VapourWiki — Audion VS Engine presets

The 28 presets of Audion VS Engine in 4 palettes: which one for which footage, what each does, its key settings and the encoder that suits it. In the window (`Start.exe`) a preset is a switch in a group of its palette; in the console it is `vs_cli.cmd run --operation <id> --set <field>=<value>`. How to work in the window and the console — USER_GUIDE_EN.md; every field with its range and choices — OPERATIONS_EN.md.

Which preset for which task

MLRT marks a neural preset: it needs vs-mlrt (section «The neural presets and MLRT»).

Footage or task	Preset	Palette / group
Interlaced source: DV, HDV, captured VHS, broadcast TS; combs on motion	`qtgmc_deinterlace`	RESTORATION / Interlace
NTSC 29.97 fps with 3:2 pulldown: film transferred to video	`tivtc_ivtc`	RESTORATION / Interlace
Light, even noise; little is known about the footage	`mild_denoise`	PRECISION / Denoise
Digital noise in the shadows of night and low-light footage	`shadow_denoise_sota`	PRECISION / Denoise
Colour noise only: colour blotches, a dirty blue channel	`chroma_cleanup`	PRECISION / Denoise
High-ISO noise; the texture of skin and fabric must stay	`mvtools_mcdegrain`	RESTORATION / Noise
Very heavy noise where BM3D smears the picture	`dpir_denoise` (MLRT)	RESTORATION / Noise

Footage or task	Preset	Palette / group
Banding in gradients: sky, walls, night light	`deband_safe` or `deband_fine_grain`	PRECISION / Deband
One preset for filmic footage: noise and banding out, a controlled grain in	`filmic_rebuild`	PRECISION / Master
A clean neutral master for the archive, no grain	`archive_clean`	PRECISION / Master
Before grading in Resolve: the least possible touch	`pregrade_prep`	PRECISION / Master
Blocks, ringing, mosquito noise of over-compressed H.264 or MPEG	`deblock_h264_artefacts`	RESTORATION / Artefacts
Rainbow and dot crawl of a composite capture: VHS, U-matic, BetaSP	`derainbow_decross`	RESTORATION / Artefacts
Haze, a flat low-contrast picture	`dehaze_local_contrast`	RESTORATION / Artefacts
Twice the size, the same on every run, no invented texture	`nnedi3_upscale_2x`	RESTORATION / Scale
Drawn animation twice the size with clean lines: 1080p to 3840×2160	`vsmlrt_realesrgan_2x` (MLRT)	RESTORATION / Scale
HD by size but soft in detail, rebuilt at twice the size	`soft_hd_rebuild_2x` (MLRT)	RESTORATION / Scale
A Kodak 35 mm stock: 250D, 500T or 50D	`film_35mm`	FILM LOOKS / Film stock
16 mm: softer, grainier, lifted blacks	`film_16mm`	FILM LOOKS / Film stock
Super 8: the heaviest grain, the faded 70s	`super8`	FILM LOOKS / Film stock
Large-format IMAX: the finest grain, a gate weave	`imax_70mm`	FILM LOOKS / Film stock
A subtle film look, safe on any footage	`cinematic`	FILM LOOKS / Look

Footage or task	Preset	Palette / group
Hard contrast and little colour (Saving Private Ryan, Se7en)	`bleach_bypass`	FILM LOOKS / Look
Cinemascope: horizontal blue flares, teal shadows	`anamorphic_scope`	FILM LOOKS / Look
A VHS tape played on a CRT TV	`vhs_crt`	RETRO / Video
A home camcorder of the 90s	`camcorder_90s`	RETRO / Video
A Polaroid instant photo of the 70s and 80s	`polaroid`	RETRO / Photo

Chains

One run takes one preset. A chain is several runs one after another: the result of a step lands in `<Target>\<Palette>\<preset>_<encoder>\`, and that folder is the Source of the next step. Write the steps in between with `prores_lt`: a lossy file between steps passes its losses on to every next step. The final encoder goes on the last step only. A step run often is worth saving as a profile (section «PROFILES»).

Scenario	Steps
Night timelapse to a film look	`shadow_denoise_sota` → `filmic_rebuild` → `cinematic`
VHS archive to a restored master	`qtgmc_deinterlace` → `derainbow_decross` → `mvtools_mcdegrain` → `archive_clean`
Old web rip to a file fit for a project	`deblock_h264_artefacts` → `dehaze_local_contrast` → `pregrade_prep`
Telecined NTSC DVD to a 24p master	`tivtc_ivtc` → `archive_clean`
Modern digital footage to a film look	`mild_denoise` → `film_35mm`
SD or HD archive to a master twice the size, with a fine grain	`vsmlrt_realesrgan_2x` (without MLRT — `nnedi3_upscale_2x`) → `deband_fine_grain`
Soft HD to a 4K film look	`soft_hd_rebuild_2x` → `film_35mm`

Grain and film looks go after the upscale: grain laid before it is enlarged together with the picture.

