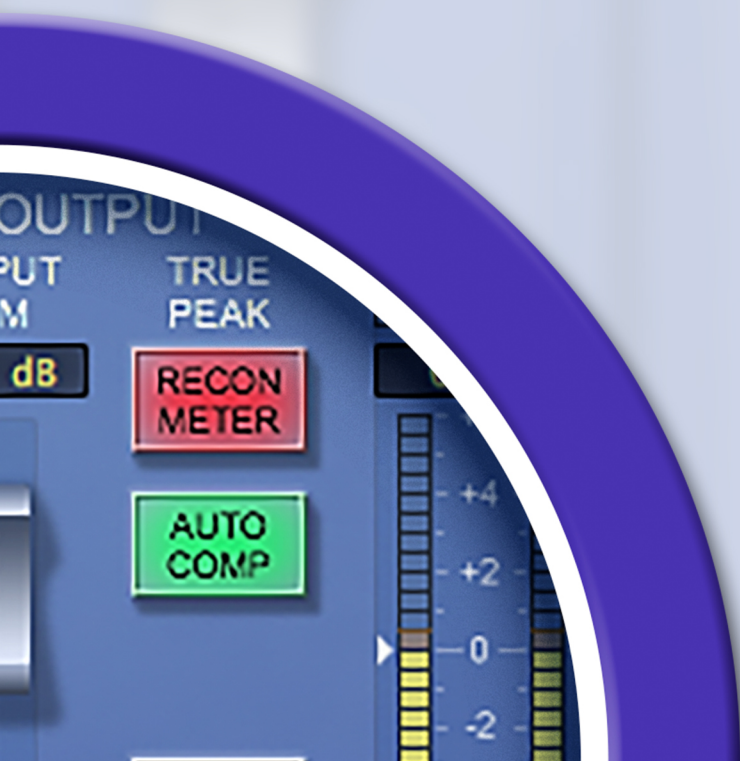




# Oxford Limiter

User Guide



# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Overview</b>	<b>6</b>
<b>3</b>	<b>Pre-Process Section</b>	<b>7</b>
3.1	Threshold . . . . .	8
<b>4</b>	<b>Enhancement Section</b>	<b>9</b>
<b>5</b>	<b>Attack Timing</b>	<b>11</b>
5.1	Using Pre-Process without the Enhance Section . . . . .	12
<b>6</b>	<b>Release Timing</b>	<b>14</b>
<b>7</b>	<b>Programme Limiting Procedures</b>	<b>16</b>
7.1	Loudness Maximisation . . . . .	16
7.2	General Gain Management . . . . .	17
<b>8</b>	<b>True Peak Reconstruction Metering</b>	<b>18</b>
8.1	Meter Operation . . . . .	19
8.2	Manual Correction . . . . .	20
8.3	Automatic Correction ('inter-sample peak' limiting) . . . . .	20
8.4	Legacy Dac Simulation Mode . . . . .	21
<b>9</b>	<b>Dither and Noise Shaping</b>	<b>22</b>
9.1	Conventional Dither . . . . .	22
9.2	Noise Shaping Dither . . . . .	24
9.3	Noise Shaping Depth Control . . . . .	27
9.4	Additional Information . . . . .	29
9.5	Output Meter Clip LEDs . . . . .	30
<b>10</b>	<b>Description of Controls</b>	<b>31</b>
10.1	Input Section . . . . .	31
10.2	Pre-Process Section . . . . .	32
10.3	Output Section . . . . .	32
10.4	Options Menu . . . . .	33

<b>11 Suggested Workflows</b>	<b>35</b>
11.1 Maximum Transparency . . . . .	35
11.2 Maximum Loudness . . . . .	36
11.3 ‘Brick-Wall’ Limiting . . . . .	36
11.4 Points to remember . . . . .	37
11.4.1 Limiting at high sample rates . . . . .	37
<b>12 Specifications</b>	<b>38</b>
12.1 Latency . . . . .	38
12.2 Pro Tools   HDX – Instances per chip . . . . .	38
<b>13 Preset Manager Toolbar</b>	<b>38</b>
<b>14 Supported Platforms</b>	<b>40</b>
<b>15 System Requirements</b>	<b>41</b>
<b>16 Copyright and Acknowledgements</b>	<b>42</b>

# 1 Introduction



The Oxford Limiter plug-in has been developed from decades of professional audio experience to provide a very high degree of quality and capability in programme loudness control and limiting functions. By employing highly accurate logarithmic side-chain processing, along with innovative adaptive timing functionality using look-ahead signal acquisition, the Limiter offers exemplary performance, whether one is seeking general transparent level control, programme loudness maximisation or heavily applied artistic audio effects.

Unique processing in the form of the Enhance function provides the sample value limiting needed to reliably avoid overloads in digital workstation environments, and allows unprecedented volume and punch to be applied to programme beyond that available from conventional limiting functions.

Comprehensive metering is included, which displays not only conventional peak sample value, but additionally allows the user to monitor the true validity of the programme in order to avoid the generation of damaging reconstruction overloads in the target equipment, which are often invisible during production (sometimes termed ‘inter sample

peaks’).

A further function allows the user to dynamically correct for reconstruction overloads in real time, thereby achieving maximum possible modulation levels without the risks of producing ‘illegal’ signals often associated with compression and limiting.

Comprehensive dithering functionality is included, with selectable and variable depth noise shaping, which ensures first class mastering output quality in either 24 bit or 16 bit modes.

## 2 Overview

The Limiter plug-in is comprised of four processes, cascaded in the following order:

- Pre-processing gain control
- Programme enhancement and overshoot control
- Reconstruction metering and compensation
- Dithering and noise shaping

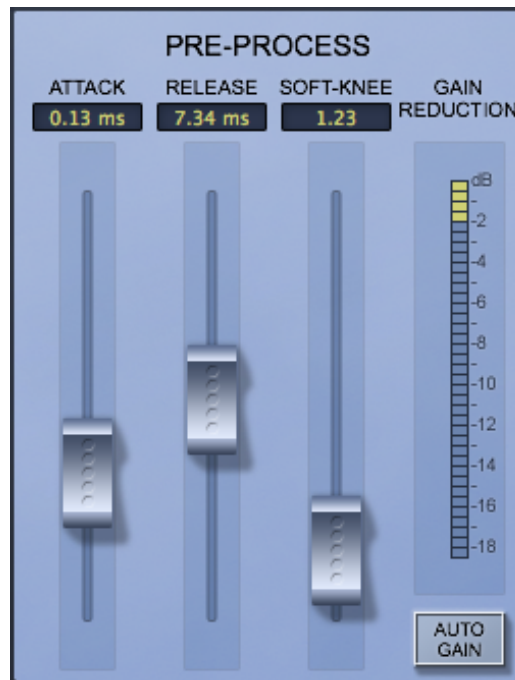
The primary purpose of the plug-in is to control peak levels and increase the volume, density and presence of musical programme **without** an excessive loss of transient and dynamic information that normally results from conventional peak limiting applications.

