# endorphin

**dual-band stereo / m/s compressor VST plugin**

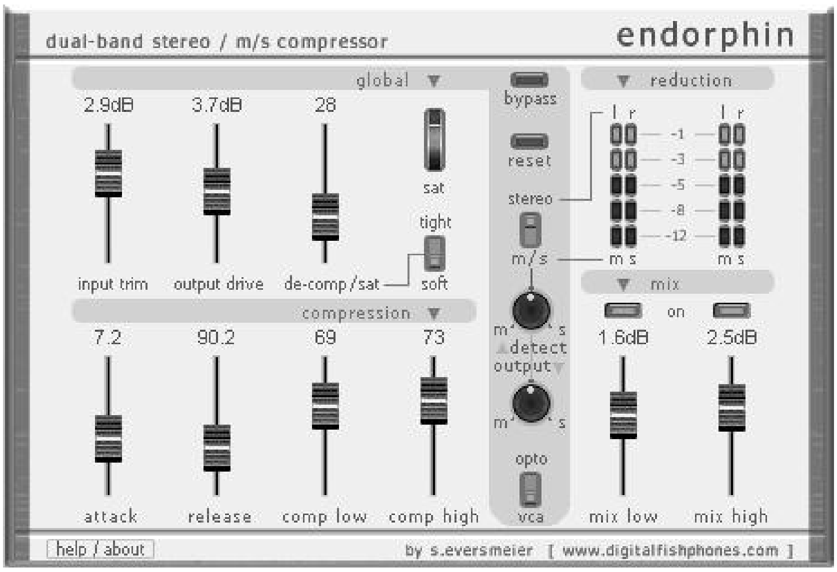


# USER'S MANUAL


Date: 2002-05-06
Author: Sascha Eversmeier, Berlin, Germany
URL: http://www.digitalfishphones.com
Email: sascha@digitalfishphones.com
Subject: endorphin, dual-band stereo / m/s compressor VST plugin, endorphin.dll
Current program version: 1.1


*This manual describes the concepts behind* endorphin *(file: endorphin.dll), its functions and the basic steps on how to use this software.*

Endorphin *is freeware and therefore free of charge. The latest program version is always available at the author's website.*

digitalfishphones.com

**Welcome to** endorphin**.** You have chosen a piece of software that comes from the heart, with its own character and soul. After 9 months of development and careful listening tests, it has reached the state of version 1. Well, let's examine that thing...

### What is endorphin?

Endorphin is a dynamics processor for complex signals, such as entire mixdowns ready for CD burning. It is capable of acting as a mastering tool within a chain of digital software to 'polish' sound material ready for CD duplication. It can be quite subtle with very little effect and coloration or really drastic when going into deep compression. You can make it breath, pump and produce a 'nailed to the wall' sound with a lot of loudness and saturation. But it can also be sonically transparent and add depth without actually sounding very 'compressed'. As usual, it depends on what you put into it and how you treat it.

One of the main intentions of endorphin is to prove that digital tools don't have to sound cold and sterile. The unit uses several methods to mimic the sonic character of real analog equipment. This includes:

- Generation of additional harmonics to the existing sound with increasing operation level caused by soft saturation.
- Analog-style compression circuit which tries to remain sonically transparent wherever possible while creating more harmonics when driven hard. Endorphin was taught to be nice, even when you treat it like a dog ;)
- Interaction. Most internal parameters react to the user's adjustments as well as to the signal itself. There is a lot going on between the stages – nearly everything affects everything else – just like would be the case on an analog circuit board. This has a very subtle effect, but contributes to the overall character.

Using the plugin should be simple and straightforward. The controls and their use might already be familiar to you from real-world compressors.

Apart from just compressing a signal, you can greatly raise the volume of your recordings far beyond what you might have guessed from the digital domain. You can definitely 'kick' endorphin to saturate at high operating levels without introducing any feared 'digital distortion'. You have a lot of control over the saturation and – with a bit of tweaking and careful listening – it will behave pretty similar to a hard-driven analog tape or a vacuum tube amplifier. But you should be aware that this says nothing about signal quality at all! Saturation is an effect. It's a bit like cooking: Your guests might praise the meal you've cooked, but certainly not because of the spice you've added...