PRECISION

Technical preparation: denoise, deband, grain. The groups follow the order of the work: Denoise, then Deband, then Master, which does both and lays grain. The grain of every preset starts from a fixed seed: one file gives the same frames on every run.

Denoise

Mild Denoise (`mild_denoise`)

DFTTest, a spectral denoiser, on the CPU: the same result on any machine, no video card or OpenCL needed. The first pick when little is known about the footage: light, even noise, a small cleanup at low risk.

Field	Values	What it changes
`strength` Strength	`light`, `medium`, `strong`; default `medium`	DFTTest sigma 4, 8 or 14.
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises the sigma with the noise measured in the source and never takes it below the setting.

Encoder: `h264_crf14` for the archive, `h264_crf17` for a final file.

Shadow Denoise SOTA (`shadow_denoise_sota`)

BM3D on 32-bit float cleans the whole frame; a smooth luma mask then keeps the cleaned luma in the shadows only, and the midtones and highlights stay as they were. The colour planes are cleaned everywhere: colour noise is ugly at any brightness. A micro-grain goes back over the cleaned shadows so that they do not look plastic. For night and low-light digital footage whose noise sits in the dark. BM3D runs on the video card or on the CPU, as the PROCESSING block says (section «BM3D on the GPU»).

Field	Values	What it changes
`sigma` BM3D sigma	`1.0` … `4.0`; default `2.5`	BM3D strength on the luma; the colour gets 0.6 of it.
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises the sigma with the measured noise, the luma and the colour each by their own.
`grain_back` Grain back	`0.0` … `2.0`; default `0.6`	Micro-grain over the cleaned shadows; `0` — none.

Field	Values	What it changes
`shadow_threshold` Shadow threshold	`0.0` … `0.5`; default `0.2`	The brightness (0 to 1) where the shadows end: higher — more of the picture is cleaned.
`transition` Mask transition	`0.01` … `0.4`; default `0.1`	How wide the fade from cleaned to untouched is: full cleaning below the threshold minus the transition, none from the threshold up.
`preview_mask` Show mask	on, off; default off	The result is the shadow mask itself, white where it cleans: for setting the threshold.

Encoder: `h264_crf17` or `prores_lt`.

Chroma Cleanup (`chroma_cleanup`)

DFTTest on the two colour planes only. The luma is not touched, so the detail stays as it was. The colour takes a heavier sigma than the luma without a visible loss, as it is stored at a lower resolution. For colour blotches and a dirty blue channel: old compact cameras, low-bitrate AVCHD.

Field	Values	What it changes
`strength` Strength	`light`, `medium`, `strong`; default `medium`	Sigma on the colour: 6, 12 or 20.
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises the sigma with the measured colour noise, the larger of U and V.

Encoder: `h264_crf17`.

Deband

Deband Safe (`deband_safe`)

neo_f3kdb with light settings on 16 bit, no grain added: the steps in skies, walls and night light go, the edges stay sharp.

Field	Values	What it changes
`range` Deband range	`8` … `32`; default `12`	How far neo_f3kdb looks for the neighbours it smooths with: larger — stronger, wider bands go.
`y` Y threshold	`0` … `128`; default `48`	How strong a step of brightness is still smoothed.

Field	Values	What it changes
`cb`, `cr` Cb and Cr thresholds	`0` … `128`; default `32`	The same for the two colour planes.

Encoder: `h264_crf17` or `h265_crf17`. Both write 10 bit, so the smoothed gradients do not band again.

Deband Fine Grain (`deband_fine_grain`)

The same neo_f3kdb, then AddGrain lays a fine grain over the smoothed gradients. The picture keeps a natural texture instead of the too-clean look of cheap compression.

Field	Values	What it changes
`range` Deband range	`8` … `32`; default `14`	How far neo_f3kdb looks for the neighbours it smooths with: larger — stronger.
`y` Y threshold	`0` … `128`; default `48`	How strong a step of brightness is still smoothed.
`cb`, `cr` Cb and Cr thresholds	`0` … `128`; default `32`	The same for the two colour planes.
`grain_var` Grain variance	`0.0` … `2.0`; default `0.6`	AddGrain variance on the luma; the colour gets 0.27 of it.

Encoder: `prores_lt` for grading.

Master

Filmic Rebuild (`filmic_rebuild`)

Three stages in one run: BM3D takes the digital noise out, neo_f3kdb the banding, then grain is laid by brightness — heaviest in the shadows, as on real film, lighter in the midtones, lightest in the highlights. The zones fade into each other over 0.1 of brightness. One preset for filmic footage. BM3D runs where the PROCESSING block says.