In order to achieve this, the application employs gain scaling and compression in the Pre-Process section, and peak overshoot control in the enhancement section. These processes are used in conjunction to enable the sound of transient and dynamic information within your programme material to reach the output of the plug-in, despite very tight control of maximum peak sample values.

The normal **ENHANCE** control setting for this action to fully occur is 100% and above in **Normal Mode**, (or any enhance setting with **SAFE MODE** selected). A variable control allows the enhance process to be adjusted or removed completely (i.e. at 0% setting), in which case the Pre-Process section may be used as a conventional programme levelling application if desired.

A **SAFE MODE** is provided which uses the enhance processing to permanently control peak levels. In this mode the **ENHANCE** control varies the perceived loudness boost of the programme by modifying the processing law.

### 3 Pre-Process Section



The Pre-Process section provides a musical programme levelling function. Its primary purpose is to control programme level over a wide range, in order to create optimum conditions for the following enhancement stage. When the **ENHANCE** control is disabled in Normal Mode (at 0% with **SAFE MODE** disabled), the Pre-Process section can be used as a conventional levelling section in its own right.

The processing threshold is set to 0 dBFS by default. Dynamic gain reduction can be achieved by increasing the **INPUT GAIN** sufficiently for the internal signal to surpass the 0 dBFS reference level. A total of 18 dB gain boost is available for this purpose, and the orange section of the **INPUT** meter illustrates the level range within which gain reduction occurs when the plug-in is active. The final output modulation level is set by the **OUTPUT LEVEL** control, which can be adjusted to compensate for dynamic conditions produced by the programme and Limiter settings, or produce a lower level mastered output if required.

Programme gain is accurately controlled by look-ahead detection (to allow action before any peaks are encountered) coupled with a logarithmic side-chain employing multiple interdependent timing functions. Timing controls are provided to modify its action

depending on the programme type and production style. In order to reduce excessive short-term gain modulation, a selectable **AUTO GAIN** function is included, which compensates for wider input level variations by imposing a longer-term time constant that underlies the peak timing. A variable progressive **SOFT KNEE** function allows varying degrees of soft limiting (for lower settings), right up to large-scale gain management active over the final 10 dB of programme dynamic range.

### 3.1 Threshold

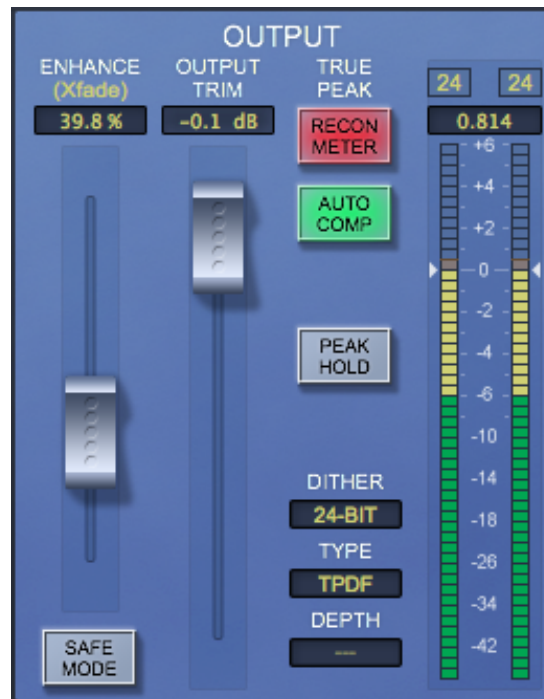
The **THRESHOLD** parameter adjusts the processing threshold of the Pre-Process and Auto Comp processing, and the **ceiling** of the Enhance/Safe Mode processing.

Sometimes it may be necessary to limit programme to less than 0 dBFS. In this case, simply set the **THRESHOLD** parameter below the default 0 dBFS setting.

One scenario in which this may be necessary is when the programme must be limited to less than 0 dBTP (dB 'True Peak') in order to comply with programme broadcast or delivery requirements. See True Peak Reconstruction Metering for details.



## 4 Enhancement Section



The purpose of the enhancement process is to provide sample value limiting and overall programme loudness improvement. The process follows the Pre-Process section in the signal path, and is controlled by a separate **ENHANCE** fader from 0% (no action) to 125% (maximum action). In Normal Mode, the range from 0% to 100% fades in the effect to full level, at which complete sample value limiting occurs. Settings from 100% to 125% further modify the process to progressively increase loudness and programme density at the expense of increasing potential distortion artefacts.

**SAFE MODE** allows absolute peak level control without excessive enhancer action, even when using slow attack settings. In Safe Mode the enhance process is set to run permanently, and the enhance fader modifies the action of the process (rather than the proportion of the effect). Setting ranges from 0% to 100% control the degree of programme loudness boost generated by the enhancer. The control region from 100% to 125% works identically to Normal Mode. It should be noted that in Safe Mode signals at all levels are being processed permanently, therefore some minor changes to the programme dynamics can occur even for a minimum setting of 0%.

The enhance process improves the perceived loudness and presence of the programme material by modifying the dynamic and harmonic content of the signal. Since the method used is different from the Pre-Processing section, it can further enhance the perceived volume of a previously processed signal, while suppressing all signal overloads. As the limiting action does not involve conventional sample value clipping, harsh distortions are avoided, and programme detail and dynamic information is largely retained.

