


SAK-BUS

*Stereo compression.
An interacting circuit.*



TECHNICAL REFERENCE MANUAL

Version 0.1.1 · Windows x64 · VST3 & standalone

Controls · signal path · calibration · evidence

01 / Overview

SAK-BUS is an original stereo mastering and mix-bus compressor. A modern trimmed complementary VCA sits between independently designed input and output magnetic stages. Its control circuit draws on the SSL 82E26/82E27 feedback/control topology. It is not an identified model of a particular vintage compressor.

The circuit retains state. A sustained bass note changes magnetic history. A transient charges the timing network. Makeup changes the level reaching the output transformer. These interactions continue when gain reduction is small and while the plugin is bypassed.

Reading this manual

Page	Subject
3	Installation, compatibility and a first session
4–5	Controls, sidechain, bypass and metering
6	Signal flow and calibration points
7	Compression and control architecture
8	Coupled Auto Release network
9	Complementary VCA model
10	Magnetic state and the electrical solve
11	Separate input/output magnetic calibration
12	Amplifiers, passives and approximations
13	Oversampling, latency and state changes
14	Headroom and level calibration
15–16	Measurements, protocols and validation
17	Evidence boundaries and primary references
18	License, source, build and version history

Published means a named source supplies the information. **Derived** means a calculation follows from a stated circuit or equation. **Calibrated / designed** means SAK selected parameters or behavior. **Measured** here means the production software was measured, unless a hardware source is explicitly identified.

A component specification is not proof of complete-unit equivalence. The source release includes the research ledger and raw data so that these distinctions can be checked.

02 / Install and begin

Windows release

The package contains a 64-bit VST3 effect and optional standalone application. Windows 11 x64 is the tested operating system. Setup accepts Windows 10 or later; Windows 10 is not separately validated. No macOS, Linux, AU, AAX or VST2 binary is supplied.

1. Close the audio host. Run **SAK-BUS-0.1.1-Setup.exe**. The installer is unsigned; Windows may show an unknown-publisher warning. The release includes SHA-256 checksums.
2. The default plugin destination is **C:\Program Files\Common Files\VST3\SAK-BUS.vst3**. This shared location requires administrator access. Setup can use a different VST3 directory if your host requires it.
3. Documentation and the optional standalone application go in **C:\Program Files\SAK Audio\SAK-BUS**. Choose the Full installation to include the standalone application.
4. Setup includes Microsoft's signed Visual C++ v14 x64 runtime and installs it if the required version is missing. Restart if that prerequisite requests it.
5. Open the host, rescan VST3 effects and insert SAK-BUS on a stereo track or bus. Main input and output must both be stereo. Float and double host audio are supported.

First session

Start with Ratio **4:1**, Attack **10 ms**, **Auto Release on**, Makeup **0 dB**, Parallel Blend **100%**, HPF **Off** and processing **Master**. Lower Threshold until there is a few decibels of reduction on stronger passages. Adjust Makeup for a useful level comparison. Auto is a suggested starting point, not the factory default: the default is manual 0.3 s.

Use slower attack settings to let more of the initial transient pass. Use the sidechain HPF when low-frequency content dominates detection. Listen for recovery in the space between phrases. Parallel blend combines the compressed path with the delayed dry input; it is not an automatic loudness match.

Portable use and removal

Extract the portable ZIP before use. Copy the **entire** SAK-BUS.vst3 directory, including Contents, to the VST3 folder. Its standalone executable can run directly with the runtime installed. The ZIP does not install the runtime automatically.

Uninstall a Setup installation through Windows Settings → Apps. Shared runtime files, host projects and user presets are retained. Delete a manually copied portable bundle from the host's scan directory to remove it. If a host does not see the plugin, verify the full bundle path, x64 host architecture and rescan/blacklist settings.

03 / Main controls

Control	Range / positions	Factory default
Threshold	−40 to +12 dBFS	−12 dBFS
Ratio	2:1, 4:1, 10:1	4:1
Attack	0.1, 0.3, 1, 3, 10, 30 ms	10 ms
Release	0.1, 0.3, 0.6, 1.2 s	0.3 s
Auto Release	Off / on	Off
Makeup	0 to +24 dB	0 dB

Threshold and ratio

Threshold sets the detector's operating level in dBFS sine-peak calibration. Lower it to increase compression. Ratio selects one of three finite-gain feedback networks. It does not select an exact straight-line dB slope at every level. The nominal local ratio is calibrated around 4 dB of reduction; actual sine curves depend on rectifier conduction, attack resistance, leakage and ripple.