Field	Values	What it changes
`sigma` BM3D sigma	`1.0` … `3.5`; default `2.0`	BM3D strength on the luma; the colour gets 0.6 of it.

Field	Values	What it changes
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises the sigma with the measured noise.
`deband_range` Deband range	`8` … `32`; default `14`	How far neo_f3kdb looks for the neighbours it smooths with.
`grain_shadow` Grain in shadows	`0.0` … `3.0`; default `1.5`	Grain variance where the picture is darker than the shadow threshold.
`grain_mid` Grain in mid-tones	`0.0` … `3.0`; default `0.9`	Grain variance between the two thresholds.
`grain_high` Grain in highlights	`0.0` … `3.0`; default `0.4`	Grain variance where the picture is brighter than the highlight threshold.
`shadow_threshold` Shadow threshold	`0.0` … `0.5`; default `0.3`	The brightness (0 to 1) where the shadows end.
`high_threshold` Highlight threshold	`0.5` … `1.0`; default `0.65`	The brightness (0 to 1) where the highlights begin.

Encoder: `prores_lt` or `h265_crf17`.

Archive Clean (`archive_clean`)

A light DFTTest and a light neo_f3kdb, no grain: noise and banding go, no texture is added. A neutral flat master for long storage, delivery or a clean intermediate. Grading works better on footage with grain: when grading is planned, take `pregrade_prep` or `filmic_rebuild`.

Field	Values	What it changes
`denoise` Denoise	`light`, `medium`, `strong`; default `light`	DFTTest sigma 4, 8 or 14.
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises the sigma with the measured noise.
`range` Deband range	`8` … `32`; default `10`	How far neo_f3kdb looks for the neighbours it smooths with.

Encoder: `prores_422hq` or `prores_422hq_mxf`. The built-in profile `archival_master` writes `prores_422hq`.

Pregrade Prep (`pregrade_prep`)

A very light DFTTest and nothing else: no deband, no grain. The footage goes to Resolve or Premiere almost as it was, only without the edge of the shot noise, so the colourist does not grade on top of it. Deband is left out on purpose: it can fight the colour tools; deband after the grade if it is needed.

Field	Values	What it changes
`strength` Strength	`light`, `medium`; default `light`	DFTTest sigma 3 or 6.
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises the sigma with the measured noise.

Encoder: `prores_422hq` (the built-in profile `resolve_handoff`) or `prores_lt_mxf` for Premiere and Avid.

FILM LOOKS

The looks of film stocks. Every look is built of the same parts: MTF softening, a halation of the highlights, grain, a gamma curve, lifted blacks, desaturation. `intensity` scales the whole look; it stands on every page of the palette.

Field	Values	What it changes
`intensity` Intensity	`0.0` … `2.0`; default `1.0`	The strength of the whole look: `0` — off, `1` — as tuned, `2` — heavy.

Film stock

Film 35mm (`film_35mm`)

Kodak 35 mm with the stock to choose. The stocks differ in grain, contrast, halation and colour cast: 250D — daylight, neutral, fine grain; 500T — tungsten, faster, a moderate grain, a cooler tilt; 50D — daylight, the cleanest grain.

Field	Values	What it changes
`stock` Stock	`250D`, `500T`, `50D`; default `250D`	The stock, as above.
`intensity` Intensity	`0.0` … `2.0`; default `1.0`	The strength of the whole look.

Encoder: `prores_lt` or `prores_lt_mxf`. The built-in profile `filmic_warm_35mm` is 250D at `intensity` 1.0 into `h264_crf14`.

Film 16mm (`film_16mm`)

The organic 16 mm feel: softer than 35 mm, as a smaller frame resolves less; more visible grain; a stronger halation; lifted blacks. For a documentary or indie look on digital footage. One field: `intensity`.

Encoder: `prores_lt`.

Super 8 (`super8`)

The most extreme of the stocks: the heaviest grain, the softest optics, the strongest halation, lifted blacks, a lowered contrast — a home movie of the 1970s. For music videos and stylised inserts. One field: `intensity`.

Encoder: `h264_crf17`: the grain is part of the look, and a lossy final file keeps it.

IMAX 70mm (`imax_70mm`)

Large format 65/70 mm, the opposite of Super 8: the resolution is kept, there is almost no grain (on a large frame each grain is tiny against the picture), a gentle halation touches the brightest highlights only, and a barely visible gate weave takes away the too-clean digital look. For shots that should feel epic rather than filmy.

Field	Values	What it changes
`intensity` Intensity	`0.0` … `2.0`; default `1.0`	The strength of the whole look.
`gate_weave` Gate weave	on, off; default on	Every frame moves a little, as film does in the gate of a projector: by 0 or 2 pixels each way, as the colour of 4:2:0 needs even steps. One file gets the same weave on every run. Off, or `intensity` 0 — the picture stands still.