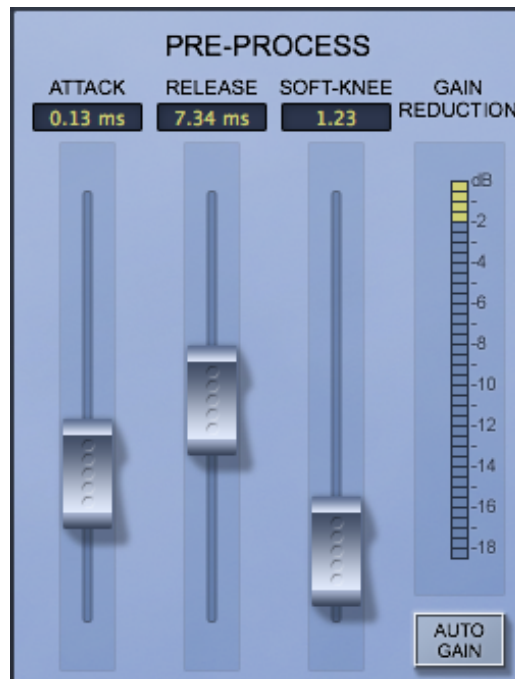
Also, since the plug-in has internal headroom, transient levels greater than notional maximum modulation can pass from the Limiter section into the enhancement stage. This means that percussive overshoots, that would normally be lost in a conventional limiter (or would give rise to overloads), may be included within the sonic results of the plug-in, producing both richer dynamic sonic detail and a useful reduction in the perceived artefacts of the limiting process, all without giving rise to any sample value overload.

*This property enables slower attack times to be used in the Pre-Process section, without creating output overloads, which would otherwise result in a need to reduce output levels.*

The enhancement section can be used effectively on its own to produce programme enrichment and peak value limiting, by using minimal gain reduction and slow timing settings in the Pre-Process section. Or it can be used to enhance highly processed content from the Pre-Process section to achieve even greater perceived loudness.

Because the enhancement process adds harmonic distortion during dynamics within the programme, under some conditions side effects may occur depending on the content of the programme material. Generally speaking side effects should be minimal when in Safe Mode and for boost settings up to 100% when mastering in the presence of most commonly occurring complex and dense composite material. However, some extra care may be needed in the case of single solo instrument tracks where there may be a predominance of sustained lower and middle frequency content. Settings above 100% are most useful when the programme type is intended to be very loud, or where extra distortion may actually prove advantageous within the style of the production.

## 5 Attack Timing



The addition of an **ATTACK** timing control is a significant departure from conventional limiter applications, and requires some explanation for the best results.

Because the level detection within the plug-in looks ahead of the gain control, peaks in the programme material are acted upon in advance of the gain reduction process. Therefore, at the fastest setting of the attack control, programme peaks are controlled within a very small margin (less than +0.25 dB, with respect to continuous sine input conditions).

The **ATTACK** fader allows the attack time to be increased to achieve a favourable improvement in the sonic qualities of the peak reduction process, by allowing peak programme transient events to escape hard gain reduction. Since the plug-in has internal headroom, these overshoot peaks are retained and **are not clipped**.

Peak overshoots resulting from a combination of the programme material and the action of the Pre-Process stage are then passed to the enhancement section where their sonic signatures can be added to the final programme sound. Providing **SAFE MODE** is selected, or the **ENHANCE** fader is set to 100% or more in Normal Mode, no output

sample value overloads will occur from the plug-in, regardless of **ATTACK** or **RELEASE** time settings.

A combination of slower attack times and the enhancement process is therefore a very powerful way to include transients in the output programme that would normally be removed by conventional limiting processes. It can create a sonic quality and impression of dynamic range that belies the degree to which the programme is actually being limited.

As an example of the difference this can make, if slower attack settings are used without the enhancer, and the output gain is reduced to accommodate the overshoots (avoiding overloads), using the same attack settings with the enhancer can create up to (and beyond) a 3 dB increase in average level, being able to legally increase the output gain setting and around another 2 dB of perceived loudness due to the enhance action itself.

***This can result in a perceived loudness increase of 5 dB to 6 dB!***

In general, very fast attack times will more readily remove extremely fine detail and short-term events, but will produce greater harmonic disturbance.

Slower attack times will progressively allow finer detail to escape the harsh sound of fast limiting, and longer term events will tend to assume a more rounded peak profile. Such settings are usually kinder to the musical programme.

The only way slower attack settings can be used without a potential need for a significant reduction of the output level control is when the ENHANCE control is set at, or above, 100% or with SAFE MODE selected. These methods will 'soft clip' transients, i.e. compress their level before the onset of clipping, providing a warmer, valve-like saturated sound. Alternatively, enable 16 bit or 24 bit dither, which will 'hard clip' over-unity transients, ie. squaring the top/bottom of waveforms and resulting in a more distorted sound, often found in transistorised circuits.


## 5.1 Using Pre-Process without the Enhance Section

You should be aware that, since the Pre-Process section is a programme gain controller rather than a simple sample clipper, programme peaks can cause a small increase in maximum output sample value, even at the fastest attack time settings. If SAFE MODE isn't selected, or the ENHANCE fader is not set at, or above, 100% in Normal Mode, these peaks will appear at the output of the plug-in. Increasing the attack times will

further increase peak overshoots, so if tight level control is required without the enhancer, it is best to leave the attack at minimum setting.

Since the plug-in has internal level headroom, the output level control can be safely used to compensate for any artistically intended overshoot without fear of causing internal signal clipping.

## 6 Release Timing



The **RELEASE** control has a very wide range to accommodate the maximum possible extent of programme and production style. The ability to set very fast release times specifically allows for the modelling of short-term peaks over restricted gain reduction ranges, up to a maximum of around 4 dB. Such settings will result in high levels of distortion for larger gain reduction ranges, and are therefore unsuitable for overall level control situations.

The **AUTO GAIN** function can be used very effectively to compensate for large level changes, while still allowing fast peak modelling for shorter peak events, as the **ATTACK** and **RELEASE** controls are permanently functional. So it is a good idea to keep **AUTO GAIN** selected under most circumstances.

Generally speaking, faster release times produce the greatest perceived loudness, since gain recovery happens quickly after peak events have passed, and average programme levels are affected only during the shortest possible periods. However, since the gain recovery can begin to occur between the waveform peaks of lower frequencies in the programme, there is a trade off to be made between the speed of release and the generation of distortion. Such distortion may be desirable under many conditions,

particularly in loud popular music productions, where some low frequency harmonics may add warmth and presence to the programme. Adjusting release timing over a wide range provides a method to ‘tune’ these effects to suit the production style.