Changing ratio also changes effective onset. The lower ratio begins acting earlier in the idealized calibration. Retune Threshold when comparing ratios rather than assuming equal numerical settings produce equal reduction.

Attack

Attack selects charging resistance. Faster settings reduce the initial transient more strongly; slower settings let more of it through. Front-panel labels are operating positions, not guaranteed 10–90% step-response times. The feedback loop accelerates charging, and its effect depends on signal level and ratio. Measurement definitions are on page 15.

Release and Auto

Manual Release selects discharge resistance. Auto replaces that discharge behavior with the coupled fast/slow capacitor network described on page 8. A short event mostly exercises the fast state; sustained reduction leaves more stored slow-state voltage. The Release dial retains its manual selection while Auto is enabled and returns to it when Auto is disabled.

Makeup

Makeup adds gain after the audio VCA and before the output drivers and transformer. It changes physical output-stage excitation, not just a final digital gain. High makeup can reach analog-model headroom before a digital output peak reaches an expected value. Compare at matched listening level and watch the output meter.

Controls are exposed to host automation and session recall. Continuous controls and switched targets are smoothed. Parameters change the running circuit without erasing its magnetic history.

04 / Utility strip and meters

Control	Positions / range	Factory default
Sidechain HPF	Off, 30, 60, 90, 120, 180 Hz	Off
Parallel Blend	0–100% wet	100%
Processing	Studio 2× / Master 4× / Render 32×	Master
Analog Noise	Off / on	Off
External Sidechain	Off / on	Off
Bypass	Off / on	Off

Sidechain HPF filters the detector feed. It reduces how strongly bass drives gain reduction without applying the same high-pass filter to the audio path. Stereo detection still uses the louder rectified channel.

Parallel Blend mixes the processed signal with latency-aligned dry input. Zero percent is dry; 100% is the circuit output. Delay alignment removes the resampling delay difference. The wet circuit still has its own frequency-dependent phase, so an intermediate blend is not a phase-neutral gain interpolation.

Processing selects the whole-circuit oversampling factor. Start in Master. Studio reduces CPU cost. Render is intended for offline work, especially intentional overload, and can be slower than real time. All modes report 248 samples latency.

Analog Noise enables circuit-referred stochastic noise. Disabling it retains deterministic VCA distortion, offsets, headroom, magnetic hysteresis and detector ripple. Noise seeds are deterministic on reset, independently assigned to the channels, and do not drift as an invented aging effect.

External Sidechain uses a host-routed mono or stereo auxiliary input. Configure the route in the host as well as enabling the switch. If the host has no enabled sidechain bus, the plugin falls back to internal detection. An enabled but silent sidechain bus intentionally supplies silent detection. Sidechain audio is not mixed into the output.

Bypass crossfades to the delayed dry input while the circuit continues to evolve. Host bypass is also supported. The host can still apply its own processing suspension policy; this manual cannot guarantee how every DAW manages inactive plugins.

Read the meters correctly

The large movement shows linked gain reduction with simulated mechanical ballistics and a 0–20 dB display range. It is a smoothed display of control reduction, not an instantaneous waveform analyzer. The numeric readout follows the movement.

The lower input/output meters show digital sample-peak levels in dBFS, with display decay and peak hold. They are not true-peak or loudness meters. A displayed level below 0 dBFS does not establish intersample headroom. Makeup, blend and downstream processing should be assessed with the metering appropriate to your delivery format.

05 / Signal flow

STEREO INPUT

- ↓ input Jiles-Atherton transformers
- ↓ 6 dB pad · finite-gain receivers · coupling capacitors
- ↓ complementary audio VCAs ← linked control voltage
- ↓ DC servos · makeup · balanced output drivers
- ↓ output Jiles-Atherton transformers
- ↓ decimation · aligned parallel blend · bypass

STEREO OUTPUT

RECEIVER FEED or EXTERNAL SIDECHAIN

- ↓ independent sidechain high-pass filters
- ↓ active full-wave rectification · loudest-channel link
- ⇌ tracking sidechain VCA · ratio / threshold network
- ⇌ diode charge path · manual RC or coupled Auto network
- shared control voltage → left and right audio VCAs

The input magnetic stage changes the waveform reaching both the audio gain cell and the internal detector. The timing voltage controls the audio cells and the tracking sidechain cell. A safeguarded same-sample solve closes this feedback-like control relation rather than feeding back a delayed ideal gain value.