Encoder: `prores_lt` or `h265_crf17`.

Look

Cinematic (`cinematic`)

A universal subtle film look, not tied to a stock: a mild MTF softening, a light bloom of the highlights, a fine grain. No heavy grain, no gate weave, no flicker. A safe first look for any modern digital footage. One field: `intensity`, which scales the grain, the halation and the softening together.

Encoder: `prores_lt` or `h264_crf17`.

Bleach Bypass (`bleach_bypass`)

The look of a print whose silver was not bleached out: a black-and-white image lies over the colour one. High contrast, silver greys, hot highlights, little colour. A strong stylistic choice. One field: `intensity` — higher gives more contrast, less colour and a little more softness.

Encoder: `prores_lt`.

Anamorphic Scope (`anamorphic_scope`)

The cinemascope look of anamorphic lenses: long horizontal blue streaks from the bright highlights only, and cool teal shadows. The work is done in 16-bit RGB, so the streaks do not clip: a highlight mask, a horizontal blur of the flare length, a blue tint, added over the picture — a streak only adds light. Then a teal shift in the shadows (brightness below 0.4) and a light grain. The frame is not cropped to 2.39:1: a true scope frame is cropped in the edit.

Field	Values	What it changes
`intensity` Intensity	`0.0` … `2.0`; default `1.0`	The strength of the whole look, the streaks and the teal included.
`flare_len` Flare length	`32` … `256`; default `96`	The length of a streak in pixels: about 48 — subtle, 160 — dramatic, 256 — extreme.
`flare_thr` Flare threshold	`0.0` … `1.0`; default `0.78`	From what brightness a highlight flares: 0.85 — only the brightest, 0.70 — much of the picture.
`teal_shadow` Teal shadows	`0.0` … `1.0`; default `0.35`	The strength of the teal tint in the shadows: `0` — none, 0.6 — heavy.

Encoder: `prores_lt` for the colourist, `h265_crf14` for the final file.

RETRO

The analog character of video and photo. Every preset has `intensity`, as in FILM LOOKS: `0.0` … `2.0`, default `1.0`.

Video

VHS CRT (`vhs_crt`)

Composite video, a VHS tape, a CRT TV: a heavy sideways bleed of the colour (the signature of the look), soft optics, a slight misalignment of the colour, mild noise. One field: `intensity`. The built-in

profile `vhs_dreamy` runs it at 1.2.

Encoder: `h264_crf17`.

Camcorder 90s (`camcorder_90s`)

A consumer camcorder of the 90s: softer than VHS, slightly overexposed, a mild colour bleed and tape noise — less a broken signal, more a home video. One field: `intensity`.

Encoder: `h264_crf17`.

Photo

Polaroid (`polaroid`)

A Polaroid SX-70 instant photo of the 70s and 80s: soft optics, a candy bloom of the highlights (even bright skin glows a little), blacks that never reach black, warm cream midtones, a little less colour, dark corners. Softer and less grainy than Super 8.

Field	Values	What it changes
`intensity` Intensity	`0.0` … `2.0`; default `1.0`	The strength of the whole look, the vignette included.
`vignette` Vignette	`0.0` … `1.0`; default `0.55`	How dark the corners go: `0` — no vignette, `1` — the corners go black.

Encoder: `h264_crf17`.

RESTORATION

What editors do coarsely or not at all: deinterlacing, inverse telecine, denoise along the motion, compression artefacts, scaling, new frames. Four presets of the palette run on neural networks (MLRT).

Interlace

QTGMC Deinterlace (`qtgmc_deinterlace`)

QTGMC of havsfunc: every field is enlarged by NNEDI3 and smoothed in time along the motion MVTools finds. The result is progressive frames without combing, flicker or ghosting on fast pans. For DV, HDV, captured VHS, S-VHS and U-matic, broadcast TS. The preset does not denoise: denoise as the next step.

Field	Values	What it changes
`field_order` Field order	`tff`, `bff`; default `tff`	TFF — top field first (broadcast, HDV, DV NTSC); BFF — bottom field first (DV PAL). The wrong order makes motion jump back and forth.
`preset` QTGMC preset	11 steps from `Draft` to `Placebo`; default `Medium`	Speed against quality: Draft is the fastest, Placebo the slowest.
`output_fps` Output frame rate	`single`, `double`; default `single`	Single — the frame rate of the source. Double — every field becomes a frame: twice the frames (25i into 50p), the smoothest motion, twice the encoding.

Encoder: `prores_lt` for further work, `h264_crf17` for a final master.