Longer release times are far more forgiving of gain changes, and allow greater overall compression, but will result in a quieter sounding output programme.

If the **AUTO GAIN** function is not selected moderate release time settings (above around 0.2 seconds) may produce audible gain ‘pumping’ due to longer and more noticeable recovery periods. If such settings are needed the aim becomes one of ‘fitting’ the release timing to the natural rhythm of events in the programme. Under these conditions better results may be achieved by increasing the input gain somewhat thereby compressing further and applying the Soft Knee function in order to compress gently over an increased portion of the dynamic range. In this way the transition in and out of compression will become gentler and less obvious.

## 7 Programme Limiting Procedures

There are many approaches to limiting within current productions trends, but most approaches fall into two categories: loudness maximisation and general gain control. A very wide range of control is provided by the Oxford plug-in to make both these situations possible with ease.

The key to successful limiting is to understand that we are much more sensitive to the **rate of change of gain** than we are to absolute level. Therefore successful limiting has a tendency to fall into an appropriate mixture of two simultaneous but conceptually separate actions:

### Fast control over small level ranges


- because they are too quick for us to notice and too small to produce damaging harmonic distortion.

### Slow control over larger level ranges

- because the gain changes are slow enough to escape obvious notice and the rate of change of level is slow enough to avoid intrusive modulation effects and distortion.


### 7.1 Loudness Maximisation

The aim of this procedure is to achieve an overall average increase in the level of the programme by reducing the size of short-term peaks, and applying extra gain to move the programme up into the extra range freed up by the removal of the peaks.



Signal before limiting





Signal after limiting

To achieve this, it is customary to select relatively fast attack and release times whilst judiciously increasing the input levels so that only the offending programme peaks are subject to reduction by the Limiter, and the average modulation level is increased.

The Sonnox Oxford Limiter can produce significantly superior results in loudness maximisation because it can fully limit the signal even when using slower attack times.

This leads to much lower distortion and less removal of dynamic programme information.

The timing controls on the Pre-Process section of the Limiter can be used freely to make subtle modification to this process, in order to achieve the best possible results.

## 7.2 General Gain Management

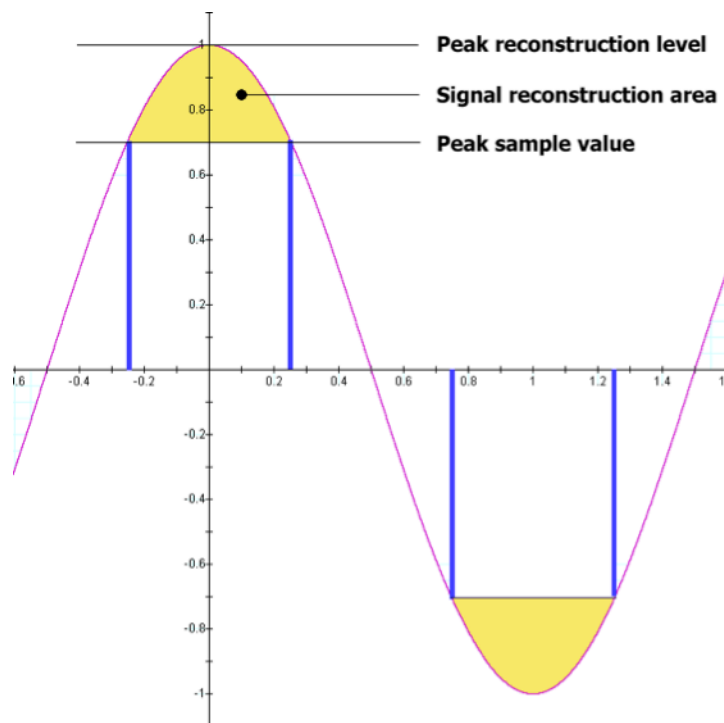
The aim of this procedure is usually to preserve the short-term dynamics of the programme as far as possible, whilst ensuring that no levels surpass maximum peak modulation. This most often entails responding to the peak level of the programme as quickly as possible, and re-scaling the gain in the longer term, such that musical dynamics are only minimally affected in the short term. The Limiter can excel in general gain management because of its wide range of control in the Pre-Process section, and the ability of the Enhance stage to control the level of short term peaks, which means that musically kinder attack and release times can be used without risk of transient overloads. Moderate attack times and slower release times usually perform best for this function.

The **AUTO GAIN** function is particularly useful in this case, as it provides a method to achieve long-term gain control whilst allowing a degree of fast gain riding over a reduced range. Moderately fast recovery times can be arranged to control short-term events with the auto gain function managing long-term level changes. Higher **SOFT KNEE** settings can also considerably improve perceived quality, as this acts in the same way as a variable ratio compressor that starts at lower levels. This allows the Limiter to preview signals at moderate levels, and reduce the rate of change of gain in the loudest peak regions.

## 8 True Peak Reconstruction Metering

An important fact, which is often overlooked, is that in any discrete time sampling system, it is possible to create sample values that may not be decoded and reconstructed correctly. Whilst it is true that, **if left unchanged**, a signal properly converted into the digital domain by a perfect ADC will always produce sample values that can be legally decoded in a perfect DAC, further processing of those samples can result in a decoded signal that is illegally high, and therefore may not be faithfully reproduced, even if no sample value limiting occurs.

The following is an illustration of one example of this situation.



With a signal frequency that is exactly 1/4 of the sample rate, and in phase (i.e. 12KHz at 48K or 11.25KHz at 44.1K) it is possible to generate full level output from a DAC from samples that barely reach 70% of full value within the digital domain.


If the level of this signal were increased towards maximum allowable sample value (close to  $-/+ 1$ ) the area representing the DAC reconstruction filter would produce signals that are considerably in excess of maximum modulation.

Since the vast majority of metering within workstation environments responds to sample value only, the above example would show a level of around -3 dBFS. However, any further increase in the level of the signal would result in a potentially illegal output level from the DAC. As this error would not be reported on metering within the workstation, in this particular case a possible 3 dB overload can result if the signal is increased to a maximum reading on the workstation meters. This phenomenon is sometimes referred to as ‘inter-sample peaking’.

