
Voxengo Elephant User Guide



Version 5.7

<https://www.voxengo.com/product/elephant/>

Contents

Introduction 3

 Features 3

 Compatibility 3

User Interface Elements 5

 DC Filter 5

 Dithering 5

 Limiter 6

 Limiter Mode Editor 7

 Stats 10

 Metering 11

 Level Meters 11

 Clipping Indication Note 12

Credits 13

 Beta-Testers 13

Questions and Answers 14

Introduction

Voxengo Elephant is a mastering limiter plug-in for professional audio and music production applications. The most remarkable feature of this signal limiter is its sonic transparency. Elephant brings sound limiting and loudness maximization without audible “fuzz” and “pumping” sonic artifacts.

Elephant is a highly competitive plug-in for mastering applications: it features a variety of limiter modes that can be deeply customized, and a comprehensive set of level metering tools, including headroom estimation, EBU R128 and K-system metering. Elephant's built-in linear-phase oversampling is an important element in achieving high-quality peak limiting and loudness maximization without inter-sample overshoots (true-peak limiting).

Elephant limiter can be used for mixing and mastering of both stereo and multi-channel music and sound material, at any standard sample rate. Limiting process can be visualized both as gain-reduction meter and real-time waveform graph.

Elephant, also featuring a DC offset filter, and a high-quality bit-depth converter with an optional noise-shaping, allows you to finish your music and audio productions with a spark!

Features

- Transparent signal limiting action
- 13 predefined limiter modes
- Limiter mode editor
- Optional release stage
- Noise-shaped bit-depth converter
- RMS, true peak clipping statistics
- EBU R128 LUFS/LU metering
- DC offset removal filter
- Waveform graph view
- Stereo and multi-channel processing
- Internal channel routing
- Channel grouping
- Up to 8x linear-phase oversampling
- 64-bit floating point processing
- Preset manager
- Undo/redo history
- A/B comparisons
- Contextual hint messages
- All sample rates support
- 24 ms compensated processing latency

Compatibility

This audio plug-in can be loaded into any audio host application that conforms to the AAX, AudioUnit, VST, or VST3 plug-in specification.

This plug-in is compatible with Windows (32- and 64-bit Windows XP, Vista, 7, 8, 10 and later versions, if not announced otherwise) and macOS (10.11 and later versions, if not announced otherwise, 64-bit Intel and Apple Silicon processor-based) computers (2.5 GHz dual-core or faster processor with at least 4 GB of system RAM required). A separate binary distribution file is available for each target computer platform and audio plug-in specification.

User Interface Elements

Note: All Voxengo plug-ins feature a highly consistent user interface. Most interface elements (buttons, labels) located at the top of the user interface are the same in all Voxengo plug-ins. For an in-depth description of these and other standard features, and user interface elements, please refer to the “Voxengo Primary User Guide”.

DC Filter

This group of controls allows you to enable and fine-tune the DC filter. DC filter is applied before all other processing stages, and is usually used for removing unwanted signal’s DC (constant) offset which reduces headroom of a sound track and may affect playback quality of this track due to non-symmetry induced by DC offsetting. The DC filter is implemented by means of a high-pass filter.

The “Freq” parameter specifies the corner frequency of the DC (high-pass) filter.

The selector below specifies DC filter’s type. The “Steep” filter types offer a steeper-cutoff frequency response, but a higher phase-shift. The “Soft” filter types have a smoother-cutoff frequency response and a lower overall phase-shift. The number before the filter type denotes dB/octave roll-off of the filter (e.g. “-18 Soft” means “-18 dB/octave soft filter”).

Note that since DC removal filter induces phase-shift, it may affect sonic character of a sound material. Another possible disadvantage of such phase-shift is that it changes the envelope of the signal, having potential of introducing higher peak levels. For best results, you may specify DC filter’s frequency as low as possible, and use non-steep filter types.

Commonly, it is suggested to use “Freq” values below 5 Hz. If your sound system does not feature a sub-woofer and its frequency response goes e.g. down to 50 Hz only, it is suggested to tune DC filter to values around 35 Hz due to absence of auditory control of such low-frequency content.

Dithering

Dithering stage – when enabled – is applied last in the processing chain of this plug-in, after the output gain is applied. Dithering allows you to perform finalizing bit-depth conversion. If you plan to perform further processing on the track where this plug-in is inserted to, or if you export the track to floating-point file format, you should leave the dithering stage disabled since the dithering effect will be lost if you apply any further processing. Note that same dithering stage settings are applied to all channel groups. Elephant uses PRVHASH dithering noise generator invented by Aleksey Vaneev, for all settings.

