


CCU user guide

CCU is a **broadcast camera's processing chain with every knob out, for Resolume Arena and Avenue**, as an FFGL effect. It does not paint a "video look" onto a clip. It runs the clip through the stages of a studio or OB camera in their fixed order, in linear light — master gain, white balance, matrix, detail, knee, gamma and black gamma, pedestal, white clip — each stage a known circuit, each knob labelled the way a camera control unit labels it. The badly set-up camera looks fall out of the order: halos the width of the detail delay, grain sharpened by a low coring, faces softened while the jacket stays sharp, highlights gone milky through the knee, a warm white that clips in one channel first, and a white balance that drifts as the camera warms up. None of it is drawn.



The repo's test card through the plugin at the defaults, rendered by the offline harness rather than captured from Resolume. The halos on the bars are the width of the detail delay, not of the edge; the top of the ramp has gone through the knee; the hair on the skin disc is softer than the weave on the jacket because of the skin window; the floor's fine texture is left alone by the coring.

Before you rely on this: released at **v0.1.0**, and honestly early. The chain is measured rather than asserted, by a harness that drives the real plugin class and reads each claim back out of the picture it made, at two rasters: every stage at its null returns the input within a bound derived from the OETF round trip in float, alpha bitwise, and again after a resize; a step's overshoot is $\text{Detail Level} \times h / 4$ for exactly the delay, whole and fractional, horizontal and vertical; an edge below Coring gets no detail at all and one above it gets the overshoot less the dead zone; above the knee point a ramp's slope is Knee Slope, continuous at the point, per channel; a ramp follows the OETF at three exponents and Black Gamma lifts only below its level; red at +6 dB through white balance and then the knee lands where that order predicts and 34,605 tolerances from the other order; the detail gain inside the skin window is $1 - \text{Skin Detail}$ and outside it is 1. Eight deliberately broken chains each fail their check, and one character changed in the shipped shader is caught. All 23 controls are shown to change the picture. It has **never been loaded into Resolume on macOS** — the one host it has run in there is the fleet's own test host, `oxbow`, for 120 frames. On Windows, a build of this source loads, registers and renders in Resolume Arena 7.27.1 on software rendering (win-lab, Mesa llvmpipe, no GPU): all 29 host controls match the declaration and all 24 that take a value move the picture, 9 of the fleet gate's 9 checks. Software rendering says nothing about a GPU or about speed. Try it on a spare layer before you put it in a show.

This codebase was created with AI assistance, directed and reviewed by a human author.

Installing

Every download carries one effect, **SW CCU**. Drop it into Resolume's effects folder and restart Resolume:

```
macOS    ~/Documents/Resolume Arena/Extra Effects/
Windows  %USERPROFILE%\Documents\Resolume Arena\Extra Effects\
```

Avenue uses the same layout under its own folder name. The effect then appears in the effects browser as **SW CCU**.

The macOS download is a universal build (Apple silicon and Intel), as a `.dmg` or a `.zip`. The Windows download is an x64 installer or a `.zip`. It is not code-signed, so the installer trips SmartScreen once: **More info** → **Run anyway**.

The one idea

The "video" look of a studio or OB camera is not a filter. It is a **processing chain in a fixed order**, and each stage is a circuit a camera engineer would recognise:

linear light → master gain → white balance → matrix → detail → knee → gamma and black gamma
→ pedestal → white clip

CCU puts the stages in that order, in linear light, and labels the knobs as a CCU does. Then the looks fall out:

- **Halos.** Detail (aperture correction) adds a high-passed copy of the picture made from pixel and line delays, so a hard edge overshoots on one side and undershoots on the other **for the width of the delay, not of the edge**. A step of height h overshoots by $\text{Detail Level} \times h / 4$.
- **Coring.** Detail below a threshold is not boosted at all, so a low coring sharpens the grain and a high one plastic-coats the picture, leaving only the big edges with their halos.
- **Skin detail.** Inside a hue window the detail gain is reduced, so faces go soft while the jacket stays sharp — and, set wrong, a red jacket goes soft while the face stays sharp.
- **Knee.** Above the knee point the highlights are compressed instead of clipped, so a sky keeps its colour and a white background goes milky. With the knee off, they clip — and because the knee follows white balance, **a warm white clips in red first**.
- **Black gamma** lifts the shadows on a curve below a quarter of video level and nowhere else.
- **Drifting white balance.** The R and B gains wander slowly on a random walk in mireds, as a warming-up camera did: visible over a minute, never a jump.