Although the above example is somewhat extreme and specific, there is plenty of potential for this to occur within the mixing environment. Combining a number of processed contributing tracks, and limiting the result to the maximum possible modulation level in order to satisfy current industry trends, when using only peak value detection and metering, is a recipe for such hidden errors.

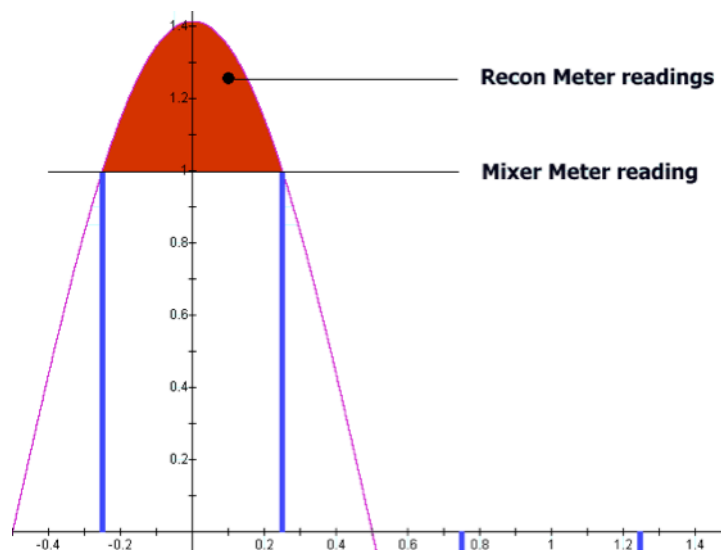
Since the very purpose of a limiting application is most often to increase average modulation levels, the Oxford Limiter provides a True Peak reconstruction meter with automatic correction processing, allowing the user to avoid or repair such errors.

## 8.1 Meter Operation



Reconstruction errors

When **RECON METER** is enabled, the Output Meter is switched from conventional peak sample value mode into ITU-R BS. 1770 compliant True Peak mode. In this mode ‘true peak’ reconstruction levels will be displayed on the meter. Levels in the red overload range of the meter represent the presence of potential reconstruction errors, as illustrated below using the previous example:



Two methods are provided to correct for these reconstruction errors.


## 8.2 Manual Correction

Since the output level fader precedes the metering, errors may be corrected manually by simply reducing the output level setting by the same amount as the maximum error level reported on the meter.

## 8.3 Automatic Correction ('inter-sample peak' limiting)

Under normal circumstances, errors that are often restricted to certain specific events are interspersed throughout the programme. The **AUTO COMP** function is provided to address the situation where it may be undesirable to reduce the level of the whole programme to avoid transient errors.

When the **AUTO COMP** button is selected, the level of the output is automatically controlled to repair reconstruction errors by the minimum amount required, and only for the duration of the error. In this way the loudness of parts of the programme unaffected by the errors remains as high as possible.



Corrected reconstruction errors

When **AUTO COMP** is enabled the corrected reconstruction errors are displayed in the output meter, greyed out to denote that they no longer present a potential error condition.

A combination of **AUTO COMP** and manual output level reduction can be used to strike a compromise, if the action of the error correction becomes intrusive in the presence of very large and intermittent error conditions.

To limit the output Programme to less than 0 dBTP, simply decrease the **THRESHOLD** parameter in the Input panel. Input Gain fader reduction may be necessary if sample value limiting is not required.

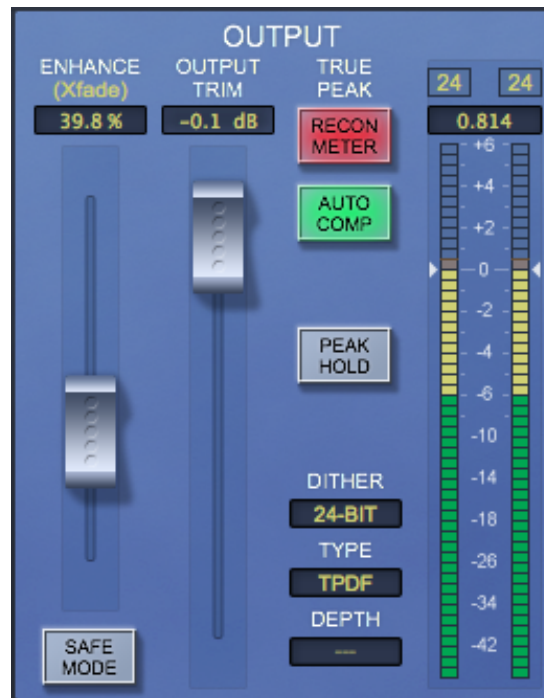
## 8.4 Legacy Dac Simulation Mode

The Oxford Limiter now defaults to using the new ITU-R BS. 1770 'True Peak' compliant mode. The previous 'Recon Meter' mode has been deprecated and is now available from the Sonnox options menu. It is selected by default for compatibility with old projects/sessions.

Legacy DAC Simulation Mode does not comply with current 'true peak' measurement standards, and will under-report reconstructed signal levels. It is therefore not recommended to use this mode except for compatibility with previous projects/sessions.

If required, the legacy mode may be enabled manually from the Sonnox options menu.

## 9 Dither and Noise Shaping



The finite mathematical precision provided by digital audio systems and the effects of dither have been a source of confusion in the audio community for some years. Such discussion may lead to possible misconceptions, which may prevent the user from achieving maximum performance from the systems in use. Therefore the dithering options included in the Oxford Limiter warrant some prior explanation.


The **DITHER** button (see right) toggles dither selection between **NONE**, **16 BIT** or **24 BIT** word length.

### 9.1 Conventional Dither


In both 24 bit and 16 bit output word length selections, high pass **TPDF** (Triangular Probability Density Function) dithering is applied to the output of the plug-in. Since any signal related error caused by finite word length limitation is turned into constant random noise with no relation to the signal itself, such dithering provides complete removal of harmonic distortion due to precision limits, which are an inescapable result of any numerical signal representation. Dithering also suppresses any possibility that the

programme will suffer loss of harmonic signal resolution due to word length restriction. The following graphs illustrate this in action.