## System requirements

Endorphin is a PC-based real-time VST plugin. So that means:

- You will need a PC machine (no Mac version available) with reasonable speed for real-time audio applications. The minimum CPU power should be 233Mhz.
- A VST-compatible software host is required, such as Steinberg Cubase VST, Emagic Logic Audio, Orion from Sonic Syndicate or hosts that are equipped with VST-to-DirectX adapters like Samplitude 6. Endorphin has been tested with the above applications. There may be others which also work, but you will have to find out for yourself.

### Installation

That's pretty easy. As you are reading this manual, I assume that you have already extracted the ZIP archive (thereby using WinZip or a similar application). Its contents are

a) the plugin (endorphin.dll) and
b) the user's manual (yes, which you are currently reading).

To install the plugin file, simply locate the folder named '\vstplugins' of your host program and copy the plugin file right into it. That's it. Now (re-)start the host. It will scan this folder and collect all plugins. When loading is completed, you should find endorphin within the list of available 'insert' plugins.

What are 'inserts'? Well, endorphin is a dynamics tool, so it definitely makes sense to use it to process an entire sum or channel signal. To process only a portion of the signal, as would be the case with a channel's AUX send, doesn't really make sense in most cases. There *may* be applications where a compressor can work on a part of a signal, but that is pretty rare and quite an advanced technique in the hands of very few people. Endorphin was built to be enough of a weapon so let's just skip that…

This plugin comes with its own user interface. You should already be familiar with your host software so that you know how to open plugins and their interfaces/editors.

### Control elements & their function

The following chart shows the signal flow within endorphin:

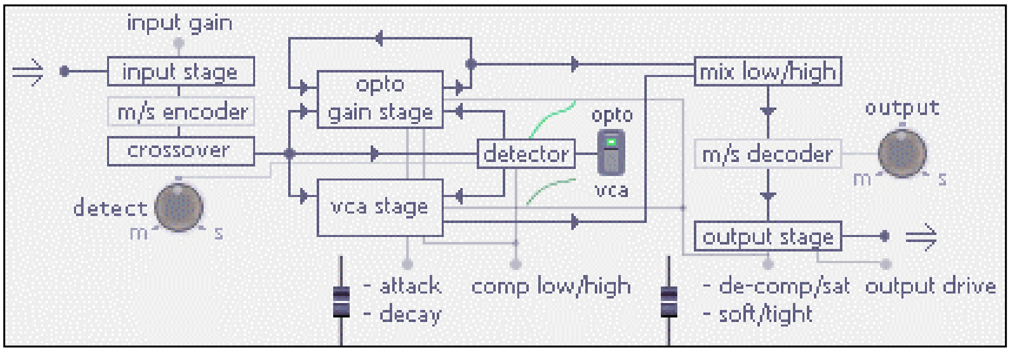


The same chart can be found by clicking the button help / about on the lower left side of the plugin interface.

The different signal stages of endorphin will be introduced on the following pages.

#### Input stage

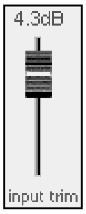


This is the very first stage of the unit. It is just a linear gain device, providing a range of +-6dB. There is nothing special going on here. Normally you wouldn't want to touch this control, unless
a) the incoming signal is either too high or too low to be processed sensibly or
b) you've chosen to feed the compressor's detector circuits with an 'extra-hot' signal. This can sometimes cause nice effects on some material. We'll discuss this later when it comes to handling detector levels.

#### Two-way frequency crossover

This separates the audio stream into low and high frequency components. The crossover is located at a center frequency of 1.2kHz. The filters do not care much about what may be common on the market. Instead, I've designed them to sound good to my ears and support the overall sonic character of endorphin. The crossover actually cuts off very gently. For instance, you'll still hear a lot of the treble region shine through when listening to the low band only.

## The compressor section (low and high band processing)

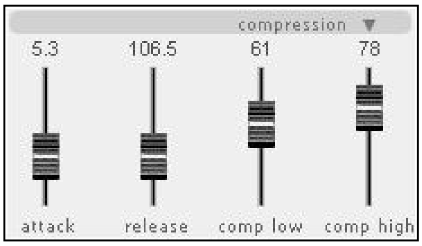


I assume that you are already familiar with basic compressor usage and know about what these devices generally do. Fine! Let's look what is common and what is different here:

The most important fact: This compressor is actually two. One is operating on the lower frequency part, the other on the 'rest', which means mid & high frequency (as stated in the previous section).