The “Bit-Depth” parameter allows you to specify dithering depth with 1-bit steps. It should be usually equal to 16-bit for CD production, 24-bit for pre-masters and storage purposes. Lower and intermediate bit-depth settings are provided for non-conventional uses. You can use the 8-bit setting to hear how different noise-shaping profiles “sound”.

The “Noise” parameter controls probability density function (PDF) of the dithering noise. The “TPDF” is a triangular PDF (2-bit peak-to-peak), the “Gauss” is an

approximation of the Gaussian PDF, the “1-Bit” is 1-bit peak-to-peak uniform PDF noise. Note that it may be hard for you to decide which one is better since the difference between the offered PDFs is hard to hear. The “Gauss” mode is the mode used in the original Elephant version 2 while the “TPDF” is an industry standard PDF for sound-signal dithering. The “1-Bit” option offers the least possible noise-floor; in comparison, the “TPDF” mode produces 1.8 dB higher noise-floor while the “Gauss” option produces 2.6 dB higher noise-floor. All these modes are equally efficient at decorrelating the rounding error from the signal. The “1-Bit” option can be selected if you need a “blanking” dither which does not add noise in silent parts of the program material; this may make a further lossless data compression more efficient. Note that when the dithering is enabled, it is important that Elephant is placed last in the FX chain, with the master bus gain set to unity (0.0 dB); also, during audio export, the audio host application’s dither should be disabled since it may affect statistical properties of Elephant’s dither.

The “Shaping” parameter selects the noise-shaping function. Noise-shaping is used to increase the signal-to-noise ratio (SNR) in the lower spectral regions at the expense of decreasing the SNR in the higher spectral regions which are usually cannot be heard by human hearing system. When left “Off” no noise-shaping will be performed. The “Equal” setting enables equal-loudness contour noise-shaping which follows human hearing threshold envelope thus greatly enhancing a subjective SNR. The “Classic” setting enables noise-shaping originally introduced in Elephant version 2.

Limiters

The “Mode” selector loads a pre-defined factory or user mode preset. Press the “Gears” button to open limiter mode editor popup window. If changes were made to a loaded preset, the name of the loaded preset will be changed to “User”.

The “In Gain” knob is used to control the volume of the signal that enters the limiter stage. Internally, audio signal is limited to 0 dBFS. So, if the audio signal's level is mostly staying below 0 dBFS threshold, you should increase the In Gain in order for the limiter to start working. Note that the “Knee” specified in the Limiter Mode Editor may decrease the threshold level at which the limiter starts working.

The “Out Gain/Ceiling” knob can be adjusted to change the output volume after the limiting was performed. This control effectively works as signal’s ceiling-like control, as it changes the post-limiter output gain. The “Out Gain/Ceil” knob has a lesser sensitivity to mouse movements so it can be changed with a higher precision than the “In Gain” knob.

The “Release” control affects release stage duration. Note that you should enable the “Rls” switch before this release stage control becomes available for adjustment. Release setting is unavailable for some algorithms like the “Clip” algorithm. The release stage in the “AIGC” algorithms produces a very strong “pumping” action at high time settings: this may or may not be desirable depending on your needs, but such operation is intentional. Longer “Release” settings reduce distortion and retain transients at the expense of both overall achieved loudness (which becomes lower) and an increased “pumping”. For best results you should probably use low “Release” values since they allow Elephant to retain transients without reducing the overall achieved loudness much.

The “Eval” switch enables limiter’s evaluation mode designed to evaluate the change in material’s dynamics applied by the limiting. When this switch is enabled, the output signal level will be approximately matched to the input signal level using 80 dB SPL equal loudness contour filtering. This way it is very simple for one to hear any distortion that is introduced by the limiter. Note that you can use the “Bypass” switch together with the “Eval” switch to compare the original input signal to the limited signal at equal loudness.

The “Graph” button opens real-time gain-reduction and waveform graphs popup window. The gain-reduction graph (“GR”) displays gain reduction happening in all channels of the current channel group. The waveform graph (“Wave”) displays peak input and output waveform levels in all channels of the current channel group.

Any change to the aforementioned parameters and most limiter mode parameters that affect output statistics resets the output statistics.

Limiter Mode Editor



This popup window offers you an opportunity to define limiter’s operation mode precisely. The “Presets” button opens the preset list where you can store or recall your own limiter mode presets. These user presets – when defined – are also displayed in the main limiter mode selector list.

Note that almost every parameter in this editor offers you a compromise between sound’s loudness and transparency.