Each transformer solves its magnetic current together with its electrical winding network. Magnetic saturation changes effective loading and primary voltage. The output transformer load also affects the driver current budget; that driver feedback uses the previous internal sample as a stated reduction.

Makeup occurs before the output magnetic stage. The parallel dry path is taken outside the interacting circuit. Bypass does not stop state evolution in the plugin engine.

Units and gain structure

Samples map to differential line voltage at **17.37068 V per digital peak unit**. A -20 dBFS sine corresponds to +4 dBu RMS. A precision 2:1 receiver pad reduces cell current; opposing output-driver legs restore approximately 6.02 dB. Each modeled amplifier leg uses ± 18 V rails and approximately ± 16.5 V available swing.

Transformer API voltages include fixed midband insertion-loss compensation. This represents receiver gain at the input. At the output, the small correction is moved mathematically across the lightly loaded transformer. Isolated passive terminal measurements explicitly undo this normalization.

06 / Compression and control

Detection and stereo linking

The detector uses active full-wave amplitude rectification, not generic RMS detection. Two independent rectifiers avoid the cancellation that an L+R audio sum can produce with opposing polarity. The larger rectified channel drives the common control voltage.

Finite rectifier bandwidth is 120 kHz and the input-referred crossover term is 20 μ V. These are conservative effective design values, not identified vintage op-amp measurements. Ripple remains in the control waveform and can modulate the audio VCA at low frequencies.

Tracking VCA and ratio behavior

The same timing voltage attenuates a tracking sidechain VCA. This produces a feedback-like compression law while detection is derived from the receiver feed. The sidechain cell is a reduced exponential control relationship; it does not replicate every audio-cell noise and distortion state.

In the constant rectified-input, negligible-leakage limit, with reduction **g** in dB:

$$\begin{aligned} g &= \text{KdB} \times [10^{(L - T_{\text{effective}} - g)/20} - 1] \\ R_{\text{local}} &= 1 + (\ln(10)/20) \times (\text{KdB} + g) \\ \text{KdB} &= (R_{\text{nominal}} - 1)/(\ln(10)/20) - 4 \end{aligned}$$

The last relation places the nominal differential ratio at 4 dB of reduction. Ratio-dependent bias aligns the idealized curves there. This is SAK calibration, not a resistor-by-resistor solve of the complete 82E26 network. Actual sine slopes also depend on conduction duty, ripple and timing resistance.

Manual timing

The attack bank uses 820 Ω , 2.7 k Ω , 8.2 k Ω , 27 k Ω , 82 k Ω and 270 k Ω feeding 0.47 μ F. Their raw RC products exceed the front-panel attack times because the feedback loop speeds charging.

Manual discharge uses 91 k Ω , 270 k Ω , 560 k Ω or 1.2 M Ω , with 3.3 M Ω across the total timing node. The 91 k Ω shortest position is a documented SAK choice: the historical scan appears to show 180 k Ω despite its 0.1 s label. The discrepancy is retained in the reference ledger.

At long attack and short release, charging/discharge loading changes the equilibrium. Treat measured compression curves and transient definitions as the operating description; do not infer them from a dial label alone.

07 / Auto Release

Auto is a physical two-state timing network: **two series-connected parallel RC branches**, sharing the same charging current. The fast branch is 0.47 μF || 91 k Ω . The slow branch is 6.8 μF || 750 k Ω . A 3.3 M Ω load spans their combined voltage.

$$u = v_f + v_s$$

$$C_f \times dv_f/dt = I_{\text{charge}} - u/R_{\text{sh}} - v_f/R_f$$

$$C_s \times dv_s/dt = I_{\text{charge}} - u/R_{\text{sh}} - v_s/R_s$$

The rectifier diode charges only when its drive exceeds the total timing voltage. Backward-Euler companion equations reduce the capacitors to an affine total-voltage relation; a safeguarded scalar solve closes the exponential tracking-VCA loop in that same internal sample.


Why history changes the release

A short event puts relatively little charge into the slow branch. A sustained passage leaves a greater slow-state contribution. The next release therefore depends on what preceded it, without measuring a program label or switching between two generic envelopes.

The individual branch voltages can briefly reverse during the slow tail while their sum remains nonnegative. Clamping both branch states to positive values would violate the coupled circuit's current balance. The final implementation preserves their signed voltages.

Quantity	Fast	Slow
Unloaded RC product	42.77 ms	5.1 s
Loaded off-state eigenmode	≈ 41.62 ms	≈ 4.1765 s

The often cited 100 ms / 12 s Auto markings are close to 90% decay of the unloaded poles. The manufacturer did not supply a measurement protocol that establishes this interpretation; the correspondence is an inference.