The clip stands in for scene light: it is linearised through the inverse of the BT.709 curve and its white is the top of what the chain ever sees. A real camera's knee has a scene two or three stops brighter than white to work on; here **Master Gain and the R and B gains are what push anything above white**, and the knee's whole character on footage comes from them.

Start here

Put SW CCU on a clip or a layer with footage that has bright edges, some highlights and, ideally, a face. At the defaults it is a hot broadcast camera, not a broken one: Detail Level 0.9 at a 2-pixel delay, coring at an edge of 0.0225, the knee on at 0.7 linear with a 0.29 slope (so white lands at 90% video and the top of the picture goes milky), Black Gamma 0.2, the Standard matrix, a 9-mired drift. Bright edges carry a halo of the delay's width; hot spots compress to a milky grey; the fine dark texture is left alone by the coring.

Then:

1. **Show Detail on.** The picture is replaced by the detail signal about mid grey: the edges the chain is boosting, and nothing else. Move **Crispeneing Freq** and watch the lines thicken; move **Coring** up and watch the fine texture vanish from the view. This is how a camera engineer sets detail, and it is the quickest way to understand the first three controls. Turn it off.
2. **Detail Level up to 1.** Every bright edge carries a hard halo. **Crispeneing Freq up** and the halo widens without the edge moving.

3. **Coring to 0.** The grain and every fine texture is sharpened with the edges. **Coring to 1** and the picture is plastic: only the biggest edges still have detail.
4. **Knee On off.** The milky highlights turn to clipped white. **Master Gain up** with the knee off and the highlights clip hard; knee back on and they compress instead.
5. **R Gain up, knee off.** The whites go warm and the highlights clip in red first: each highlight's core goes white inside a warm surround, because red reaches the ceiling before green and blue do. Knee on and the red is compressed rather than clipped, so the highlight stays warm all the way up.
6. **Gamma** to the bottom for a flat, crushed-highlight picture; to the top for a bright, contrasty one. **Black Gamma** up to lift the shadows without touching anything above a quarter level.
7. **Matrix** through its four presets, and **Saturation** about the middle.

Every slider is declared to the host as 0 to 1. The value each position stands for is given with each control below, and every null (0 dB, a gain of 1, a saturation of 1) is exact.

The Exposure group

Master Gain — head-end gain, in **linear light, before white balance**: $-6 + 24 v$ dB, from -6 dB at 0 to +18 dB at 1, exactly **0 dB at 0.25**, the default. A CCU's dB gain is sensor gain, and a gained-up highlight is what the knee exists to catch, so it sits at the front of the chain rather than after gamma as a video gain would. It is the control that pushes a clip's white above 1.0 and hands the knee something to compress.

Slider	gain	what it does at the defaults
0	-6 dB	the picture at half, nothing near the knee
0.25	0 dB	the default
0.5	+6 dB	whites at 2× linear: everything bright is in the knee
0.75	+12 dB	mid-tones in the knee, a milky picture
1	+18 dB	8× linear

Master Black — the pedestal, added on video after gamma: $0.2 (v - 0.25)$, from -0.05 at 0 to +0.15 at 1, exactly **0 at 0.25**, the default. Above 0 the blacks lift to grey; below it the shadows crush.

