FEATURES & TECHNICAL DETAILS

HLConvolver 2.15.7 for macOS

Native DSD & WavPack | Multichannel FIR | CoreAudio & UPnP

HIGH-PERFORMANCE FIR CONVOLUTION | UNIVERSAL 2 FOR APPLE SILICON & INTEL

HLConvolver is a high-performance FIR convolution application and plug-in suite for stereo and multichannel playback. It is part of the Hang Loose DSP Software Suite and runs on Apple Silicon and Intel Macs as a Universal 2 Standalone application, Audio Unit (AUv2) and VST3.

One convolver - local or networked audio

The Standalone application can sit between a music application and local hardware, between an application and a network renderer, or between an incoming UPnP source and either destination.

CoreAudio -> CoreAudio	CoreAudio -> UPnP
UPnP -> UPnP	UPnP -> CoreAudio

For example, a Mac player can send PCM through BlackHole into HLConvolver for processing and output to a USB DAC, or send the convolved result to a UPnP renderer. BlackHole is an optional third-party virtual device and is not included.

Software formats

- **macOS:** Standalone 2.15.7, AUv2 2.15.7 and VST3 2.15.7, with native arm64/x86_64 code.
- **Suite components:** HLConvolver AAX 1.2.6 and HLHost 2.1.6 remain separately maintained components.
- **Interfaces:** Native UI, local browser-based Web UI, optional headless operation and macOS Background Service.

HLConvolver processes host-supplied plug-in PCM or Standalone route audio; supported resolution, sample rate and channel count depend on the end-to-end path, hardware and FIR workload.

Native DSD input for Standalone UPnP playback

The Standalone UPnP-input routes accept native DSD from Sony DSF and uncompressed DSDIFF/DFF, plus positively identified DoP-over-WAV. Supported input rates are DSD64, DSD128, DSD256 and DSD512 in the 44.1-kHz family.

- Native DSD is unpacked/decoded and normalized to **352.8 kHz planar floating-point PCM** before speaker-position routing and FIR processing.
- DoP-over-WAV is recognized by a coherent alternating marker sequence across channels; ordinary PCM WAV is not mislabeled as DoP.
- DSF, DSDIFF and WAVE_FORMAT_EXTENSIBLE use supported speaker metadata. Unsupported or ambiguous multichannel layouts are not guessed.
- DST-compressed DSDIFF/DFF and DSD1024 are outside this input capability. PCM output remains the default; selecting native DSD input does not automatically select DSD output.

Lossless WavPack and other UPnP sources

Lossless integer PCM and supported 32-bit floating-point **WavPack (.wv)** sources are decoded by the Standalone UPnP path. Engineering validation includes 12-channel, 352.8-kHz / 24-bit DXD WavPack. Lossy/hybrid WavPack and WavPack DSD are not supported by this path.

The Standalone UPnP input formats include:

PCM & decoded formats	DSD / DoP formats
WAV, AIFF/AIFC, FLAC	Sony DSF
Lossless WavPack (.wv)	Uncompressed DSDIFF/DFF
Supported float WavPack	DoP-over-WAV with verified markers

UPnP output formats

UPnP PCM output is selectable at **16-, 24- or 32-bit** precision. DSD/DoP output is a separate configurable capability on supported UPnP-output routes. CoreAudio physical output routes deliver PCM. Input-decoding support must not be confused with native DSD passthrough in every output topology.

Interoperability notes

The UPnP ConnectionManager advertises supported FLAC, AIFF, native DSF and uncompressed DFF input types. Discovery and control support newer versioned renderer service types as well as compatible v1 devices, but individual player/renderer behavior remains device-specific.

Multichannel speaker-position routing

Multichannel UPnP sources can be mapped into a configured fixed 7.1.4 system bus using trusted source speaker positions. Routing is normalization, **not upmixing**: HLConvolver does not invent, duplicate, sum or increase channels simply because a larger system bus is present. Unused destination speakers remain silent.

Example fixed 7.1.4 output order

1 L	2 R	3 C	4 LFE
5 Ls	6 Rs	7 Lrs	8 Rrs
9 Ltf	10 Rtf	11 Ltr	12 Rtr

In **Auto** mode, HLC uses source-format layout information when it is authoritative. The explicit **Assume music layout from channel count** mode authorizes the documented 2.0, 4.0, 5.1, 7.1, 5.1.4 and 7.1.4 music layouts when metadata alone is insufficient.

Examples from system routing

Input source	Effective output positions
5.1 music: L, R, C, LFE, Lrs, Rrs	1, 2, 3, 4, 7, 8
7.1.4 source: L, R, C, LFE, Ls, Rs, Lrs, Rrs, Ltf, Rtf, Ltr, Rtr	Identity map, positions 1-12
5.0 Sony DSF: L, R, C, Lrs, Rrs, with authoritative metadata	1, 2, 3, 7, 8; absent destinations silent

Routing status and channel masks

The Web UI reports the active source layout, its metadata/assumption basis, the system layout and the effective source-to-speaker map. The CoreAudio backend separately supports validated physical-channel selections on multichannel devices; a logical 8-channel choice need not mean physical channels 1-8.

Where no input device has yet been configured, Standalone Audio Settings remains accessible for output-only setup. Physical device switching and edge-case reconnect behavior are hardware dependent and remain subject to field acceptance.