The “Alg” selector allows you to select the limiting algorithm used by the current limiter mode. All algorithms offer a brick-wall limiting operation meaning input signal you are limiting would not go over 0 dBFS level (given the “Out Gain” parameter stays at or below 0 dB). Note that the algorithm has no direct correspondence to the mode you choose via the “Mode” selector (“mode” is a complex intermix of parameters while the algorithm is one of the parameters of the mode). Each algorithm has its own distinctive features:

- The “Clip” algorithm works as a “plain” hard saturation limiter: it sounds cleaner than the ordinary clipping you can get overdriving the signal in a sound editor. The “Clip” algorithm is not a suggested type of limiting since it can create much distortion and many harmonics. At the same time, this algorithm allows you to preserve the frequency balance of the sound material you are working on, and it also tends not to smooth out the attack transients. The “Clip” algorithm is best used with oversampling enabled.
- The set of “AIGC” algorithms represents algorithms with a built-in “automatic input gain control” (AIGC), each with a specific dynamic response. The “AIGC” algorithm internally adjusts the specified “In Gain” value striving to deliver a cleaner limiting sound. Even though AIGC can give a cleaner sound, at high “In Gain” settings it can create a “pumping” sound. The “AIGC-3”

- algorithm is a “cleaner” variant of the “AIGC-1” algorithm; the “AIGC-4” algorithm is a “cleaner” variant of the “AIGC-2” algorithm.
- The “EL-1” algorithm is a legacy limiter algorithm which was introduced in the very first version of Elephant mastering limiter. With this algorithm you can usually achieve high loudness, but at the expense of a relatively high distortion. This algorithm is not recommended for regular use.
 - The “EL UNI” algorithm is a “unified” limiting algorithm, implementing best limiting techniques and offering the highest level of flexibility. This algorithm implements distortion minimization technique and also contains several dynamic response types that adjust the dynamic response of the limiter.
 - The “EL-4” algorithm which was introduced in Elephant version 4 is a reworked version of the “EL UNI” algorithm. The “EL-4” algorithm is based on a time-proven topology of the “EL UNI” algorithm, but with more clarity added to it via novel control elements that reduce distortion. Note that the “EL-4” algorithm may sound subjectively “quieter” and with less “bite” than the “EL UNI” algorithm even if the objective signal’s RMS is about the same: this is because “EL-4” saturates the sound less while the saturation of the “EL UNI” algorithm increases subjective “bite” of the sound. The “bite” can be increased by reducing the “Lookahead” parameter (see below).
 - The “EL-C” algorithm which is the newest addition to Elephant version 4, uses a novel approach to limiting, based on compression found in compressors, but applied in a special simultaneous feed-forward and feedback manner. This algorithm has fixed internal timing constants making it very easy to setup and use. This algorithm is not meant to completely surpass and replace the “EL-4” algorithm, but in some music genres, the “EL-C” algorithm produces less coloration and less transient distortion. It also has a different release character, and sounds “smoother” than the “EL-4” algorithm.
 - The “EL INS” is a clipper-like (“instant”) limiter algorithm, without release timing elements, but with a look-ahead transient smoothing taking place. It is similar to the original “EL-1” algorithm, but is more computationally-efficient. This algorithm is effective for EDM production where signal is limited to bring more energy to it. This algorithm, like the “Clip” algorithm, produces more distortion in comparison to other algorithms. The “Lookahead” parameter affects the look-ahead time and the degree of transient smoothing taking place.
 - The “EL-5” algorithm introduced in Elephant version 5 takes the best from “EL-4” and classic “AIGC” algorithms. The “EL-5” is not as “hard” as the “EL-4” algorithm, but it is much more precise than the “AIGC” algorithms while producing an alike smooth sound. This algorithm better retains dynamical structure of the program material while at the same time adds a unique “sparkling and fresh mojo” to the resulting sound. This algorithm is responsive to the “Timing” and “Rls/Atk Ratio” parameters that can be adjusted to tune the algorithm further.

The “DRC” selector enables the “Dynamic Release Control” limiter processing stage which dynamically adjusts limiter’s release envelope so that it better resembles the original unlimited sound’s release form, offering a stronger program material dependence (the “Release” stage does not need to be enabled for the “DRC” stage to work). This stage audibly reduces loudness swings and “pumping” artifacts generated by strong limiting thus producing a much more stable, smooth and “punchy” sound. Additionally, this stage reduces harmonic distortion. This stage is universal and can be applied to any algorithm. For backwards compatibility and taste-related reasons this stage can be disabled, but it is suggested to enable this stage whenever possible, because it makes the limiting statistically better in many aspects.