TIVTC IVTC (`tivtc_ivtc`)

Inverse telecine: TFM matches the film fields back into frames, TDecimate drops the repeated frame of every cycle. NTSC 29.97 fps with 3:2 pulldown goes back to the progressive 23.976 fps film master. For NTSC DVDs, broadcast prints, digitised film. The frame count drops by `rdrop` frames of every `cycle`: by a fifth at the defaults; the length of the clip stays. A source below 29 fps is refused: it runs at film rate already, and decimation would drop every fifth real frame.

Field	Values	What it changes
`pp` TFM post-processing	`0` … `7`; default `6`	`0` — off … `7` — deinterlace whatever does not match; `6` deinterlaces the frames TFM could not match cleanly.
`cycle` Cycle	`2` … `25`; default `5`	The TDecimate cycle in frames; 5 for a standard NTSC telecine.
`rdrop` Drop per cycle	`1` … `5`; default `1`	Frames dropped from each cycle; 1 of 5 for NTSC.

Encoder: `prores_422` or `prores_422hq_mxf`.

Noise

MVTools MCDegrain (`mvtools_mcdegrain`)

MVTools follows every block across the frames before and after, and MDegrain averages each pixel with its neighbours along the motion. The noise goes; still texture stays sharp, and so does moving

texture, as it is averaged along its motion, not across it. For high-ISO and low-light footage where skin and fabric must keep their texture.

Field	Values	What it changes
`radius` Radius	`1` … `3`; default `2`	How many frames before and after are taken: 2 — a window of 5 frames; higher — stronger and slower.
`thsad` THSAD	`100` … `400`; default `200`	The threshold of a block's difference: lower — only confident matches are averaged, higher — more averaging.
`noise` Noise	`auto`, `fixed`; default `auto`	Auto raises `thsad` and the scene-change threshold with the measured luma noise.
`blksize` Block size	`8`, `16`; default `16`	The block of the motion search: 16 for SD and HD, faster; 8 for fine detail and 4K, slower.

Encoder: `prores_lt` or `h265_crf17`.

DPIR Denoise (`dpir_denoise`)

A neural preset: it needs MLRT. DPIR, a deep denoise network, keeps the texture of skin and fabric on very dirty footage where BM3D and DFTTest smear it: ISO 6400 and higher, a phone in the dark. The network is told how heavy the noise is (`strength`). The frame size and the frame count stay.

Field	Values	What it changes
`strength` Strength	`1.0` … `50.0`; default `10.0`	The noise level the network expects: 5 — light shot noise, 10 — medium (ISO 1600–3200), 15 — heavy (ISO 6400 and higher), 25 — very heavy (ISO 12800, a phone in the dark), 50 — extreme.
`model` Model	`drunet_color` Colour (default), `drunet_gray` Grey, `drunet_deblocking_color` Blocks, colour, `drunet_deblocking_gray` Blocks, grey	Colour — colour video; Blocks — the blocks of JPEG and MPEG as well; Grey — the brightness alone, the colour kept as it is (black-and-white footage).
`tile` Tile	`256` … `512`, step 16; default `384`	The piece of the frame the network takes at a time, in pixels: smaller — less video memory, slower.

Encoder: `prores_lt` for the colourist, `h264_crf14` for the final file.

Artefacts

Deblock H.264 Artefacts (`deblock_h264_artefacts`)

Deblock_QED of havsfunc, for the artefacts of block compression: the edges of 8×8 blocks, ringing around edges, mosquito noise. It finds the real edges first and smooths the borders of the blocks only, so the true detail stays. For old web rips, messenger re-encodes (WhatsApp, Telegram), low-bitrate SD broadcast.

Field	Values	What it changes
`quant1` Quant 1	`16` … `32`; default `24`	The strength on the edges of the blocks: higher — stronger.
`quant2` Quant 2	`16` … `32`; default `26`	The strength inside the blocks.
`aoffset` Alpha offset	`0` … `2`; default `1`	How sensitive the filter is to edges.
`boffset` Beta offset	`0` … `2`; default `1`	How sensitive the filter is inside the blocks.

Encoder: `h264_crf17`.

Derainbow Decross (`derainbow_decross`)

The rainbow, dot crawl and cross colour of NTSC composite are a fine colour flicker that moves from frame to frame. The motion is found on the luma, and the colour planes are averaged along it over 3 frames — the frame, the one before and the one after. The moving artefacts cancel out, the real colour stays. The luma is not touched, so the sharpness stays. For digitised VHS, U-matic, BetaSP, broadcast masters, S-Video into SDI.

Field	Values	What it changes
`strength` Strength	`0.0` … `1.0`; default `0.6`	How strongly the colour is smoothed: `0` — as it was, `1` — fully.
`blksize` Block size	`8`, `16`; default `16`	The block of the motion search: 16 for SD and HD, 8 for fine detail and 4K.