The whole compressor section can either operate as a stereo or M/S encoding & decoding device. If set to 'stereo', the detectors take both channels to determine the current signal energy. This mode is active by default because it is the traditional way of compression. To break with traditions, you should read the M/S section of this manual and prepare to dive into some advanced usage. More on M/S processing later on...

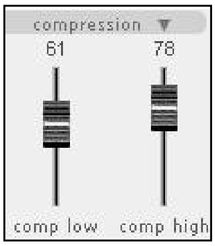


The interface provides separate sliders for the amount of compression of the two frequency bands (**comp low** & **comp high**).
Are you missing the 'usual' threshold & ratio controls? Well, each comp slider includes both parameters. When you move up one of the **comp** sliders, you are actually lowering the threshold of this frequency band (where compression comes into play) and also lowering the compression ratio: a high setting results in more compression.

You might ask 'why no traditional controls?'. This is because
a) endorphin is a *soft-knee* compressor and
b) not supposed to be an academic but a straightforward device instead.

If you've never heard of the term *soft-knee*, you were probably using a *hard-knee* unit. In this case, the ratio you've dialled in is taking place exactly above the threshold point. Below the threshold, the signal stays unaffected. Imagine a transition curve from 'no compression' to 'maximum compression': There would be a sharp edge. Above this point, the ratio tells us about how much we are actually reducing the signal. In traditional design, the ratio will stay constant (that means linear).
Many digital compressors are using a hard knee, since this is easy to implement and takes very little processing power.
Using a hard knee may work fine on material such as drums or anything where it is okay to compress as an *effect*, it is mostly audible to the listener. But if you are aiming at inaudible, transparent compression or working on complex sources like a bus/submix signal or an entire mixdown, things become complicated and a hard-knee compressor can hardly handle such a task.
Endorphin was designed for complex signals and is therefor providing a soft-knee transition from 'uncompressed' to 'fully compressed'. Furthermore, one cannot speak correctly of a dialled-in ratio because here, the process of gain reduction is not linear.

In the **vca** mode, gain reduction increases with rising input energy:

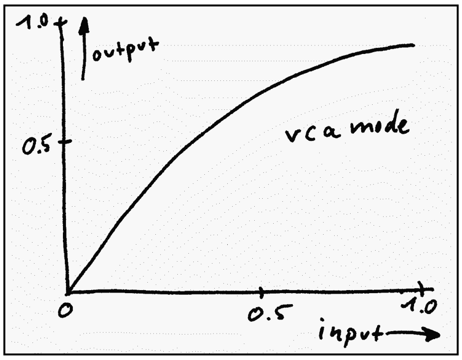


VCA is a design found on most of today's analog dynamics compressors. The VCA (Voltage Controlled Amplifier) is a device that amplifies / scales the audio by the amount of control voltage. This voltage (the measured RMS) is generated by the audio at the input (here: the crossover output connectors). In the real world, the VCA is a solid-state device (many semiconductors or one single IC) which is capable of doing very deep compression (wide dynamic range) and has a quick and accurate response. Because the input signal is directly taken for measuring, this design is called 'feed forward'.
This mode has an instant response to signal changes. If you want to, you can actually go into 'sound design' here and make things smash and pump. I would recommend the **vca** setting for drum loops as well as for dance tracks (but you should find out yourself).

In the **opto** mode, the transition curve looks a bit different. Its behaviour also differs from a hard-knee design, partly because of the curve, partly because of a totally different circuit design:

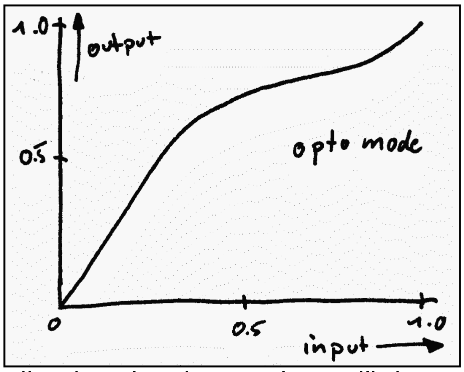


'**opto**': Choosing this mode makes endorphin act like a 'vintage' compressor using an opto-electrical gain control element (a photo resistor attached to a light source). This design was very common with manufacturers of the old 'vacuum tube' age. To limit a signal, one would have chosen a 'variable-mu' (tube) design because sufficient VCA circuits where not available at that time. To achieve an open compression with density and natural character, opto-electrical devices were preferably chosen. They have many sonic advantages, apart from having a soft-knee reactance (and a typical opto-electric transition curve which does less compression at high signal levels), such a device has *memory*: The more you go for compression, the more you will force the opto device to light up. Likewise, the photo resistor will decrease its resistance (which is being used for the gain stage, the 'volume pot', so to speak). But this action is somehow sluggish, the less resistance there is, the more will it take the device to recover. Therefor, a large impulse leads to a long release time, low input levels have a shorter release (in fact our ears do something similar every day). This sounds very 'natural' compared to 'static' time settings and is usually best to achieve transparent compression without audible artefacts.

The time sliders on endorphin still stay intact in the **opto** mode, but they are partly affected by this 'automatic' behaviour. You should use them as a coarse direction.
The 'opto' compressor has a 'feed-back'-circuit. That means, that the output of the gain control element is being used to drive the detector (which would be driven by the input on a feed-forward board). This has a great 'smoothing' or 'stabilising' effect, so it should not confuse you that the compressor controls do not deliver the quick reactance you can expect from the **vca** circuit.

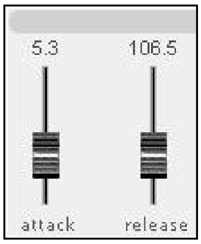


The **attack** & **release** sliders set the response time of both compressor bands (low & high). They are internally scaled to sensible values, the displays above the sliders show averaged values. The reason is that if you use more than a single band in a compressor, it is very important to adapt the time constants of each band, otherwise there would be serious distortion on the low end when you try to use short response settings (like the high band will often need). Remember the wave length...

The detectors were modelled by thinking the 'analog' way: To set a threshold, the input is raised, but non-linearly. There is a saturator on each detector that will smooth out the peaks. The more compression you are aiming at, the more they will get a saturated signal. This greatly smoothes out sudden volume changes at deep compression and is one of the reasons why endorphin is that gentle. In other words, if you crank up the input, the detectors will be blind to the peaks because you have simply distorted them. This might help a lot to achieve some 'analog' feeling though it is not sensible at every time.

**Mix stage**

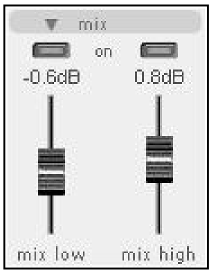


This is pretty unspectacular, as this section is really just mixing both compressor outputs. The sliders have a range of +-6dB. There are applications where you'll need to make some adjustments here, e.g. if you had set up very little compression on the low band, but quite much on the high band. Though both **comp** sliders have an 'auto makeup' feature, things will most certainly need some adaptation here.
The mix section provides two on/off switches that let you quickly monitor what is going on in each of the frequency bands without having to re-adjust the mix faders. Both switches are in the **on** position by default.

**De-comp/sat section**

Well, this stage is actually not isolated by itself (which is why the interface says it's a 'global' setting). It controls the sonic behaviour of the compressor as well as the output stage. Most of the 'analog' modelling in this software is taking place here.

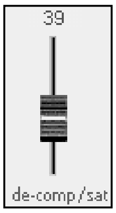


Raising the **de-comp/sat** slider changes the output clipping behaviour from hard clip to soft clip. This introduces early harmonics. The new harmonic structure is partly mapped directly to the output stage (to raise the level without serious audible clipping), another great part affects the compressor: More gain reduction leads to more overall saturation. Because the saturation directly increases the whole output signal, you can overcome the well-known effect of over-compression: It's like expanding it all again, but the altered harmonic structure leads to a fatter sound. What sounded flat and lifeless before now breathes again.