You may disable this stage if your goal is to maximally flatten the dynamic response of the sound material by means of limiting. As the “DRC” stage closely interoperates with the selected algorithm and the optional Release stage, in overall the existence of the “DRC” stage effectively doubles the number of available operation modes in Elephant. The “DRC Exp” mode enables the expanded control range which further reduces loudness swings at the expense of higher transient distortion by additionally adjusting the attack envelope; this mode also provides a higher overall “punch”.

The “EL Dyn” control selects additional dynamic response type used when the “EL-1”, “EL UNI” or “EL-4” algorithm is engaged. Depending on the selected type Elephant will reveal or hide subtle features of the sound when limiting happens. The “Off” option can be used to disable this additional dynamic response. Note that sonic difference between all the offered response types is subtle – you may choose the type in accordance with your subjective feelings alone. E.g. one “EL Dyn” mode may recover a little bit faster after gain reduction happened, another mode may recover later: this – in turn – produces a little bit different subjective “feel” of the sound. The “EL Dyn” setting is most apparent when heavy compression is going on. If gain reduction hits 1-2 dB only you won't hear much difference between the offered “EL Dyn” modes.

The “Ch. Link” knob controls the amount of link between channels in the group. In a fully unlinked mode Elephant limits all channels of a group independently of each other. This mode can give a higher signal gain in comparison to the linked mode. However, unlinked mode can also damage the stereo-field information. For example, if you have a floor tom panned to the left with some signal in the right channel, unlinked limiting can bring the right channel up, destroying the original panning of a floor tom. Fully linked mode limits all channels in a group simultaneously, preserving the stereo-field information.

The following “Knee”, “Timing”, “Rls/Atk Ratio” and “Rls Shape” parameters are available for the “EL-1”, “EL UNI”, and “EL-4” algorithms only.

The “Knee” parameter allows you to control soft knee of the limiter’s transfer function. A wider knee allows you to achieve a smoother overall sound, lower knee values are best used for harder sound.

The “Timing” parameter which is specified in percent adjusts (multiplies) limiting algorithm’s internal timing constants. The value of “40” maximally reduces them while the value of “500” maximally increases them. Higher values produce a less “compacted”, more expanded sound at the expense of higher transient distortion which is immediately obvious when mid-frequency range sound like piano or acapella is processed. Please use values below 250 if you want to get sound which is as transparent as most other competing limiters. This parameter does not affect the “EL Dyn” modes.

The “Rls/Atk Ratio” parameter adjusts internal release timing constant’s ratio to attack timing constant. The lower ratio values that are close to “1:1” produce louder and brighter overall sound, with non-pronounced lower frequencies. The higher values produce quieter overall sound with better pronounced lower frequencies and a bit dimmed higher frequencies.

The “Rls Shape” parameter adjusts algorithm's internal release shape: values close to -1 produce a “fatter”, louder, more compressed sound, with fast release; values close to 1 and beyond produce a “thinner”, quieter, more expanded sound, with delayed release and reduced loudness swings. The “Rls Shape” values near 2 offer an even

higher transparency, and additionally produce a bit “louder” overall sound impression due to a close resemblance of the S-curve which has a central region with a relatively fast release. The attack shape cannot be changed, and it always resembles a “fast” logarithmic shape. Note that this parameter together with the “Rls/Atk Ratio” does not affect the operation of the optional “Release” switch-enabled stage.

A set of the “Trans” controls collectively affects limiter’s transient form. Using these controls you can trade between sonic cleanness of limiting and maximal achieved loudness.

- The “Trans” switch selects transient form. This switch affects overall coloration and transient response. The “Classic” option implements transient response originally introduced in the Elephant version 2. The “Sharp” option offers a sharper and more instant transient response which, however, adds more harmonic coloration in comparison to the “Classic” option. The “Asymm” option offers you an asymmetric transient form which can be also called an “analog” transient form since it features a causal response (in terms of digital signal processing); the coloration strength of this option is in-between the “Classic” and “Sharp” options.
- The “Lookahead” knob adjusts transient’s duration (the time it takes for the limiter to go from “no limiting” to “full limiting”). A lower time value usually sounds “harsher”, but provides a higher loudness boost, and a higher overall “punch”. Higher look-ahead values sound smoother, retain transients’ shape and their coloration better, but provide a lower overall loudness boosts, and may sound a bit “buried in the mix”.
- The “Tr. Shape” parameter adjusts the transient form. Values below “0” produce a smoother form, values above “0” produce a sharper form.

The “Trans View” displays an approximate look of limiter’s transient. Transient's look does not depend on the “Lookahead” parameter: transient stretches with the “Lookahead” parameter, but the overall look of the transient does not change.

Stats

This panel displays historic (since the last reset, playback start, or parameter change) statistical information for each output channel of the currently selected channel group.