The first graph shows the damaged spectral result of passing a low-level 1KHz sine wave through a 16 bit undithered truncation:




The next graph shows the exact same signal and truncation to 16 bits, but with the **HP TPDF** dither applied:



As it can be seen, all the harmonic errors have been removed. Also, since the FFT analysis method provides an enhanced view of the signal below the noise floor, it can also be seen that there is effectively no low level floor below which a signal will fail to pass.

The graph below illustrates this fact by showing a 1KHz signal at  $-120$  dBFS passing through a dithered 16 bit system. This corresponds to a signal 24 dB below the level of the least significant bit; the effective channel SNR (Signal to Noise Ratio) is added in blue for illustration purposes.




This illustrates that dither turns a quantised numerical signal conduit into the equivalent of a naturally continuous (un-quantised) system, which exhibits a finite signal to noise ratio with **no practical limit** to harmonic signal resolution. In other words the inescapable presence of quantisation in numerical systems **does not** forcibly lead to ‘discontinuity’ or ‘resolution loss’ in the signal. The misunderstanding of this fact underpins many of the most damaging misconceptions surrounding digital audio systems. It can also be deduced from the above graphs that any undithered digital representation of an audio signal is effectively illegal.

## 9.2 Noise Shaping Dither

If for some reason SNR figures of 93 dB at 16 bits (or 143 dB at 24 bits) prove insufficient, noise shaping can create an apparent increase in SNR, but there are some potentially hidden costs. Noise shaped dithering is a mechanism that aims to reduce the perceived loudness of the noise of a dithered signal by either forcing the spectrum of the noise out of the audible range or placing it into frequency ranges to which we are less sensitive. In this way the noise at very low levels may be reduced and even lost entirely, if they are at the limit of our hearing within ambient noise conditions. The following graph illustrates this process.





The red line shows the original 16 bit dithered output with the  $-120$  dBr signal passing. The blue line shows the effect of noise shaping (**TYPE 1** at 100%) on the same signal transfer.




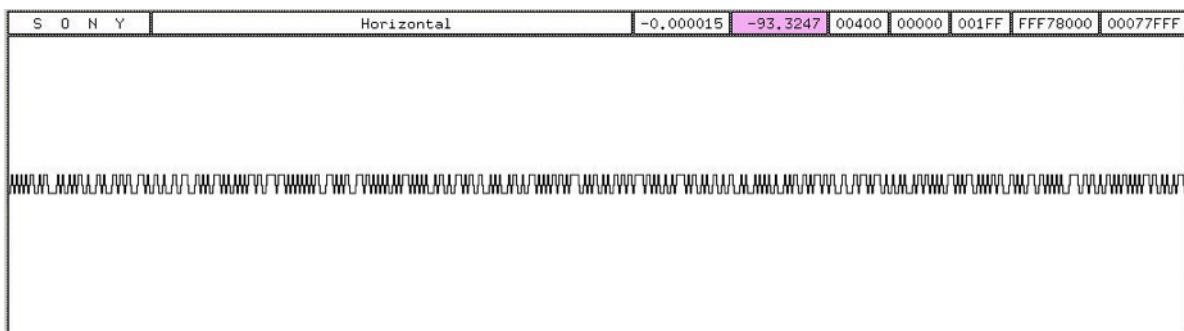
In use, the **DITHER TYPE** button (see right) toggles **TPDF**, **TYPE 1**, **TYPE 2**, **TYPE 3**, and **TYPE 4**.

It can be seen that the noise has been substantially reduced in the regions up to around 8KHz where we are most sensitive, at the expense of extra noise energy in the higher

ranges above 10KHz at which we are less sensitive. Such processing can have a dramatic effect on the perceived intrusion of low-level noise.

However, as is always the case, one cannot get something for nothing and it can be seen from the above graph that the total noise power across the whole range must remain constant to satisfy the dithering requirement.

This means the noise level necessarily increases in some ranges of the spectrum. A level increase anywhere in the spectrum must be accommodated by an *increase* in total peak noise level. The following graphs illustrate this in action, under the previous test conditions.




The first graph above shows a sample value plot of the conventional **TPDF** dithered signal. The second graph shows the increase in values caused by the application of the **TYPE 1** noise shaping at 100%. Focussing the dither energy into a more restricted range than would naturally occur causes the level to increase.

From the effective level values highlighted in the graphs it can be seen that the application of noise shaping has increased the effective noise level from around  $-93$  dBr to  $-80$  dBr, an increase of roughly 13 dB.

From this it can be understood that the design of a suitable noise shaping frequency curve is a trade off between the perceived loudness of the noise under certain conditions and the increase in overall level of the dither signal; much of this trade off relies on what we can hear (psycho-acoustics). Significant research has been carried out over the years into various approaches to this issue, and several accepted curves are in use around the industry.

The Oxford Limiter includes four noise-shaping curves: Types 1 and 3 are fifth order and Types 2 and 4 are third order designs, representing a varied set of trade-offs to suit most programme types, as illustrated below:



Whilst it is understood that the selection of noise shaping type is largely a matter of user preference, generally speaking Types 1 and 2 produce the most dramatic reduction in overall noise loudness, with Type 1 being the most effective of all. Types 3 and 4 provide gentler responses, which under some circumstances may produce less intrusive sounding spectrums, at the expense of higher audible residual noise. Type 3 also provides greater noise attenuation in the range between 10KHz and 16KHz, at the expense of higher noise levels in the mid ranges.

### 9.3 Noise Shaping Depth Control


From the previous section it can be seen that noise shaping can potentially cause unwanted effects in equipment and processes down line, particularly if the programme is to be further modified, such as in mastering situations. Some unwanted effects may include:

- Marked increase in noise levels if the file is not transferred intact bit for bit, (ie. if further processing is implemented)
- Premature meter readings in silence
- Premature peak level overloads (as increased dither levels add to peak signal value)
- Unwanted low-level behaviour in dynamics processing
- Disturbance caused to lossy encoding processes such as MP3, AAC, WMA etc.
- Increased audibility (unmasking) of various errors that may occur in play-out systems

Generally speaking, high levels of noise shaping render a signal that is more fragile. Almost any change to the produced audio file after noise shaping could potentially result in unwanted effects.