48 kHz / Master. 1 kHz carrier, -6 dBFS program, -24 dBFS threshold, 4:1, 0.1 ms attack; then -60 dBFS. Histories: 80 ms, 2 s and 12 s. Curves come from the saved production measurements.

08 / Complementary VCA

The audio cell uses a Blackmer-style complementary log/antilog reduction based on Gary Hebert's published analysis and modern THAT 2180A terminal data. Finite bias current, symmetry error and residual resistance share a branch-current model; distortion and offset are related consequences.

For input current $i = V_{in} / 20 \text{ k}\Omega$, gain $G = \exp(-CV/b)$ and $I_b = 20 \text{ }\mu\text{A}$:

```
b = 0.0061 / (ln(10)/20)
q = sqrt(i2/4 + Ib2/G)
i_positive = q + i/2
i_negative = q - i/2
j = G × i_branch × symmetry × exp[r × (i_branch - j)/b]
Vout = 20 kΩ × (j_positive - j_negative)
```

Three Newton refinements solve each positive output branch. With perfect symmetry and zero residual resistance, the nonlinear branch terms cancel to ideal exponential gain. Imbalance generates even harmonics and gain-dependent offset within the same relationship.

Calibrated imperfections

Residual resistance **0.395 Ω** represents what remains after internal parasitic cancellation. Symmetry is **$1 \pm 0.0003475/2$** . These are fitted model coefficients, not measurements of a specific die. Unity and attenuation distortion anchors constrain the fit; positive gain provides a consistency check.

The CV scale is **6.1 mV/dB** at fixed reference temperature. There is no warm-up drift. A shared 1.83 mA current budget, with a smooth knee beginning at 1.4 mA, limits the cell. This overload law is an effective approximation rather than an identified transistor-level current constraint.

Output slew follows **$24G/(1+G) \text{ V}/\mu\text{s}$** . A 20 kΩ / 22 pF transimpedance pole supplies modeled output bandwidth. Fixed output bias, gain-scaled input bias and mismatch-derived offset feed the coupling and DC-servo states.

Noise and identification limits

Noise uses the four gain-dependent terms of Hebert equation 24. The shot contribution follows instantaneous total branch current. A 1.233 power factor aligns the nominal calculation with the later packaged-device unity-noise specification. Generation scales with internal sample rate and passes through the circuit and decimator.

The selected device's complete trimmed transistor netlist, process models and parasitics are unavailable. This analytical model preserves observable mechanisms. It is neither a transistor-perfect die simulation nor an identified model of a particular historical dbx 202 module.

09 / Magnetic state

SAK-BUS uses the shared SI Jiles–Atherton solver extracted unchanged from SAK-TUBE. It supplies **new input and output parameters**, not the amplifier's steel-core configuration. The solver's lineage establishes the method; it does not validate the new transformer calibrations.

The core retains reversible and irreversible magnetization, Langevin anhysteretic response, direction gating, molecular-field coupling, effective gap reluctance and eddy-current loss. The prior committed state is held fixed during trial evaluations. Only the converged electrical root commits the next magnetic state.

Magnetics inside the winding network

Copper resistance, source/load impedance, leakage inductance and effective differential capacitance are part of a loaded transformer equivalent. Trapezoidal companions reduce the linear network to an affine load-current relation. The primary root is:

```
Iload(v) = a × v + b
λ_trial = λ_old + Δt/2 × (v_old + v)
F(v) = v + Rsource_total × [Imag(λ_trial) + n × Iload(v)]
      - Vsource = 0
```

Here λ is flux linkage and n is turns ratio. Positive incremental reluctance supplies a monotone electrical root. Bracketed Newton steps prevent an overload trial from simply escaping its solution interval.

Saturation increases magnetizing current and reduces incremental winding impedance. Source resistance and loading turn this into voltage and phase changes. A static waveshaper after an ideal transformer would omit this interaction and its dependence on prior history.

History and scope

Startup phase, low-frequency bursts and DC excitation can leave long decays. Measurements must specify settling and flux history. Reset returns the engine to deterministic initial state; ordinary parameter changes do not reset the core.

The legacy nominal-inductance field does not set Jiles–Atherton incremental inductance. The actual value follows material slope, turns, area, path and gap. Final input/output demagnetized inductances are calculated and verified as approximately 300 H and 40 H.