Encoder: `prores_lt` or `h264_crf17`.

Dehaze Local Contrast (`dehaze_local_contrast`)

Local contrast on the luma alone: the picture minus its blur is added back, weighted by a brightness mask that peaks in the midtones and fades out at black and white. The highlights do not burn, the shadows do not go muddy. The colour is not touched, so it does not shift; the math is 16-bit, without halos. For haze and a flat low-contrast picture.

Field	Values	What it changes
`strength` Strength	`0.0` … `2.0`; default `1.0`	How much local contrast is added back: `0` — the picture as it was, `2` — aggressive.
`radius` Radius	`1` … `16`; default `8`	The size of the blur in pixels: about 8 — clarity of fine structure, 12–16 — haze over large structures.

Encoder: `prores_lt` for the colourist, `h264_crf17` for the final file.

Scale

NNEDI3 Upscale 2X (`nnedi3_upscale_2x`)

NNEDI3 (znedi3, through `nnedi3_resample`) doubles the width and the height: clean edges, the same result on every run, no network to invent texture, no MLRT needed. For SD and HD archive that must be enlarged predictably. 1920×1080 becomes 3840×2160.

Field	Values	What it changes
`quality` Quality	`fast`, `balanced`, `best`; default `balanced`	Fast — a lighter NNEDI3; balanced — as tuned; best — slower, cleaner edges.
`chroma` Colour upscale	`spline36`, `nnedi3`; default `spline36`	Spline36 is the safer one; NNEDI3 is sharper and slower.

Encoder: `prores_lt` for handoff, `h264_crf14` or `h265_crf14` for the final file.

Real-ESRGAN 2x (`vsmlrt_realesrgan_2x`)

A neural preset: it needs MLRT. Real-ESRGAN doubles the width and the height (1080p into 3840×2160) with cleaner, sharper lines than Lanczos or Spline36. Its models in vs-mlrt are made for drawn animation: on live action they can over-sharpen and smooth the texture away — for live action `nnedi3_upscale_2x` (no MLRT needed) is the safer upscale.

Field	Values	What it changes
`model` Model	`animevideov3` Anime (default), `animejanaiV2-L1` … `animejanaiV2-L3` AnimeJaNai 1–3	Anime is the lightest, AnimeJaNai 1–3 go from lighter to stronger; every one is made for drawn animation.
`tile` Tile	`256` … `768`, step 16; default `384`	The piece of the frame the network takes at a time, in pixels: smaller — less video memory, slower.
`tile_pad` Tile overlap	`4` … `64`; default `16`	How far the pieces overlap, in pixels, so that their seams do not show.

vs-mlrt carries no general model of Real-ESRGAN (`general-x4v3`, `general-wdn-x4v3`): the preset does not offer them, and a model the installed vs-mlrt lacks is refused by name, never run as another.

Encoder: `prores_lt` for further work, `h265_crf14` for the final file.

Soft HD Rebuild 2x (`soft_hd_rebuild_2x`)

A neural preset: it needs MLRT. For footage that is HD by size but holds far less detail: soft optics, pixel binning, a strong anti-aliasing filter, old codecs, generations of archive copies. A light DFTTest cleanup comes first, so that the network does not enlarge the noise and the blocks; then Real-ESRGAN at twice the width and height; then a careful sharpening of the luma.

Field	Values	What it changes
`rebuild` Rebuild	`conservative` Careful, `balanced` Balanced (default), `aggressive` Strong	How much detail comes back after the network: Careful — no halos, Strong — for very soft sources.
`cleanup` Cleanup	`light`, `medium`, `strong`; default `medium`	The DFTTest cleanup before the network: sigma 2, 4 or 7.
`model`, `tile`, `tile_pad`	as in Real-ESRGAN 2x	The model, the tile and its overlap.

Encoder: `prores_lt` for grading, `h265_crf14` for the final file.

PROFILES

A profile is a preset with its settings and its encoder, saved under a name in `config\profiles`. It runs on a file or on a whole folder and keeps a step of a chain ready.

Run a profile (`profile_run`)

Field	Values	What it changes
`profile` Profile	the profiles of `config\profiles`	The profile to run; choosing one sets its encoder on the cards.

The PROCESSING block stands on the page and moves whatever the preset of the profile has. The result goes into `<Target>\Profiles\<profile>_<encoder>\`.

Save a profile (`profile_save`)

The button «Save a profile» on the page of a preset saves the settings of that page and the encoder of the cards.

Field	Values	What it changes
`profile_name` Name	Latin letters, digits, `-` and `_`; required	The name of the profile.
`profile_description` Description	a line	What the profile is for.