For these reasons the **DEPTH** selector is provided to put you in charge of the degree to which noise shaping is applied. When any of the noise shape curves are selected, the depth selector varies the degree of noise shaping from 0% to 100% in 10% steps. At 0% the dither is conventional HP TPDF dither (as if noise shaping were not selected), at 100% full noise shaping is applied. All control positions within the range produce legal proportions of dither.

The action of the depth control is illustrated below with **TYPE 1** selected:



There is no technical (or philosophical) advantage to noise shaping above and beyond that which can be actually heard directly. Therefore the decision to use noise shaping

(and to what extent) is basically determined by what might actually be heard in practice. If conventional TPDF dither provides sufficient audible dynamic range such that noise never intrudes within the programme material, it is safest to avoid noise-shaping altogether.

Because of the potential fragility of a noise shaped signal, it is better to ensure that it is carried out only at the *final stages* of mastering, immediately prior to release. If your mix is not already a final master, it is technically preferable to send a TPDF dithered 24 bit file to the mastering facility rather than a 16 bit noise shaped file.

Another important factor is that the effectiveness of psycho-acoustic noise shaping relies heavily on our sensitivity to noise spectra at the *threshold of hearing*. Therefore, if noise shaped dither actually gets to be heard directly, it will often sound quite strange and intrusive and may detract from the listener's experience of the programme. The most effective and safest approach is to apply noise shaping at the minimum amount necessary to render the noise inaudible within the conditions the programme is destined to be auditioned.

## 9.4 Additional Information

Generally speaking, a signal must be dithered if it is to be output via any form of reduced bit-depth signal (for example, you are listening to it via a DAC) or to any limited resolution media. The only time you can avoid dithering your material is if you know for certain that dithering will be applied later (for example your material will be sent to a mastering studio).

The Dither section, if enabled, must 'hard clip' your signal if it has not already been controlled to legal sample values. This is because dither must be referenced to a digital maximum, and any sample value beyond that maximum must be limited to the maximum.

If there are active clips with Dither enabled, you should monitor your output levels on the Output Meters with Dither Mode set to NONE and Recon Meter set to OFF. This will indicate when the Dither section would have clipped the output signal. In combination with the Peak Hold button, playing your entire track will show how far over unity your maximum sample value is. The output meter LED's are arranged to be exactly 0.5 dB per LED above 0dB. You can then use the output fader to compensate by the same amount. Finally, you would normally revert Dither mode back to ON after doing this.

If you are preparing a file or a track that will be burned directly to CD, it is imperative that you set the Oxford Limiter Dither mode to 16 bits. You can also experiment with the other dither shapes to see if this improves the sound of your mix, especially if your tracks

contain particularly quiet material. Noise shaping reduces the level of dither noise in the regions where the ear is most sensitive, at the expense of pushing the noise into regions of less sensitivity.

## 9.5 Output Meter Clip LEDs

The Clip LEDs in the output meter always display whether the output signal is clipping, and are very specific as to exactly what they indicate:

### **Dither = None**

An active clip LED indicates that the sample value has reached a value greater than 0 dBFS for one or more samples. In this case, the output clip LEDs should be in agreement with the workstation clip LEDs. When dither is disabled, the clip LEDs are drawn Orange, because a floating point 'over' value does not necessarily indicate a clipped sample value.

### **Dither = 16 bit**

An active clip LED indicates that the truncated output sample value has hit the maximum possible value in the 16 bit integer domain for one or more samples. These values are not the same as 0.0 dBFS, being very slightly smaller, and so the workstation clip LEDs may not indicate an agreement with the LEDs on the Oxford Limiter plug-in. However, since the 16 bit integer value has hit digital max, and the samples are destined to be written to 16 bit fixed precision media, the clip LEDs are drawn with an elevated danger level of Red with a '16' indication in them to remind the user they have a particular meaning.

### **Dither = 24 bit**

An active clip LED indicates that the truncated output sample value has hit the maximum possible value in the 24 bit integer domain for one or more samples. These values are not exactly the same as 0.0 dBFS (although it is closer than for 16 bit dither), so the workstation clip LEDs will not necessarily indicate an agreement with the LEDs on the Oxford Limiter plug-in. As for the 16 bit mode, since the 24 bit integer value has hit digital max, and the samples are destined to be written to 24 bit fixed precision media, the clip LEDs are drawn with an elevated danger level of Red with a '24' indication in them to remind the user they have a particular meaning.

## 10 Description of Controls



The plug-in's user interface is divided into three main areas: an Input section, the Pre-Process section, and an Output section.

Note that the parameter fields provide continuous feedback of settings values, and have type-in fields so you can change values directly from your keyboard.

### 10.1 Input Section

#### INPUT meter

Displays the effective input level from  $-42$  dBFS to  $+18$  dBFS. Gain settings producing effective programme levels above the THRESHOLD level are subject to dynamic gain control. The orange overload indication *operates on input signal levels prior to the Limiter processing*, regardless of the gain setting, in order to provide input overload warning at all times.

#### THRESHOLD value

Adjusts the THRESHOLD of the PRE-PROCESS and AUTO COMP, and the ceiling

of the ENHANCE/SAFE MODE processor.

**INPUT GAIN fader**

Adjusts the INPUT GAIN into the Limiter process between -18 dB and +18 dB. By default, the THRESHOLD for the onset of dynamic range reduction is 0 dBFS. When using the Limiter to enhance loudness, use the INPUT GAIN to increase the gain and exceed the THRESHOLD.

**INPUT GAIN display**

Displays effective gain at all times and allows settings to be typed in manually.

## 10.2 Pre-Process Section

**ATTACK fader**

Adjusts the ATTACK time between 0.052 ms and 1.04 ms.

**RELEASE fader**

Adjusts the RELEASE time between 0.052 ms and 1038 ms.

**SOFT KNEE fader**

Adjusts the response curve between 0 dB (hard limiting) and 10 dB (softer limiting).

**GAIN REDUCTION meter**

Permanently displays the total peak gain reduction due to Pre-Process dynamic gain control action.

**AUTO GAIN button**

Engages an additional long-term gain scaling process. This automatic gain scaling helps to accommodate wider programme level ranges.

**SAFE MODE button**

Engages SAFE MODE which uses the ENHANCE section to control peak output levels at all times. In SAFE MODE, the ENHANCE fader adjusts the degree of dynamic loudness boost, rather than the amount of enhancement.