This is a differential two-channel audio model. Common-mode interference, shielding, hum-field pickup, random aging and temperature drift are not synthesized. Public terminal specifications do not uniquely identify the material loop, geometry or winding construction.

10 / Two magnetic stages

The input and output are independently designed 1:1 high-nickel-equivalent transformers. The input deliberately targets about **6 dB more 20 Hz headroom** than the Jensen JT-11P-1 reference. The output uses JT-11-BMCF-like terminal targets. Neither is claimed to reproduce the construction of a Jensen product.

Parameter	Input	Output
Saturation polarization	0.75 T	0.75 T
Effective turns	4200	2100
Effective core area	53.98 mm ²	131.99 mm ²
Effective magnetic path	0.12 m	0.15 m
Effective series gap	3.104 μm	16.605 μm
Demagnetized inductance	300 H	40 H
Primary / secondary copper	1450 / 1550 Ω	40 / 40 Ω
Production source / load	600 Ω / 10 kΩ	2 Ω / 10 kΩ
Reference source / load	600 Ω / 10 kΩ	0 Ω / 600 Ω
Leakage / differential capacitance	18 mH / 55 pF	20 μH / 90 pF

The gap represents effective reluctance of joints and geometry, not an identified machined air gap. Geometry and material coefficients are underdetermined design parameters. Effective differential capacitance is not the sum of published winding-to-shield capacitances.

Accepted terminal calibration

20 Hz condition	Target	Model result
Input, +4 dBu source	0.005% THD, SAK choice	0.004979%
Input, +26 dBu source	1% THD, SAK choice	0.99780%
Output, ≈+4 dBu into 600 Ω	0.002% THD	0.001996%
Output, +26.993 dBu into 600 Ω	≈1% THD	1.00224%

These isolated tests use 192 kHz, coherent cosine excitation, a two-cycle smooth onset, twelve settling cycles and six measured cycles. Passive secondary voltage is measured with API normalization undone; harmonics through order 32 are included when below Nyquist.

A rejected fit met terminal anchors but produced excessive intrinsic coercivity. The accepted constrained output fit retained the anchors and reduced model H_c at 0.7 T to about 2.88 A/m. Positive loop loss and incremental reluctance were checked. These are model-loop checks, not comparisons with measured Jensen B–H loops.

The dense output overload knee remains steep and lacks a published dense hardware curve for comparison. Changing the load to the production 10 kΩ fixture changes the result. Isolated transformer headroom must not be reported as complete-plugin headroom.

11 / Amplifiers and passives

Receivers and drivers use a reduced finite-amplifier model anchored to published NE5532-family scales. It retains finite DC gain, bandwidth, slew, voltage compliance, load-current limiting and a bounded recovery state.

Effective parameter	Model value / interpretation
Open-loop DC gain	100,000
Gain-bandwidth product	12 MHz
Slew	5 V/ μ s
Supply / available leg swing	± 18 V / approximately ± 16.5 V
Excess-charge recovery	40 μ s; effective approximation
Input coupling	47 μ F / 20 k Ω high-pass
Output DC servo	0.15 Hz closed-loop high-pass equivalent

What is reduced

The dominant pole uses an exact held-input solution. This avoids introducing a spurious Nyquist zero by discretizing an irrelevant megahertz pole. A smooth complementary compliance transition replaces a hard clipping corner; its four-thermal-voltage width is a reduced driver-junction assumption.

The 40 μ s excess-charge recovery and small rail-excess contribution have not been identified from a measured overload trace of this complete design. They are bounded effective mechanisms. At normal audio-band level, several volts per microsecond of slew capability need not create audible slew distortion.

Output drivers consume the transformer's previous internal-sample load current. The winding magnetic KCL root itself is current-sample implicit. This split is an explicit approximation near current limit; nonlinear sample-rate convergence and overload tests help assess its consequences.

The two opposing output legs share a differential equivalent. Common-mode mismatch is omitted. Rails are finite but regulated: dramatic power-supply sag is not added without supporting evidence.

Noise, tolerances and coupling

Analog Noise enables resistor thermal and amplifier voltage-noise contributions alongside VCA noise. The same filters and physical path shape their spectra. Disabling noise does not remove deterministic DC offsets, hysteresis or rail behavior.

Input coupling and the output servo let gain-change offsets produce finite recovery transients. The servo's closed-loop high-pass equivalent retains that interaction without a full transistor-level servo amplifier.