FIR convolution and resampling

HLConvolver supports long FIR filters, rate-specific configurations, multiple filterbanks and selectable convolution processing. With the bundled **r8brain** resampler, filter data can be resampled when a matching sample-rate FIR has not been provided. Direct stereo WAV FIR loading also benefits from this normalization in the 2.15.x line.

Processing choices

Convolution mode	Description
Zero	Uniform partitioning, zero-latency setting
4096	Non-uniform convolution
8192	Non-uniform convolution
16384	Non-uniform convolution

Filterbanks have independent **transition-band, stopband-attenuation and filter-phase** resampling controls in the Native and Web UI. FIR latency is reported in samples and milliseconds, and plug-ins report applicable latency to hosts for compensation.

Measured non-uniform convolution examples

During earlier macOS 2.12 engineering on an M1 Mac mini, a stereo 65,536-tap filter at 48 kHz resampled to about 963,000 taps per channel at 705.6 kHz. The average real-time DSP workload was 65.0% with Zero and 29.9% with non-uniform 16384 - approximately a 54% reduction in average processing time.

A 12-in/12-out configuration at 352.8 kHz with 23 active convolution paths and over 11 million resampled FIR coefficients produced:

Convolution mode	Average real-time load	Observed result
Zero	132%	Unable to sustain clean playback
Non-uniform 8192	93.6%	Clean playback
Non-uniform 16384	77.5%	Clean playback

Performance caveat: These are measured examples from one historical test machine, not universal product limits. Performance changes with processor, OS, scheduling, filter length and topology, sample rate, route, host, channels and device buffering.

Automatic CoreAudio sample-rate following

The Standalone CoreAudio input can follow valid source sample-rate changes without requiring the user to manually reselect a rate for each track. Actual acceptance depends on the physical input and output endpoints, selected channel counts, filter readiness and host behavior.

Split input and output ownership

CoreAudio -> CoreAudio	CoreAudio -> UPnP
Local physical output is owned by HLC. The split engine coordinates the source-rate transition, DSP preparation, local DAC reopen, ASRC/output buffering and output-clock readiness.	The physical local output is not the clock owner. The input/DSP path supplies post-convolution PCM to the UPnP /live stream, without reopening an unrelated DAC to admit the new source rate.

This separation is important if, for example, a local HDMI output supports only 48 kHz while a downstream UPnP renderer accepts an 88.2-kHz stream. It does not remove the downstream renderer's own startup and rate-change behavior.

UPnP transport and renderer control

The Standalone provides an HTTP **/live** WAV stream for network-renderer output, with rate-appropriate WAV headers and PCM after FIR. It supports compatible renderer discovery, AVTransport SetURI/Play and relevant track-transition control, including improved new-rate replay bootstrap and queued next-track handling.

UPnP controller/renderer implementations differ. Some devices can take longer to acknowledge transport commands, rediscover a renderer, or open the replacement stream. These behaviors are being measured independently of CoreAudio and FIR throughput; no generic one-second, gapless or start-attack guarantee is made for every renderer.

Control and service operation

- Native application UI and Web UI at **http://127.0.0.1:8081** during local Standalone operation.
- Headless operation and macOS Background Service for unattended use; service and manually launched application must not compete for the same audio resources.
- Web controls for routes, CoreAudio devices, channel masks, convolution mode, filterbanks, FIR resampling, UPnP renderer and output format, speaker mapping, and diagnostics.

Diagnostics, support logging and disk safeguards

HLConvolver 2.15.x adds managed, bounded diagnostic logs and a one-click Support ZIP intended to retain the logs needed for rate, route and device investigations. Signal-path diagnostics are managed outside the real-time audio callback wherever supported.

Managed logging property	2.15.7 policy
Log segment size	8 MiB maximum per recognized managed file segment
Rotation history	Two rotated segments per active managed log
Managed-log directory budget	500 MiB rolling byte budget for recognized managed logs
Support ZIP	Disk-backed archive with manifest and included/omitted counts
Default location	~/Library/Logs/HLConvolver

Scope: The 500 MiB value is not an all-disk guarantee. Unsupported custom log names, other applications, alternate log folders and historical files written by an older version can fall outside managed retention. Existing giant 2.14.0 logs may require safe, stopped-process cleanup; the companion README documents the read-only check and manual recovery steps.

Installed components and deployment

The HL DSP Suite macOS package is intended to install Standalone, AUv2 and VST3 as Universal 2 binaries. AAX 1.2.6 and HLHost 2.1.6 are separately signed components. HLConvolver is a perpetual-license product; pricing, activation rights and the current public installer should always be confirmed from the Accurate Sound download page.

Technical scope and compatibility

- Native DSD/WavPack input applies to Standalone UPnP-input routes; hosted plug-ins receive audio in their host-provided PCM format.
- UPnP controller and renderer interoperability varies. Some combinations may need additional time for discovery, playback starts or sample-rate changes; gapless operation is not guaranteed with every third-party endpoint.
- Audio-device capabilities, filter readiness and route configuration determine which sample rates and channel mappings can be used in a particular system. Consult the README for installation and safe log-management instructions.

Free trial and documentation: A 14-day trial, six-activation perpetual license, current macOS download and additional documentation are available at accuratesound.ca/hang-loose-dsp-software-suite/. The Suite includes free future updates and is not subscription-based.