The built-in profiles:

Profile	Preset	Settings	Encoder
`shadow_clean_quick`	`shadow_denoise_sota`	`sigma` 2.5, `grain_back` 0.6, `shadow_threshold` 0.20, `transition` 0.10	`h264_crf14`
`filmic_warm_35mm`	`film_35mm`	`stock` 250D, `intensity` 1.0	`h264_crf14`
`archival_master`	`archive_clean`	`denoise` light, `range` 10	`prores_422hq`
`resolve_handoff`	`pregrade_prep`	`strength` light	`prores_422hq`
`vhs_dreamy`	`vhs_crt`	`intensity` 1.2	`h264_crf14`

Denoise strength by measured noise

Seven denoise presets — `mild_denoise`, `chroma_cleanup`, `archive_clean`, `pregrade_prep`, `shadow_denoise_sota`, `filmic_rebuild`, `mvtools_mcdegrain` — measure the noise of the source before they start (`noise=auto`, the default):

- 5 frames spread over the clip, flat areas only: texture, edges and areas within 3 % of black or white are left out;
- the result is a standard deviation in 8-bit codes, the luma and the colour each their own.

Their settings (`strength`, `sigma`, `thsad`) are made for noise of about 2 codes. Measured noise above that raises them: BM3D's sigma and MVTools' thresholds grow with the noise, DFTTest's sigma with its square. A setting is never lowered, so clean footage gets exactly what the setting says. MVTools' scene-change threshold rises with the noise too: on heavy noise every block would exceed the default one, the whole frame would count as a new scene and be left unfiltered.

`noise=fixed` (`--set noise=fixed`) takes the settings as they are. What was measured and what the preset made of it go into the report of the run (the field `noise` of each file), for example:

```
[NOISE] mild_denoise: noise 5.34 / 3.44 / 2.75 (8-bit, Y/U/V) -> DFTTest sigma 57.1
```

BM3D on the GPU

`shadow_denoise_sota` and `filmic_rebuild` are mostly BM3D. `bm3dvk` runs BM3D on a Vulkan device of the VapourSynth core — NVIDIA, AMD or Intel, with the ordinary graphics driver; no CUDA Toolkit is needed. Without a Vulkan device BM3D runs on the CPU (`bm3d`).

The PROCESSING block (`vs_accel`) of these two presets:

Value	Where BM3D runs
`auto` (default)	On the video card when VapourSynth sees a working Vulkan device, else on the CPU.
`cpu`	On the CPU: `bm3d` (or `bm3dcpu`, where it loads).
`gpu`	On the video card through Vulkan (`bm3dvk`), any maker; the CPU when there is no Vulkan device.

- On the video card two runs of one file agree within about 65 dB PSNR: the card sums in a varying order. `bm3d` on the CPU gives the same frames byte for byte on every run.
- The other filters run on the CPU, and the video card does not help them; so PROCESSING is shown only where it moves something: the two BM3D presets, the neural presets and profiles.
- SERVICE → «BM3D: video card or CPU» (`bench_gpu`) measures frames a second of each on the first video of the Source; «Doctor» makes a live call and names the device. For a bare `vspipe` run of a `.vpy`, `AUDION_VS_BM3D=vulkan` or `AUDION_VS_BM3D=bm3d` names one of them.

The neural presets and MLRT

Three presets run on vs-mlrt, neural networks inside VapourSynth: `dpir_denoise`, `vsmlrt_realesrgan_2x`, `soft_hd_rebuild_2x`. RIFE Interpolation is out: on an RTX 5070 TensorRT builds wrong frames between the frames, DirectML is open only through ONNX Runtime of the old plugin interface, which VapourSynth R80+ does not load, and on the processor a new 4K frame takes 1.28 s — about half an hour for a minute of 24p doubled to 48p (Decisions).

- **Installation.** The program looks for an NVIDIA RTX card with `nvidia-smi`. With one, it offers MLRT among the additional components at the start and in SETUP (`install_mlrt`): TensorRT, ONNX Runtime and the models, about 3.5 GB to download, trimmed to this card. MLRT goes into the folder of the engine; nothing is written into Windows. Until it is installed, the page of a neural preset says MLRT IS NOT INSTALLED.
- **Without an RTX card and without MLRT** the window does not show these presets, and the console refuses them before the start with `[PREFLIGHT] FAIL MLRT` (exit code 2). MLRT installed on a machine with RTX keeps working when the folder moves to another machine: there the networks run on DirectML or on the CPU.
- **Where the network runs** — the PROCESSING block (`vs_accel`):

Value	Where the network runs
`auto` (default)	TensorRT on an NVIDIA RTX card, else DirectML on the video card, else the CPU. Without an RTX card TensorRT is not tried at all.
`cpu`	ONNX Runtime on the CPU: it works everywhere and is many times slower.
`gpu`	TensorRT on NVIDIA RTX, DirectML on another card.