White Clip — the ceiling on video: $1 + 0.3 (v - 0.5)$, from 0.85 at 0 to 1.15 at 1, exactly **1 at 0.5**, the default. Below 1 the picture is clipped at less than white; above 1 the clip is off the top of the range and nothing is clipped by this stage. The floor at 0 is always applied.

The White group

R Gain, B Gain — the red and blue gains about green, in linear light: **± 6 dB**, 12 ($v - 0.5$) dB, exactly **0 dB at 0.5**. Green is the reference, as on a CCU. Because these come **before the knee**, a warm white (R Gain up) is compressed in red by the knee where a gain after the knee would clip it — the harness measures that order.

Drift — a slow random walk on the R and B gains, in mireds: an Ornstein–Uhlenbeck walk with a 20-second time constant, **$60 \times v$ mireds RMS**. Default **0.15, which is 9 mireds RMS**: visible over a minute, never a jump. A mired moves the R and B gains in opposite directions by about 0.4%. The walk runs on the host's clock, so it drifts at the speed the show plays at, and it is seeded so a slider move never jumps the state. 0 is a camera that has warmed up.

The Matrix group

Matrix — a 3×3 in linear light, applied after white balance. The presets are derivations with generic names, not any manufacturer's:

- **Identity**: nothing.
- **Standard** (the default): SMPTE 170M primaries to BT.709 primaries, derived from the chromaticities — a small red-and-green correction that most cameras carry as their base matrix.
- **High Saturation**: a saturation of 1.35 about BT.709 luma.
- **Film-like**: a saturation of 0.82 with a 4° hue rotation about the grey axis.

The luma of the matrixed picture is what the detail and skin stages work on.

Saturation — **$2 v$** about BT.709 luma, from 0 (monochrome) to 2, exactly **1 at 0.5**, the default. Multiplies whatever the preset did.

The Detail group

This is the aperture corrector, and the controls are the ones on a CCU's detail page.

Detail Level — the gain on the detail signal: **$3 v$** , from 0 (no detail at all) to 3. Default **0.3, which is 0.9**. The detail signal is the ($-\frac{1}{4}, \frac{1}{2}, -\frac{1}{4}$) kernel on the luma at $\pm$ the delay, so **a step of height h overshoots by $\text{Level} \times h / 4$** on one side and undershoots by the same on the other.

Crispening Freq — the delay, in pixels: **$1 + 8 v$** , from 1 pixel at 0 to 9 at 1, a whole number at every eighth of the slider. Default **0.125, which is 2 pixels**. This is the width of the halo: the overshoot lasts exactly the delay, whatever the edge's own width. A fractional delay lerps between two whole-pixel taps. The name is the CCU's; it is at most 16 characters because that is the FFGL limit.

H/V Ratio — horizontal against vertical detail: $H = \min(1, 2v)$, $V = \min(1, 2(1-v))$, so vertical only at 0, horizontal only at 1, **both at full at 0.5**, the default. The vertical detail uses the same delay in lines.

Coring — the edge height below which there is **no detail at all**: $0.25v^2$ in linear light, from 0 to 0.25, squared for resolution at the bottom. Default **0.3, which is an edge of 0.0225**. The dead zone on the detail signal is a quarter of that (a step of h peaks at $h/4$). Below the coring, grain and fine texture are left exactly alone; above it, the overshoot is the full one less the dead zone.

Slider	edge height	what it looks like
0	0	every grain sharpened with the edges
0.3	0.0225	the default: the noise floor left alone, texture kept
0.6	0.09	only strong edges have detail; skin and fabric go smooth
1	0.25	plastic: only the hardest edges carry a halo

Level Dependence — how far the detail is reduced in the shadows: the detail gain is scaled by $1 - v(1 - \min(1, Y/0.2))$, so at 1 there is no detail at black and full detail from a luma of 0.2 up; at 0 the shadows are sharpened like everything else. Default **0.5**. Cameras have this because shadow noise is where detail hurts most.

Skin Detail — the detail gain inside the skin window is $1 - v$. Default **0.4**, so a face gets 60% of the detail the rest of the picture gets. At 1 there is no detail on skin at all; at 0 the window does nothing.