Timing-capacitor ESR is negligible against even the smallest 820 Ω attack resistor in this model. Component tolerances are fixed; matched stereo is a design target. Random tolerances, dielectric absorption, unmeasured voltage coefficients and arbitrary thermal drift are not invented to create motion.

12 / Quality and latency

Mode	Whole-circuit rate	Intended use
Studio	2× host rate	Lower CPU auditioning
Master	4× host rate	Normal-level mixing and mastering; default
Render	32× host rate	Offline rendering and intentional overload

All interacting circuit computation uses double precision at the selected internal rate. This includes transformers, VCA, detector and timing feedback, rather than only isolated saturation blocks. The host wrapper accepts both 32-bit float and 64-bit double buffers.

Linear interpolation/decimation uses 257-tap symmetric Kaiser half-band stages, beta 11, alternate zero taps and symmetric convolution. Every mode compensates to **248 host samples**. The wet circuit retains physical frequency-dependent phase.

Host sample rate	248-sample delay
44.1 kHz	5.624 ms
48 kHz	5.167 ms
96 kHz	2.583 ms
192 kHz	1.292 ms

These delays are arithmetic conversions, not additional measured delays. Dry blend and bypass use the same external alignment. Audio-device and host-buffer latency is separate.

Quality changes and runtime behavior

Changing quality transfers magnetic/timing history to another preallocated engine, warms its resamplers and crossfades. Both engines run briefly during that transition, so temporary CPU cost can increase. Host latency stays fixed.

The audio callback performs no allocation, lock, file access or coefficient generation. Parameters are smoothed; Auto/manual changes preserve total control voltage. Non-finite samples and parameter payloads are sanitized, and tiny recursive states are zeroed to avoid denormal stalls.

Render can exceed real-time capacity. Its greater overload alias suppression does not imply zero aliasing for every signal. The measured residual includes aliases and other nonharmonic state modulation; the exact experiment is described on page 16.

Tail and stopping playback

The host is told to retain **16 seconds** of signal-induced decay. The specified +20 dBu / 20 Hz burst settles below −120 dBFS after about 8.15 s; adding 0.1 digital unit of DC extends this to about 11.22 s. These noise-off results use 10 ms peak windows. Optional stationary analog noise continues while processing is active. Host suspension can affect how a tail is retained.

13 / Headroom and calibration

SAK-BUS maps digital audio to a differential line-voltage scale. It does not treat ± 1 as a forced analog limiter. Samples above digital full scale are processed through the circuit.


```
Vpeak = digital_peak × 17.37068 V
Vrms(sine) = Vpeak / sqrt(2)
dBu = 20 log10(Vrms / 0.775 V)
-20 dBFS peak sine ↔ +4 dBu RMS
0 dBFS peak sine ↔ +24 dBu RMS
```

For noise, calculate integrated RMS voltage directly. Applying the peak-sine calibration to digital noise RMS introduces a 3.0103 dB error. Threshold labels use the sine-peak calibration, while the meters show digital sample peaks.

Several limits, not one ceiling

The receiver pad keeps normal cell currents inside the VCA's working range. Output gain is restored by the differential drivers. Transformer magnetizing current, VCA current budget, driver compliance and load interact; any of them can set the practical headroom before a nominal rail limit.

In the complete-path sweep, the first 1%-THD input crossings were **+26 dBu at 20 Hz** and **+29.5 dBu at 1 kHz**, using 0.5 dB steps. These are sampled model results, not exact continuous thresholds or a specification across every frequency. Compression was held inactive with Threshold +12 dBFS.



48 kHz / Master 4x, noise off, unity makeup and fully wet. THD sums harmonics 2–20 below Nyquist. Input level is sine RMS dBu; the horizontal rule marks 1% THD. Values come from the saved production sweep.

Working level

Use normal mix levels first. Gain reduction and magnetic excitation are separate controls on behavior: lowering Threshold increases reduction, while raising input or Makeup changes the physical stages' drive. If the objective is a fair bypass comparison, match listening level manually. No auto-gain or output limiter is implied.

14 / Measurements

The production engine is exercised through a C ABI bridge; the measurement scripts do not substitute another compressor. The supplied final run at **48 kHz / Master 4x** passed **81/81** acceptance and regression checks. Separate C++/JUICE, pluginval and robustness tests cover other behaviors.