- **The first run of a model** builds its TensorRT engine: 30–120 s. The engine is kept on disk, and later runs start at once. A new tile size or a new frame size builds a new engine.

- **The tile** is the piece of the frame the network takes at a time. A smaller tile needs less video memory and runs slower: when the card runs out of memory, lower `tile`. `tile_pad` overlaps the pieces so that their seams do not show.
- **Size and frames.** `vsmlrt_realesrgan_2x` and `soft_hd_rebuild_2x` double the width and the height. `dpir_denoise` keeps both.

```
vs_cli.cmd run --operation vsmlrt_realesrgan_2x --set model=animejanaiV2-L2 --set vs_accel=gpu --encoder prores_lt
```

Choosing an encoder

The quality ladder is the same for the software encoders, NVENC, QuickSync and AMF: **14 → 17 → 21**, three steps that are clearly apart. The window's default is x.264 on the CPU at 14 in MP4, that is `h264_crf14`.

Encoder	When
`h264_crf14` / `h265_crf14`	Near lossless: the archive, a master for further work.
`h264_crf17` / `h265_crf17`	Losses almost invisible: a final master.
`h264_crf21` / `h265_crf21`	Web, proxy, preview.
`prores_lt`	ProRes LT, 10-bit 4:2:2 in MOV: grading and round trips, no keying. The ProRes default of the window.
`prores_lt_mxf`	ProRes LT in MXF: round trips with Premiere and Avid.
`prores_422` / `prores_422_mxf`	A little more bitrate than LT; MXF for Premiere and Avid.
`prores_422hq` / `prores_422hq_mxf`	Keying and repeated passes; HQ in MXF for broadcast and Avid finishing.
`dnxhr_lb` / `_sq` / `_hq` / `_hqx`	DNxHR: LB — low bandwidth, SQ — standard, HQ — high quality, all 8-bit 4:2:2; HQX — 10 bit.
`h264_nvenc_q14` / `_q17` / `_q21`	NVIDIA hardware H.264, 8-bit 4:2:0. CQ stands in for CRF. It takes the encoding off the CPU, so the filters of VapourSynth do not wait for it.

Encoder	When
`h265_nvenc_q14` / `_q17` / `_q21`	NVIDIA hardware HEVC, 10 bit. Filters on the CPU or the video card and the encoding on NVENC: the CPU does not stall.
`h264_qsv_q14` / `_q17` / `_q21`	Intel QuickSync H.264: batch and proxy work on Intel graphics.
`h265_qsv_q14` / `_q17` / `_q21`	Intel QuickSync HEVC, 10 bit where the graphics supports it.
`h264_amf_q14` / `_q17` / `_q21`	AMD AMF H.264 through CQP.
`h265_amf_q14` / `_q17` / `_q21`	AMD AMF HEVC, 10 bit.

x.264 and x.265 on the CPU write 10-bit 4:2:0, in MP4 or MKV. A hardware encoder this machine has not passed (SERVICE → «Hardware capabilities») is off in the window.

Rule. For any step that is not the last — `prores_lt` or `prores_lt_mxf`, or DNxHR where the receiving program prefers it. For the final file — `h264_crf14` or `h264_crf17` (software, the best quality per bit), or the hardware ladder of the machine (`nvenc`, `qsv`, `amf`) when time matters.

Hardware or software

Scenario	Encoder
A heavy pipeline (`filmic_rebuild`, `mvtools_mcdegrain`) on NVIDIA	**NVENC**: the CPU is freed for the filters, the whole run is 30–50 % shorter.
A final file where the quality per bit matters most	**libx264 / libx265 CRF**: the software encoder is still slightly more efficient at low CRF.
Intel graphics, a laptop batch	**QuickSync**: good throughput at a low CPU load.
AMD Radeon	**AMF**: the same 14 / 17 / 21 ladder, no encoding on the CPU.
A batch of hundreds of files	**NVENC, QuickSync or AMF**: the saved time adds up.
No working hardware encoder	Software only: `h264_crf*`, `h265_crf*`.

NVENC on RTX 40 and RTX 50 cards gives quality comparable to libx264 medium at the same CQ. On older cards (GTX 10, RTX 20) the software CRF is slightly more efficient at the same quality, and NVENC is still many times faster.

Audio

Only the video goes through the preset; the audio is handled by itself:

- the audio of the source is copied as it is;
- audio that MP4 does not take as it is — PCM, for example, of a ProRes source going into x.264 or x.265 — is encoded to AAC at 384 kbit/s;
- into MXF the sound goes as PCM 24 bit at 48 kHz: its muxer writes no other.

The output file carries the colour range of the source (`-color_range`): a full-range source stays full range.