The “Integrated” indicator displays integrated signal power estimation. The overall multi-channel signal power is also displayed with a larger font. Note that if the audio host application provides input channel names to the plug-in, the overall power estimation excludes the “LFE” channel, and applies 1.5 dB correction to “Ls” and “Rs” channels.

The “Max Crest Factor” indicator shows maximum crest factor (difference) between the integrated RMS and peak RMS values reached. Peak RMS value is not displayed anywhere on the user interface. 50 ms time window is used to estimate the peak RMS value. You may add the “Integrated” and “Max Crest Factor” values together to obtain a peak RMS value. Note that when comparing the “Max Crest Factor” value in Elephant to plug-ins of other companies the peak RMS time windows should be matched in the plug-ins in comparison or otherwise the readings will be different, with shorter peak RMS time window (e.g. instant peak value) showing a greater “Max Crest Factor” value.

The “Max Gain Reduction” indicator shows a maximal instant gain-reduction value reached.

The “True Peak”/“Peak” indicator displays 1-sample output true peak/sample peak level. Basically, this indicator detects samples that will be clipped if signal is converted to a limited bit-depth. The “TP” switch changes between “true peak” and “sample peak” level and clipping indication. Note that this switch does not enable the “true peak limiting” – at least 4X oversampling is still required. The true peak clipping indication is also useful when using noise-shaping at low bit-depth dithering settings since noise becomes powerful enough to go over 0 dBFS. When clipping indicator grows, it is a sign that you should lower the “Out Gain” setting.

The “Reset” button can be used to reset counters and start statistical accumulation over.

Metering

The “Metering” selector allows you to choose the meter biasing mode. The “dBFS”, “dBFS.30” and “dBFS.15” modes apply no biasing. These three modes have a differing level meter range. The “dBFS+3” mode adds 3 dB to level meter indicator, the “Integrated” values (peak level indication is unchanged in comparison to the “dBFS” modes).

The “K-20”, “K-14” and “K-12” metering modes offer various K-system biasing modes originally offered by Bob Katz. The “K-20 C”, “K-14 C” and “K-12 C” implement level calibration modes. In these “K” modes, level meter’s RMS integration and release time constants are fixed to 600 ms as per K-system specification, level meter’s integration time is not affected by plug-in’s “Settings” window. Weighting is not used in all these “K” modes.

The “LUFS” and “LU” modes implement EBU-R128 K-weighted metering standard, by following ITU-R BS.1770-4 and Tech 3341-2016 notes (most elements of the “EBU Mode” are offered, except EBU +18 scale). The modes with the “ML” suffix display the momentary loudness on the level meter, with 0.4-second integration window, the “SL” modes reflect the short-term loudness, with 3-second integration window. Both modes also display integrated loudness on the “Integrated” statistics panel. In these modes, level meter’s integration time is fixed and is not affected by plug-in’s “Settings” window. Integrated loudness level is updated only twice a second, and only considers the recent 5 minutes of loudness measurements since reset/restart, due to real-time CPU restrictions. The “Max Crest Factor” calculation in LUFS/LU modes uses the same technique as in the other metering modes; it is, however, performed using RMS integrated loudness level, not EBU-R128 loudness level, but it uses K-weighting.

Level Meters

Elephant features three RMS level meters, with scales shown in decibel. Peak level indication is present on all meters. The “In/Thrs” meter displays input level meter with the limiting threshold mark that is tied to the “In Gain” knob. The “GR” is an absolute gain-reduction meter showing limiter’s operation. The third meter shows plug-in’s master output level. You may take a look at the “out/in” indicator to see the average loudness change taking place, which may be accounted for with the “Out Gain” knob, or via clicking on the numerical value of this indicator.

Clipping Indication Note

Even though Elephant was designed as a “brickwall limiter”, the clipping indicator in both the plug-in and audio host application may increase in some cases: when you change limiter’s parameters – in this case the clipping is only transient; when oversampling is used, or when dithering noise is very loud (at low bit-depth) – in this case clipping can be permanent and will require Out Gain parameter reduction.

The reason clipping indication may increase when oversampling is used is the filtering performed by oversampling process which may induce new peak levels. This is especially possible if the audio material being limited contains strong audio content above 20 kHz or when the “Lookahead” values below 0.50 are used. The only way to avoid clipping in such situation is to decrease the Out Gain parameter.

When oversampling is not in use, Elephant is guaranteed to deliver the “brickwall” limiting performance, based on sample peak measurements (“true peak” readings may differ).

Credits

DSP algorithms, internal signal routing code, user interface layout by Aleksey Vaneev.