Complete-path measurement	Result	Conditions
Response spread, 20 Hz–20 kHz	0.065867 dB	+4 dBu; compression inactive
THD, 1 kHz	0.003799%	+4 dBu; compression inactive
Unity output noise	−89.533 dBu	Noise on; 20 Hz–20 kHz integration
Stereo gain difference	0.005849 dB	20 dB channel input imbalance
In-phase / anti-phase GR difference	0 dB	Matched polarity experiment
1 kHz small-signal gain spread	0.000006865 dB	44.1 / 48 / 96 / 192 kHz

The results above are software measurements, not evidence that the design has matched an SSL unit. Manufacturer whole-unit noise and THD+N figures can use different bandwidths, loading, balancing circuits and bypass states. SAK's THD result excludes noise.

Steady-state protocol

Tones are coherent with a 250 ms analysis window: 4 Hz spacing, including exact 20 Hz and 1 kHz. Cosine startup reduces initial DC flux bias. Lead-in is at least 350 ms and eight cycles; compression curves use 600 ms. Fundamental amplitude/phase use complex projection. THD sums RMS harmonics 2–20 below host Nyquist relative to the fundamental. Near Nyquist fewer harmonics are available.

Noise uses Welch PSD integration over 20 Hz–20 kHz, after the first second of a three-second zero-input run. Deterministic seeds make the enabled-noise measurement reproducible. The disabled result may include residual settling and numerical error.

Static and transient protocols

Static curves use 1 kHz, Threshold −24 dBFS, fastest attack and every ratio. Curvature means an effective fitted slope need not equal its panel label. Attack steps the carrier from −48 to −6 dBFS and reports 63.2%, 90% and 10–90% times. The known interpolation delay is removed for circuit-time interpretation.

Release steps down and measures 36.8% and 10% remaining reduction. Auto uses 80 ms, 2 s and 12 s histories and 15 s recovery. These are control-reduction experiments, not unloaded capacitor measurements. Raw-network tests separately verify common charging current, loaded poles and passive recovery.

15 / Validation in context

Bare VCA comparisons

Input / gain, 1 kHz	Model THD	Published 2180A typical / max
0 dBV / 0 dB	0.005006%	0.005 / 0.010%
+10 dBV / -15 dB	0.020628%	0.020 / 0.030%
-5 dBV / +15 dB	0.020582%	0.020 / 0.030%

The reference uses $R_{in} = R_{out} = 20\text{ k}\Omega$, $\pm 15\text{ V}$ and $25\text{ }^{\circ}\text{C}$. Model unity noise is -98.086 dBV versus -98 dBV typical / -97 dBV maximum; at $+15\text{ dB}$ gain it is -87.950 dBV versus $-88 / -86\text{ dBV}$. Noise bandwidth is 20 Hz – 20 kHz . These are component checks, separate from the loaded complete signal path.

Nonharmonic residual and overload

Mode	+21.5 dBu / 7 kHz	+31.96 dBu / 7 kHz
Studio 2×	-121.76 dBc	-43.93 dBc
Master 4×	-147.02 dBc	-70.44 dBc
Render 32×	-171.23 dBc	-125.64 dBc

The 48 kHz experiment uses a 2 s cosine and final 500 ms Hann window. DC and legitimate in-band harmonics have $\pm 20\text{ Hz}$ exclusions; power is corrected for the window's equivalent noise bandwidth. The metric includes aliasing, state modulation and residual leakage. It is not a universal alias floor or audibility claim. The overloaded case explains why Render matters when intentionally exceeding analog headroom.

Regression and host scope

The engine tests cover overload, non-finite input, reset repeatability, delayed bypass, ratio ordering, sidechain routing and timing invariants. JUCE tests check state serialization, malformed payloads, block partition independence, float/double behavior, live quality changes, fixed latency and parity with the measurement engine.

Pluginval strictness 10 passed on the supplied final VST3 with seed 5917445, at 44.1/48/96 kHz and block sizes 64/257/1024. Its GUI tests were disabled; separate editor tests cover geometry and gestures. The external Steinberg validator was unavailable and skipped. This is not a claim of testing every DAW.

The nonlinear convergence/recovery audit passed 3/3 checks. Magnetic terminal and forced-flux tests checked positive loop loss, positive incremental reluctance and electrical residual below 10^{-10} V in the dense sweep. They bound exercised states rather than proving all possible histories. No hardware null or external-DAW listening evaluation is claimed.