Skin Hue — the centre of the window, **$360v$ degrees** of hue on the linear colour, 0 at red and 60 at yellow. Default **0.0556, which is 20°**: skin. The two ends of the slider are the same hue.

Skin Width — the window's half-width, **$5 + 55v$ degrees**, from 5° to 60°. Default **0.2727, which is 20°**. The window is flat inside three-quarters of its width and smoothstepped to nothing at the edge, and it is gated on chroma — a neutral is never skin, however its hue reads.

The Knee group

Knee On — on by default. Off, the picture above white clips; on, it compresses.

Knee Point — where the knee starts, **$0.4 + 0.6v$ in linear light**, from 0.4 to 1.0. Default **0.5, which is 0.7**. Below the point the chain is straight; above it the slope changes, continuous at the point.

Knee Slope — the slope above the point, **$0.05 + 0.95v$** , from 0.05 to exactly 1 (no knee) at

1. Default **0.25, which is 0.29**: white at 1.0 linear lands at 0.787 linear, about 90% video, and that is why the defaults look milky at the top. The knee is **per channel**, in linear light, so a coloured highlight desaturates as it compresses — a per-channel knee is what most cameras call "knee".

With the point at 0.7 and the slope at 0.29, a linear value L above 0.7 becomes $0.7 + 0.29 (L - 0.7)$. At +6 dB ($L = 2$ for white) that is 1.08, still above the white clip; the picture is a bright grey with the highlights flattened, not a clipped one.

The Gamma group

Gamma — the opto-electronic transfer function's **exponent**, **$0.35 + 0.2 v$** , from 0.35 to 0.55, **BT.709's 0.45 at 0.5**, the default. The curve is a family: the linear segment's slope and the offset follow from the exponent by continuity at 0.018, and at 0.45 they are the standard's 1.099 and 4.5 to its printed precision. Lower is flatter with crushed highlights; higher is brighter and more contrasty. The clip was linearised through the same family at 0.45, so at the default the round trip is the identity.

Black Gamma — a lift below a quarter of video level, and nowhere else: a smooth bump on video of **$0.056 \times v$** at its peak (a third of the way up to the level), zero at and above 0.25 video and with zero slope there. Default **0.2**, a lift of 0.011 at most. Lifts the shadows without touching the mid-tones or the highlights.

The Output group

Mix — the processed picture against the untouched clip, 0 to 1; **1 by default**. Zero is the clip as it arrived. The output carries the clip's own alpha throughout.

Show Detail — off by default. Instead of the picture, draws **the detail signal about mid grey**: what the aperture corrector is adding, after coring, level dependence and the skin window, before it is added. Flat areas are mid grey, edges are bright-and-dark pairs the width of the delay. Use it to set the delay and coring by eye, the way a camera engineer does on a monitor.

How it works

Once a frame, two passes on the GPU and a few constants in double on the CPU:

1. **Linear**. The clip is linearised (the inverse BT.709 curve, as a family at exponent 0.45), gained, white-balanced with the drift, and matrixed, into an RGBA32F buffer at picture size. The luma of the result goes into the buffer's alpha, because detail and skin work on it.
2. **Process**. Each output pixel reads its own linear value and eight neighbours (two per far tap, lerped by hand for a fractional delay), builds the detail signal, cores it, scales it by level and skin, adds it to all three channels, applies the knee, the OETF, black gamma, the pedestal and the white clip, and mixes with the source.

Every tap is a `texelFetch` at an integer coordinate, so nothing rests on the texture unit's filtering. The OETF's constants, the matrix and the drift's step are computed in double once a frame and handed over as floats. Nothing carries across frames except the drift, a double on the CPU stepped by the host's clock.