The saturation itself uses an emphasis and de-emphasis stage. If you know how tape-reel machines work (or once played a bit around with a Dolby cassette), you'll certainly know this: Things are being filtered, processed and filtered again reversibly (to get back to the original response).
Creating harmonics with endorphin is working similarly. It is important not to distort the bass frequencies when saturating (this simply sounds awful). But a lot of saturation can be done to mid and high frequency signals (that's generally one of the secrets of 'warmth' perception) instead.
So, when you move up the slider, endorphin knows it has to saturate more right now, so it raises the frequency point where it all should happen. Technically speaking, a high-pass filter is being swept up, using a -6dB low shelf with a corner frequency from 200 up to 600Hz. The saturation is applied to the whole signal, but the bass and lower mid frequencies are less affected because they have been dampened by the filter. This technique prevents intermodulation products that easily appear when you saturate a signal with much of low frequency content.
The de-emphasis stage takes care of reversing the filter process. The theoretical frequency response is flat again, but the actual harmonic structure has changed.

Now, *of which type* are the additional harmonics and *when* and *how much* are created?

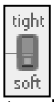


The amount of harmonics will be affected by the **decomp/sat** slider. The spectral structure itself is determined by the setting of the small switch next to the slider. There may be situations where switching between the two states might not introduce huge audible differences while on some other track the effect can instantly be perceived. But as the switch as well as the slider control most of the sound character of endorphin, I'd definitely encourage you to get used to this section, even if things first appear pretty subtle.

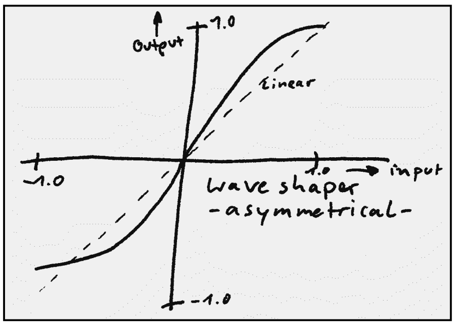


In the **soft** setting, asymmetrical saturation is being applied to the signal. That means the positive branch of the waveform is shaped differently than the negative branch. By doing so, the harmonic spectrum will be enriched with odd & even harmonics. The transformation curve is shaped in a way that saturation will occur quite early. So you don't need very high signal levels to gain additional harmonics here.
The working and the sonic result can be compared to a vacuum tube amplifier. In general, it's not good to saturate complex signals in an asymmetrical way. That's why this mode uses only very little asymmetry. The primary usage should be the 'decompression' effect: By doing so, the saturation is most prominent on the signal's transient. You should avoid aiming at constant high levels with the **soft** setting turned on (thereby radically increasing the **output drive** slider or keeping the **mix** levels set high). Some material will tend to loose definition, though it might sound incredible 'warm' at first.
As a basic rule, you should stop the slider right at the point where you start to perceive some additional 'roundness'. This setting might be the right candidate for tracks which lack just a bit of warmth and don't need to be extremely loud.

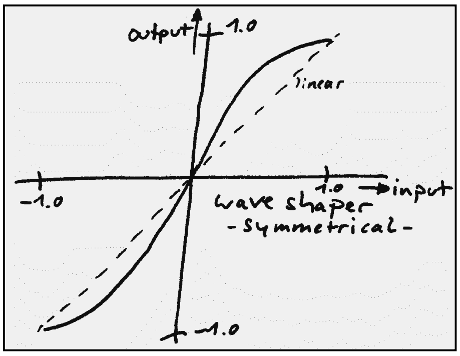


If you'd like to achieve high output levels and want some extra punch from the 'decompression' effect, try the **tight** setting. Here, the saturation takes place much later which means that mostly high signal levels contain extra harmonics. This helps a lot in keeping things clear because it tends to work only on signal peaks (whatever those are after some drastic compression...).
The saturation is strictly symmetrical which creates only odd harmonics. By playing with this setting, you might realise that the thinking 'even harmonics are always better' is a common preconception and very rarely true on complex sources like a mixdown.

The **tight** setting is able to produce much more of the **de-comp** effect as more saturation is generally possible here. The risk of getting a muddy and confused sound at high output levels is not that high as with the **soft** setting. My recommendation is to try this type of saturation on loud & dynamic material.

## Output drive

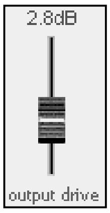


Like already mentioned, the audio runs through a saturator stage before it leaves the (virtual) connectors. The amount of extra gain can be set using this slider. It can add up to 6dB of loudness, which means you can double the volume here. But you should listen closely.

The **drive** algorithm will use the type of saturation you've set up in the **decomp/sat** section.