16 / Evidence boundaries

Evidence class	What it supports	What it does not establish
Historical schematics and operating guides	Control topology, switched values, loudest-channel behavior	Exact surviving-unit transient or knee curves
Published VCA equations and specifications	Branch reduction, noise/gain behavior, terminal anchors	Transistor-perfect trimmed die or historical 202 identity
Jensen terminal specifications	Selected winding/loading scales and output targets	Unique geometry, material loops or SAK input headroom
Derived timing equations	Coupled states and loaded/unloaded poles	A manufacturer's unstated step-response protocol
SAK calibration and measurement	Behavior of this implementation under stated tests	Exact full-unit hardware equivalence

Unavailable reference data include calibrated same-unit vintage attack/release traces, compression-knee curves, THD-versus-reduction and IMD sweeps; identified transformer minor loops and geometry; and a complete selected-device transistor netlist. These gaps remain open.

The initial MODEL_REVIEW.md records findings against an earlier implementation. The final architecture, calibration and measurements supersede it. Corrected VCA output-slew scaling, gain-dependent noise, signed Auto branch states, magnetic fitting and state handling are described here in their final form. The review remains in the source for provenance, not as a list of current defects.

Primary reference index

1. Gary K. Hebert, [An Improved Monolithic Voltage-Controlled Amplifier, AES 4055 \(1995\)](#): complementary cell, output slew and noise.
2. THAT, [2180 Series, 600029 Rev 02 \(2008\)](#): grade-A THD/noise anchors and test conditions.
3. SSL, [CF82E27 timing schematic, scan hosted by Gyraf](#): timing components and topology. CIRCUIT_REFERENCES.md records the wider 82E26/82E27 study.
4. SSL, [XLogic Stereo Compressor Owner's Manual](#), Appendix C; [THE BUS+ Guide 1.1](#), pp. 10, 32: operating context and separately scoped unit specifications.
5. Jensen, [JT-11P-1](#) and [JT-11-BMCF](#): terminal reference conditions.
6. TI, [NE5534 SLOS070D](#) and [NE5532 SLOS075K](#): amplifier-family scales; revisions are distinguished in the research ledger.

Manufacturer names identify references, not affiliation or endorsement. Their publications retain their publishers' rights; the source release links to them rather than relicensing them.

17 / Source, license and history

Free software

Copyright © 2026 SAK Audio. Original SAK-BUS source, interface artwork, build/installation scripts and this manual are released under **GNU AGPL version 3 only (AGPL-3.0-only)**, following SAK-TUBE. The full license accompanies every package. The software is provided without warranty. This software license does not impose a license on music processed with it.

JUCE 8.0.6 is used under its AGPLv3 option; its bundled VST3 SDK uses the GPLv3 option. Third-party notices and matching JUCE sources are included. The shared magnetic extraction retains its provenance and reference notices. Microsoft's runtime is a system prerequisite distributed under Microsoft's terms. Website code is separately MIT-licensed.

Corresponding source and rebuild

The matching **SAK-Audio-0.1.1-Source.zip** accompanies the installer and portable download. It includes production DSP/editor source, Blender scene and generator, rendered artwork, shared magnetics, CMake, tests, research, raw validation outputs, this editable manual and installer scripts. Development build trees and intermediate binaries are excluded.

Install CMake 3.22+, Visual Studio 2022 C++ desktop tools and the Windows SDK. The archive includes its JUCE checkout under `third_party/JUCE`. From the extracted project:

```
cmake -S . -B build -G "Visual Studio 17 2022" -A x64 `
  -DSAK_JUCE_DIR=third_party/JUCE
cmake --build build --config Release --parallel 4
ctest --test-dir build -C Release --output-on-failure
```

For measurements, install `tools/requirements-measurements.txt` and run `tools/measure.py` with the built `sakbus_bridge.dll` and `--verify`. Keep `reference_data.json` and the numerical baseline distinct. Do not regenerate a baseline merely to hide failures.

The manual is maintained in `manual/reference.md`. Install its locked Node dependencies, install Playwright Chromium, then run `manual/generate.mjs`. `installer/build-release.ps1` stages Release binaries, verifies the Microsoft runtime signature, builds Inno Setup, and writes versioned packages with SHA-256 manifests. `RELEASE.md` documents the exact commands.

Version record

0.1.0 · 22 September 2026 — Initial SAK-BUS Windows release: stereo VST3 and standalone, original circuit-informed control/VCA/magnetic architecture, Studio/Master/Render modes, documented research and validation. This release preparation adds the installer, matching dependency source, reference manual and SAK Audio product pages. The final product interface and DSP are preserved.

Use the versioned release manifest to identify exact binaries and source. Maintain binary, source, license and manual access together when redistributing. See `RELEASE-VALIDATION.md` for the packaging and installation test record.