Graphics user interface code by Vladimir Stolytko. Graphics elements by Vladimir Stolytko and Scott Kane.

This plug-in is implemented in multi-platform C++ code form and uses “zlib” compression library (written by Jean-loup Gailly and Mark Adler), “base64” code by Jouni Malinen, FFT algorithm by Takuya Ooura, filter design equations by Robert Bristow-Johnson, VST plug-in technology by Steinberg, AudioUnit plug-in SDK by Apple, Inc., AAX plug-in SDK by Avid Technology, Inc., Intel IPP and run-time library by Intel Corporation (used under the corresponding licenses granted by these parties).

Voxengo Elephant Copyright © 2003-2025 Aleksey Vaneev.

VST is a trademark and software of Steinberg Media Technologies GmbH.

Beta-Testers

Alan Willey

Ben Williams

Dave Huizing

Edward Rendini

Eugene Medvedev

gl.tter

Jay Key

Julien K/BIDI

Lesha

Matthew Fagg

Michael Anthony

Mike Roland

Murray McDowall

Niklas Silen

Roopesh Patel

Steffen Bluemm

Special thanks go to Niklas Silen for several plug-in preset ideas.

Questions and Answers

Q. Why does this plug-in have latency?

A. This plug-in has latency due to the look-ahead analysis algorithms it utilizes. The latency depends on the project's sample rate. At 44.1kHz it equals to 1041 samples, and grows steadily with the sample rate.

Q. Is it just my ears, or does Elephant actually filter out low-end frequencies to attain the higher levels? It seems to have “thinning out” effect.

A. It's probably not precise to call that a “thinning out effect”, or at least it is not intentional. It's simply a normal effect of limiting where lower frequencies are more heavily affected (lowered in volume due to the high amplitude) than higher frequencies (which are, to begin with, are usually lower in amplitude than the lower frequencies).

If you are ending up losing some low-end frequencies, it is suggested to reduce the “In Gain”, or insert a pre-equalizer with some additional gain set at the lower frequencies.

With limiters, pre-equalization is the only way to get the desirable frequency balance after limiting is applied. You can also try to use the “Clip” limiting mode which tends to preserve the lower frequencies at the expense of higher overall harmonic distortion.

Q. How much version 3 of Elephant is different from version 2, sound-wise?

A. Version 3 (and above) of Elephant recreates sound of version 2 for the most part except minor corrections and enhancements like the “Knee” control and the “EL Dyn” type selection.

It should be also noted that preset format of version 3 is incompatible with version 2 presets, but both version 2 and version 3 plug-ins can be used in the same project without conflicting each other, allowing for a smooth transition from version 2 to version 3.

Q. Do you have any recommendations on the most usable “Lookahead” and “Tr. Shape” values?

A. These settings must be selected according to the sound material you are processing. The best way to select some appropriate values is to first enable the “Eval” switch and overload the Elephant with a high “In Gain” setting. There will be much distortion coming from the plug-in overall, but this is not what you are looking for. In this overloaded mode you can select both the “Lookahead” and “Tr. Shape” values and hear the difference obviously. Values that sound best for the high gain will sound good for the normal gain, too. The “Release” knob can be tuned last, after you have set the “In Gain” to its final position.

You can also use this general tuning sequence:

- Enable the “Eval” switch.
- Select a desired “In Gain” value.
- Select an appropriate “Lookahead” value. For example, heavy and punchy music needs a lower “Lookahead” value. Ballads and soft orchestral-like music usually need higher values to preserve transparency.
- Select an appropriate “Tr. Shape” value. Negative values sound “softer” while positive values sound “harder”.
- Set the “Out Gain” value so that the output signal level generally does not go above -0.1 dB.
- Select an appropriate “Release” value (if this parameter is available in the selected limiter mode). At this stage an additional clarity can be achieved. Higher “Release” values usually provide better overall clarity.
- Disable the “Eval” switch.

Q. I wanted to ask, have you ever considered a next version of Elephant that would have an absolute ceiling (like the L2 does)?

A. The “Out Gain” knob works exactly like a ceiling control, the only difference is that it is a knob – not a slider.

Q. Is AIGC supposed to set the optimal Input gain? Should I take it, that it means it sort of turns the gain down if you overdo it? Is it sort of an intelligent “In Gain” mode?

A. AIGC adjusts the “In Gain” internally, and does it adaptively. It is not supposed to visually adjust the “In Gain” value you specify; it strives to reach that “In Gain” value internally, but at the same time it strives to produce less distortion. In other words, it is exactly an intelligent “In Gain” mode.