If you are seriously aiming at a higher output than the initial input level, you should keep an eye on the 'sat' meter which displays the amount of saturation the whole unit introduced. It does not measure distortion that

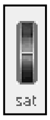


occurred before the plugin but it can be helpful to determine the right amount of decompression, especially when using the **soft** setting. Generally, it is no major problem if the needle occasionally hits the red region. But if it resides here for a longer time or stays like being nailed to the top there's certainly something gone wrong with your music.

While releasing the audio, endorphin needs to take care of clearing up the signal, because larger amounts of saturation always create a DC offset, where asymmetrical saturation introduces the most. In order to ensure a maximum headroom (limiting always relies on a symmetrical waveform), the signal will be DC coupled here by applying a $2^{nd}$-order high-pass filter at 30Hz.

The overall output is terminated by a (fixed) saturation/soft clipping stage which is of symmetrical type and is just to stop the signal peaks that may have slipped through the previous sections. Otherwise, they would seriously clip the next stage in the signal chain from the host application's view.
The output takes care that the signal does not exceed -0.1dBFS. This is the maximum level that endorphin is capable of.
This level is high enough to use the maximum of available headroom while being compatible to older CD players (that might have problems with 0dBFS) and prevents you from cutting tracks the CD plant would certainly reject because of too much clipping with all bits set.

In the case of mastering for CD applications, you should arrange things so that either endorphin is the last unit within your chain of processing or that following stages maintain an unaltered signal level.

## Bypass / reset buttons

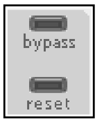


By clicking the **bypass** button you will only hear the unprocessed audio just as if endorphin was not loaded at all. However, your input signal is still being processed in its entirety – all control elements are still functioning. This lets you quickly switch between the 'dry' and processed signal state. I highly recommend that you make regular A/B comparisons to make sure you are still on the right track.

The **reset** button instantly recalls the unit's default position for each of the control elements. This is the same as the *'endorphin init'* preset and is the easiest way to start from scratch.

### Advanced technique: M/S processing mode

So far we have looked at using endorphin in the classical stereo-compressor mode. This mode usually works fine. However, a situation might occur where you'd like to add extra density to a stereo track but a massively pumping rhythm section coming from the center of the stereo image keeps on affecting other instruments and room/ambience information. As an example of a common problem, imagine the drums pulling a synth pad down on every $4^{th}$ beat. With intensive compression it all might sound as if the pad was intentionally 'ducked'. Even with multiband compressors this may be a problem if two completely different instruments are sitting very close together in the frequency spectrum. They compete in level and the compressor's detector only has the overall signal energy to refer to.
This is where M/S processing might help us out.

**Note**: Switching endorphin to M/S mode will result in increase of CPU power consumption. The M/S technique relies on a true stereo signal, so it does not make any sense to feed this mode with a monaural source.

### What's M/S?

Apart from being an ancient and serious microphone technique, M/S is basically about encoding and decoding a stereo signal. The magic lies in between.

To encode a stereo source that consists of left & right channel information, we derive the so-called **sum and difference** signals from the source. To get the sum, one simply 'wires' both channels together, the difference signal is one channel subtracted from the other (it doesn't matter which one).
The sum signal is the **M** portion which stands for *monaural*, *mono* or sometimes *mid signal*. The **S** portion is the source's *stereo information*, or *side signal*.
At the end of the compression path (right before the output stage), endorphin will apply the signal's M/S decoding process. To decode such a signal, one takes the side signal and a copy of it, reverses the copy's polarity (out of phase, 180 degrees), mixes both portions in a 1:1 ratio and attaches the resulting signal to the mono (sum) signal.

### How to use M/S compression with endorphin

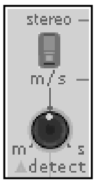


By activating the M/S mode on endorphin, you tell the plugin to derive mid field and side information out of the stereo signal.
You can control the level of signal detection by using the balance knob labelled **detect**. Turning it towards **M** directs both the low & high band detectors to use just the audio coming from the mid field of your stereo track. Likewise, moving the knob towards **s** will result in reacting only to the side signal's energy. But if your source is sensibly correlated, there won't be any need to make huge adjustments here.

