -NoLogo -NoProfile -ExecutionPolicy Bypass -File
install\Bench-AllPresets.ps1 -ProjectRoot . -InputFile <video> -Frames 1 -Cuda sweep
```


Encoders — my patterns

(Audion default workflow per discussion 2026-04-28)

Quality ladder (same 14 / 17 / 21 tiers across software, NVENC, QuickSync, and AMF): **14 → 17 → 21**, three distinguishable steps without overlap.

Profile	When to use
<code>h264_crf14</code> / <code>h265_crf14</code> ★ default	Semi-lossless archive, master for downstream work. Default Audion since 2026-04-28.
<code>h264_crf17</code> / <code>h265_crf17</code>	"Almost-invisible" lossy, final master
<code>h264_crf21</code> / <code>h265_crf21</code>	Web preview / proxy
<code>prores_lt</code> ★	Audion default ProRes — no keying, for grading and round-trip
<code>prores_lt_mxf</code> ★	ProRes LT in MXF wrapper — for Adobe Premiere / Avid round-trip
<code>prores_422</code>	When you need a bit more bitrate than LT
<code>prores_422_mxf</code>	ProRes 422 in MXF wrapper for Adobe / Avid handoff
<code>prores_422hq</code>	Only if keying is planned
<code>prores_422hq_mxf</code>	HQ + MXF for broadcast / Avid finishing
<code>dnxhr_lb</code> / <code>_sq</code> / <code>_hq</code> / <code>_hqx</code>	Avid-style DNxHR (LB low / SQ standard / HQ high / HQX 10-bit)
<code>h264_nvenc_q14</code> / <code>_q17</code> / <code>_q21</code>	NVENC hardware H.264 on NVIDIA. Removes CPU bottleneck from libx264 — VS filtering doesn't fight encode. CQ — analog of CRF. 8-bit yuv420p.
<code>h265_nvenc_q14</code> / <code>_q17</code> / <code>_q21</code>	NVENC hardware HEVC, 10-bit p010le. Ideal for full pipeline on NVIDIA: VS filtering on CPU/CUDA + encode on NVENC = zero CPU stall.
<code>h264_qsv_q14</code> / <code>_q17</code> / <code>_q21</code>	Intel QuickSync H.264 for Intel iGPU batch/proxy work.
<code>h265_qsv_q14</code> / <code>_q17</code> / <code>_q21</code>	Intel QuickSync HEVC, p010le where supported.
<code>h264_amf_q14</code> / <code>_q17</code> / <code>_q21</code>	AMD AMF H.264 via CQP.

Profile	When to use
<code>h265_amf_q14</code> / <code>_q17</code> / <code>_q21</code>	AMD AMF HEVC, p010le output.

Rule: for any non-final processing → `prores_lt` / `prores_lt_mxf` or DNxHR if the receiver prefers it. For final delivery → `h264_crf14` / `h264_crf17` (software, best density per bit) OR the matching hardware ladder for the host (`nvenc`, `qsv`, `amf`) when wall-time matters.

When to use hardware encode vs software

Scenario	Encoder
Heavy VS pipeline (filmic_rebuild, mvtools_mcdegrain) on NVIDIA	NVENC — CPU is freed for VS filter, total wall-time drops 30-50%
Final delivery where bit density matters most	libx264/libx265 CRF — software encoder is still slightly more accurate at low CRF
Intel iGPU host / laptop batch	QuickSync — good throughput with low CPU pressure
AMD/Radeon host	AMF — same 14/17/21 CQP ladder, avoids CPU encode bottlenecks
Large batch of hundreds of files	NVENC / QSV / AMF — savings accumulate
No working hardware encoder	Software only (<code>h264_crf*</code> / <code>h265_crf*</code>)

NVENC requires: NVIDIA driver R525+. Quality: NVENC on Ada/Blackwell (RTX 40/50) is **comparable** to libx264 medium at the same CQ; on older generations (Pascal/Turing) software CRF wins slightly on bitrate efficiency at the same quality, but NVENC is still many times faster.

More — full reference

- English per-preset reference: docstring in `system_core/presets/<palette>/<preset>.vpy`
- CLI flags: `runtime\python.exe system_core\main.py run --help`
- Profiles (cross-palette combinations): `runtime\python.exe system_core\main.py list-profiles`
- Recursive batch with mirror folder structure: `apply-profile-batch --recursive`

- Stack health check: `runtime\python.exe system_core\doctor.py`
-

Last updated: 2026-05-27 (Phase 18.B + NNEDI3 2x — 22 presets across 4 palettes; CPU fallback smoke `22/22 PASS`, `Frames=1`, `Cuda=off`; RTX 5070 CUDA smoke `30/30 PASS`, `Frames=1`, `Cuda=sweep`)