Q. I was using the plug-in in the “Clip” mode... Clipping was almost identical to just overloading my DAW!

A. The “Clip” mode is best used with oversampling enabled: this way it will sound much better than a simple overloading you can get from your DAW.

Q. If there are “AIGC” modes available, what is the point of having “EL” modes? Are there certain types of music “EL” modes would work better with or something? Or some sort of other advantage?

A. The “EL” modes can still be useful. As this user guide states, the “AIGC” modes may give some slight pumping. The “EL” mode may work better for the cases where you dislike “AIGC” mode pumping.

Q. How does Elephant handle inter-sample overshoots?

A. Inter-sample overshoots can be handled by Elephant only when oversampling is active. At “4x” oversampling Elephant handles these very well.

Q. I would like to know what the averaging window Elephant uses to calculate the RMS power?

A. RMS power on the “Statistics” panel is measured for the whole duration since the moment the meters were reset or playback was restarted. So, after adjusting various controls you should reset the meters to see the recent RMS power. When you change certain plug-in’s parameters the reset happens automatically.

Q. Is it a bad idea to dither twice? I believe it does actually sound better, though I may be fooling myself.

A. An added dithering noise can sound nice at times. Probably, there can be no general advice regarding dithering and especially noise-shaping: if you like how it sounds then use dithering twice.

Q. When setting the bit-depth for dithering in Elephant, what setting is best to use? Specifically, I mix down or render at 24 bit and later to 16 bit for CD burning. Because I ultimately end up at 16 bit, would the 16 bit dither setting in Elephant be appropriate?

A. Yes, 16-bit setting would be appropriate, especially if you are not going to use any other dithering/bit-depth conversion process after Elephant.

Q. Can I use Elephant as a dithering plug-in, switching off all other processing?

A. Limiter processing cannot be switched off in Elephant. But as long as your material does not have peaks above 0 dBFS and oversampling was left disabled, limiter does nothing except consuming CPU a little. Note that you should also set the “Knee” parameter to 0% so that no soft-knee limiting occurs.

Q. Is dithering applied after oversampling?

A. Yes, it is applied after the oversampling stage.

Q. I'm looking for help with DC Filters. I would like to understand applications for DC filters.

A. DC filtering is technically useful to apply below 20 Hz. This removes unnecessary low-frequency oscillation and DC offsetting, and in some cases allows one to achieve higher undistorted loudness boosts.

Do not forget that a human person generally can't hear anything below 20 Hz and that is the reason various weird things may happen in that region due to the lack of audible control. If you are not using a subwoofer which can reproduce frequencies below 50 Hz during mixing or mastering, it may be useful to unconditionally remove low frequencies below 40 Hz, because you do not have control over these lower frequencies. Otherwise, when later playing back your recording on a subwoofer-powered audio system, some unexpected sounds may pop-up.

Q. I noticed you can choose any DC filter frequency desired. Is this a “shelf” number and works as a filter to remove anything below that number or is it a target Hz to remove that frequency only? What is a good practice or starting point for DC filtering?

A. It is a corner frequency, and everything below it will be removed. DC filter is also known as the “High-pass” filter. The “-24 Soft” at 10 Hz is a pretty good starting choice in my opinion.

Q. Does Elephant use linear-phase filters for DC filtering?

A. The filters available for selection in the “DC Filter” window are not linear-phase, but they offer differing degrees of minimum phase shift depending on the slope and type. The higher the slope, the higher the overall phase shifting. All “Soft” filter types produce less amount of phase shift in comparison to the “Steep” filter types.

Q. DC offset measurement tools show that Elephant introduces DC offset when it is limiting a signal.

A. This is to be expected since all dynamic range adjustment processes introduce a small amount of DC offset. This DC offset appears by the virtue of mathematics, not due to an error in Elephant. This DC offset is not constant and is added only when limiting (dynamic range compression) actually happens on asymmetric peak levels. Note that you won’t be able to remove this offset without introducing new signal level overshoots.

Q. Can I use Elephant also as protection for my monitors and ears? I mean protection against loud sounds when certain other plugs start to scream?

A. Yes, Elephant can be used for this kind of task. For such task, make sure oversampling setting stays at “1x” and the “Knee”, “In Gain”, and “Out Gain” parameters stay at 0: this way no CPU power consuming processing other than peak suppression will be taking place.

Q. K metering – what is it?

A. The “K” signal level metering system was introduced by Bob Katz. You may get more information at <https://www.digido.com/>. Before using this metering system, the sound pressure level of your monitoring environment should be acoustically calibrated according to the specified guidelines. Note that when K metering is in use the meters have fixed 600 ms integration time and 600 ms release time constants.