The advantage of M/S usage will reveal in a situation like the on described above. When raising the **comp** sliders, take a look at the reduction meters:

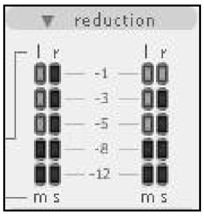


You'll certainly notice that the left bar of each level meter (labelled as **m** at the bottom) shows more current gain reduction than the right bar (labelled **s**). In the example, the mid signal simply contains the most energy. So it is the stereo mid field information that is actually being compressed more.
It's in the nature of M/S processing that in such a case the perceived stereo width image widens, simply because the level in the center temporarily dropped. If you overdo it, it might become a problem, but normally it is not. Very often, the resulting effect is less disturbing than ordinary stereo compression where center instruments are always modulating the rest of the signal.

If you encounter a drastic change in the width of the stereo image by using M/S compression, you can control the balance between the mid and the side signal with the **output** knob of this section:

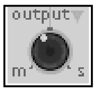


Most times you will have to move the knob slightly in the **m** direction to regain the original stereo width.
But if you'd like to widen up your track, try moving the knob towards the **s** position. Yes, endorphin includes a true stereo width enhancer!
And even more, this enhancer works with saturation (independently on the M & S signal). This can help to support the perception of depth and 'life' in your music, even if you calibrate the knob for an 'unchanged' stereo image. For instance, the saturation can help to raise small portions of the side signal, such as room information or other details.
And as the saturation works more on the mid and high frequency, this results in an slight effect that's similar to tube exciters though the algorithm does never add any additional brilliance to the source.

You don't have to worry about de-correlation and phase problems. M/S is by nature 100% compatible to monaural sources.
This seems obvious because there is no signal delay and therefor no phase shift at all. When summing an M/S-processed track to mono, you will only end up with the mid signal because the 'extra width' has been generated using an out-of-phase signal from either channel. So when you take only the **s** portion, the left & right signals will exactly sum up to 0.
Try this by turning the knob fully to the right and render the track to a single mono file. You will hear... silence.

**A final note:**

- Always trust your ears, don't just rely on previous settings of a similar source or project.
- Do not join the 'I got the hottest CD'-competition. Dynamics compressors can be easily abused to make loud but weak and fatiguing tracks, so take great care. There are times when less is more.
- Press the bypass switch from time to time, just to re-calibrate your ears.
- Don't be fooled by the term 'analog-style'. This is not the analog world. Never ever. Even if we're aiming at it with great effort, true analog circuits still sound different and keep a certain magic. This is not a story of good or bad, it is the question of 'What do I want to achieve? What is my way of working? What is the weakest thing in my chain?'. We shouldn't forget that.
- Give me some feedback on endorphin, of whatever kind. In the end, it all helps to improve this software and influence my future developments.

- **Have fun**.

*Sascha Eversmeier*
Berlin, May 6, 2002
http://www.digitalfishphones.com



digitalfishphones.com

**PLEASE DO NOT HOST THIS SOFTWARE ON THE NET FOR YOURSELF. ALWAYS POINT A LINK TO MY PAGE (http://www.digitalfishphones.com) TO MAINTAIN UPDATED PROGRAM VERSIONS!!! Thank you ;)**

**LICENSE TERMS FOR ENDORPHIN DUAL-BAND STEREO COMPRESSOR PLUGIN SOFTWARE (endorphin.dll, in the following: 'software')**

The software is provided to the user 'as is'.

I make no warranties, either express or implied, with respect to the software and associated materials provided to the user, including but not limited to any warranty of fitness for a particular purpose. This software was purely written to meet my requirements, I do not warrant that the functions contained will meet yours, nor that the operation of it will be uninterrupted or error-free, or that defects in the software will be corrected.

I do not warrant or make any representations regarding the use or the results of the use of the software or any documentation provided therewith in terms of their correctness, accuracy, reliability, or otherwise. No information or advice given by me shall create a warranty or in any way increase the scope of this warranty.

I am not liable for any claims or damages whatsoever, including property damage, personal injury, intellectual property infringement, loss of profits, or interruption of business, or for any special, consequential or incidental damages, however caused.

At present this software is still under development and not to be a public release. It is freely given to people to test its functionality and allow feedback to me as the author for future development, though I do not warrant that there will be such.
The user is not allowed to distribute the program without my permission. Further, the user may not modify or may not decompile nor debug the endorphin software.

**THIS PROGRAM IS FREEWARE.**

digitalfishphones.com