Performance

Measured by the offline harness on an M4 Max at the defaults, best of three runs of 60 frames after a warm-up, on a GPU shared with other work:

	ms/frame	% of a 60 fps frame
1280×720	0.06	0.3%
1920×1080	0.17	1.0%
3840×2160	0.68–0.70	4.1%

Two passes and eight texel fetches a pixel; it is cheap. GPU memory is one RGBA32F buffer at picture size: 32 MB at 1080p, 127 MB at 4K. Nothing was timed inside Resolume, and nothing was timed on Windows.

If it looks wrong

Everything bright has gone milky grey. That is the knee at the defaults doing what a hot camera does. Raise Knee Point, raise Knee Slope toward 1, or turn Knee On off and let it clip.

The highlights clip hard and go white. Knee On is off, or Master Gain is high with the slope at

1. Turn the knee on.

The grain is sharpened. Coring is low. Raise it until the noise leaves the Show Detail view.

The picture is plastic. Coring is high, or Skin Detail is high with a wide window covering the whole picture's hue. Lower Coring; check Skin Hue and Skin Width.

A red or orange object has gone soft while the face is sharp. Skin Hue is on the wrong hue. Turn Show Detail on and move Skin Hue until the face's texture drops out and the object's comes back.

Warm highlights with white cores. A warm white balance (R Gain up) with the knee off: red reaches the ceiling first and the core of every highlight loses its warmth. Turn the knee on.

The colour slowly wanders. Drift is above 0. That is the point; set it to 0 for a warmed-up camera.

The blacks are grey. Master Black is above a quarter, or Black Gamma is high.

The picture is a grey field with bright-and-dark lines on it. Show Detail is on.

SW CCU is not in the effects browser. Check the folder under Installing, and that Resolume was restarted.

The effect does nothing at all. A shader that will not compile looks exactly like that, and the real message is in the log:

```
macOS    ~/Library/Logs/ccu/ccu.YYYY-MM-DD.log
Windows  %LOCALAPPDATA%\ccu\logs\ccu.YYYY-MM-DD.log
```

It records the build that was loaded, the GL vendor, renderer and version at load, which pass failed to compile if one did, and a buffer that could not be allocated.

Known limits

- **Never loaded into Resolume on macOS**, and nothing has driven the controls in a host there. How the twenty-three controls read in the inspector is untested.
- **The clip stands in for scene light.** An 8-bit clip's white is the top of what the chain sees; only Master Gain and the R and B gains push anything above it, and the knee's character on footage comes from them.
- **The knee is per channel only.** There is no luma knee with chroma preserved.
- **The skin window is on the linear colour's hue**, not on a camera's post-gamma chroma vector.
- **The vertical detail uses the same delay as the horizontal**, in lines, lerp'd for a fraction; a real camera's line delays are whole lines.
- **The drift's mired-to-gain law is a judged constant** (0.4% per mired), not a Planckian locus, and the drift has never been watched over a minute in a host.
- **No audio input, no presets** and no OpenFX version.
- **Only ever run on an Apple M4 Max and on win-lab's software renderer**, although the macOS build contains an Intel slice. On Windows, see the note at the top of this guide.
- **There is a browser demo** at ccu-demo.stoatworks-labs.com. It runs the plugin's own shaders in WebGL2 with the CPU half ported to JavaScript; the page lists what it does not reproduce.
- **Checked at up to 1280×720**, and only timed at 4K.

About

The last group, **About**, carries the plugin's name, version, licence and maker, and buttons that open this user guide (stoatworks-labs.com/software/ccu/guide/), the project page, the source on GitHub

and the support page in your browser.

Reporting something

github.com/stoatworks-labs/ccu/issues. A screenshot, the Detail Level, Crispening Freq, Coring, Knee Point and Knee Slope settings, and the composition's resolution are usually enough. If the effect did nothing, attach the log.

CCU · guide updated 2026-09-24 · stoatworks-labs.com/software/ccu/

Latest version of this guide: stoatworks-labs.com/software/ccu/guide/ · Source and issues: <https://github.com/stoatworks-labs/ccu>