Q. In bouncing a few projects now using Elephant in the “Clip” mode I've noticed that the peaks are perfectly flat-topped. Doesn't the “Clip” mode soft clip the peak, modulating the samples around 0dB? Shouldn't there be some variation in the waveform there?

A. The “Clip” mode in Elephant performs “hard” clipping. But Elephant also uses a bit of compression in this mode in order to minimize distortion.

Q. Is there a sense to put Elephant on each track to prevent overloading?

A. Generally, no. Elephant is a brick-wall limiter and its main purpose is to limit peaks on the final “limited bit-depth” output, where values above 0 dBFS do not exist, and to perform limiting in a sonically-pleasant manner. Since modern audio host applications have floating-point signal paths on each track there is no need to use limiters on each track as values above 0 dBFS can be passed safely within host’s signal paths.

Q. I am wondering if it is correct to assume that if Elephant is used more lightly on mixes (say, reducing 2-3 db) as opposed to reducing 4-6 db's, there is more “purity” of sound, i.e., the sound is ultimately less distorted – would that be correct to assume? Is this also part of the "Loudness Wars" debate?

A. Of course, it is always best to apply the lowest suitable gain; working without any additional gain is the best variant even. On the other hand, do not overestimate ear response; in my opinion, it gets accustomed to dynamics in a song and loses sensitivity to dynamics: from this point of view loudness maximization is not so “evil”. What's bad in it is distortion and change of timbre it usually produces. This point is exactly a part of the “Loudness Wars” debate.

Q. Just wanted to ask about the release control in Elephant. When in the disabled state, what does that mean exactly, does the release control just change to a default setting? Or is that actually some kind of auto release mode, or is it the equivalent of release knob level being set to a value of 1.00?

A. When the release stage is left disabled, no release stage is used at all, and so dynamics become more instant. This usually produces a considerably different dynamic behavior in the selected limiter algorithm in comparison to the state with the release stage enabled.

Q. Should I be using “8x” oversampling when doing final mixes?

A. This can be beneficial, but “4x” is usually enough.

Q. Could you provide some description on what the “Max” and “Inflate” modes do?

A. The “Max” mode usually provides highest overall loudness, but at the expense of higher distortion. The “Inflate” mode provides a modest, but a clean loudness boost with a large “knee”.

Q. Why do I get negative “Max Gain Reduction” readings even if there's nothing to limit (peak @ -7 dBFS for example)?

A. Please make sure you have set the “Knee” parameter to 0%. In this case the limiter will not be adjusting the input signal below 0 dBFS at all and so the “Max Gain Reduction” reading will not be showing gain reduction happening.

Q. I'm not really into loudness at all, I'm more into vintage music, organic and natural sounds, I wanted to know your opinion about settings for this.

A. For softer music genres using a limiter mode with a higher "Knee" setting could be preferable.

Q. Am I right in that I hear that limiting shrinks stereo-field a bit?

A. Yes, you are right with this observation, it is called the "channel intermodulation". This is easily described mathematically. Widest stereo-field signal can be defined as having zero correlation between left and right audio channels. When stereo-linked signal limiting happens, the audio signal in both channels is multiplied by the transfer function that is based on a maximum of both signals. This in turn increases correlation between channels leading to an impression of a narrower stereo-field. This effect cannot be avoided other than by decreasing the "Channel Linking" parameter in the limiter mode editor.

Q. Whenever I change the "Ch. Link" or "Lookahead" parameters I see overshoot on the host's output level meter. I also seem to get signal level overshoots over 0 dBFS when not using the "Clip" mode. I was expecting Elephant to suppress metering overshoots completely.

A. Overshoots may still be present if the level meter you are using detects the inter-sample overshoots. To fully suppress the inter-sample overshoots you should use at least 4x oversampling in Elephant. Beside that, short-time overshoots when changing some of limiter's parameters are to be expected: this aspect is mentioned in this guide.

Also make sure the master output gain in the audio host application is set to 0.0 dB or lower. If this setting is higher than 0.0 dB, an additional signal gain may be introduced which is beyond control of the plug-in.

Another reason for overshoots is the signal that goes straight to sound card's output without going through the plug-in first.

Q. I'm trying to figure out how to use Elephant in side-chain or mid-side modes.

A. Simply put, Elephant does not feature side-chain nor mid-side processing functions. The routing presets you see dimmed are available to other Voxengo plug-ins having these functions.

Q. Is the Elephant limiter a true peak limiter? I know that it can show true peaks, but does it stop true peaks?

A. The "True peak" statistics is not directly related to the limiter algorithm. You should enable the "4x" oversampling: in this case Elephant will act as a "true peak" limiter.

Happy Mastering!