## 10.3 Output Section

**ENHANCE fader**

Adjusts the amount of dynamic enhancement from 0-125%. In NORMAL MODE, full sample value limiting occurs at 100% or higher. When SAFE MODE is engaged,



sample value limiting occurs throughout the range. The ENHANCE fader then controls the degree of dynamic loudness boost.

**OUTPUT TRIM fader**

Adjusts the OUTPUT TRIM level between -18 dBr and 0 dBr.

**DITHER menu**

Switches between no dithering (NONE) and 16-BIT and 24-BIT output word lengths.

**DITHER TYPE menu**

Switches between output DITHER TYPE, including the conventional TPDF and four different noise shaping options.

**DEPTH menu**

Adjusts the amount of DITHER processing, between 0-100%, for any of the above noise shaping options.

**OUTPUT meter**

Displays the plug-in's final output level.

**OUTPUT Clip LEDs**

Indicates sample value > 0 dBFS (orange LED) for None (no dither), or maximum values in 16 bit and 24 bit dither modes (Red LED with '16' or '24' as appropriate).

**PEAK HOLD button**

Engages a HOLD option for the peak meters, these will not clear until manually done so. Without the PEAK HOLD engaged, peak meters hold for 2 seconds.

**RECON METER button**

Switches the output meter to display actual reconstructed programme signal levels. These are the signals that will be present after decoding with a Digital-to-Analogue converter. Without correction, the overload region is displayed in red. To correct for overloads in the reconstructed signal, enable AUTO COMP.

**AUTO COMP button**

Engages an additional processing stage that dynamically corrects 'true peaks' which breach the THRESHOLD. Uncorrected overloads are shown in red, corrected overloads are shown in orange.

## 10.4 Options Menu

Clicking the Sonnox button produces a drop-down options menu.

**Enable Tooltips**

Toggle display of helpful tooltips when hovering over a control.

**Clip Lights**

These options determine the approximate time that an overload indicator will stay on for when the plug-in has detected a full-level sample at either its input or output.

**Enable Sonnox Toolbar**

Displays or hides the Sonnox Preset Manager Toolbar (see Section 4).

**Show Preset Name Path**

Shows the hierarchical folder path for Presets that are stored in sub-folders of the default Preset folder.

**About Sonnox Oxford Limiter**

Displays the date, version and build number of the plug-in.

**True Peak Recon Meter Mode**

Switch the reconstructed detection and output level metering between ITU-R BS.1770 compliant 'True Peak' mode and the 'Legacy DAC Simulation' mode. See **Section 7**(True Peak Reconstruction Metering) for details.

## 11 Suggested Workflows

A step-by-step guide for three different approaches.

### 11.1 Maximum Transparency

Setting up the Oxford Limiter for the cleanest, most transparent output that is suitable for output to a file or a DAC.

1. Set the attack and release times to relatively slow (for example about midway) so that you like the sound. (Try starting out with the preset 'Clean gentle no enh')
2. Push the input gain to about +6 dB.
3. Set Safe Mode Off, and Enhance fader to 0%
4. Set Dither to None.
5. Set Peak hold On, Auto-Comp off, Recon off.
6. Run through your entire material. At the end, note how much over unity your peak value is.
- 7.

#### **Either**

Reduce the Output Fader by the same amount. If there were no over unity peaks, you can increase your output gain by the amount the peak value was under-unity.

**Or** Reduce the Input fader by the same amount. The less you push the input gain, the less compressed the sound (or the more its dynamic characteristics remain untouched.)

8. If you changed the Input level in stage 7, you must repeat stage 6 and 7 to check where the peak sample turns out to be. This is because changing the input level will have non-linear consequences to the output level.
9. Set dither mode to 16 bits (for CD) or 24 bits (for a good DAC feed by a digital output), enable Auto-Comp, and rerun through your mix to check that the output does not clip. If it does, reduce the Output fader by a fraction, say 0.0008 dB.
10. You are now ready to burn to CD.

## 11.2 Maximum Loudness

Setting up the Oxford Limiter for the Loudest Possible Impact.

1. Set Attack and Release times to relatively slow (for example about midway) so that you like the sound. (Start with preset 'Hot and Pumpy' or 'Slammer Safe').
2. Set Safe Mode On, and set Enhance fader to 0%.
3. Set dither mode to 16 bits (for CD) or 24 bits (for a good DAC feed by a digital output).
4. Set Peak Hold On and Auto-Comp On.
5. Set Output fader to -0.0008 dB.
6. You can now experiment with pushing the input level while listening to the sound, and pushing the Enhance level while listening to the sound, or both. The more Enhance you have, the bigger and fatter the sound. The more input you have, the louder the perceived result. No matter what you do to the input fader or enhance fader, you should not get hard clipping.

## 11.3 'Brick-Wall' Limiting

Setting up the Oxford Limiter as a 'Brick-wall' Limiter.

1. Set Attack and Release times to relatively fast (near the bottom) so that you like the sound without distortion appearing. (Start with preset 'Brickwall Enh').
2. Set Safe Mode On, and set Enhance fader to 0
3. Set dither mode to 16 bits (for CD) or 24 bits (for a good DAC feed by a digital output).
4. Set Peak Hold On and Auto Comp On.
5. You can now experiment with pushing the input level while listening to the sound. The more input you have, the more 'compressed' the perceived result, up to a point. You can also experiment with pushing the enhance fader to fatten your sound. No matter what you do to the input fader or enhance fader, you should not get hard clipping, but you may hear soft distortion if either are pushed too much.
6. For a cleaner sound, you may want to turn the Safe Mode off, and proceed as for Recipe 1 above, but with faster attack time.

## 11.4 Points to remember

- If you enable **SAFE MODE**, you will never see a clip.
- If you enable **AUTO COMP** with **True Peak Mode**, you will never see a clip.
- If you enable **Dither Mode**, *you will never see any more clips than you already have.*

### 11.4.1 Limiting at high sample rates

If you are mastering at higher than the delivery sample rate, it is imperative to be aware of the effects of the subsequent down-sampling process.

Downsampling will almost always increase peak level. So, clipping may occur if the input to the resampler peaks at or close to 0 dBFS.

Therefore, it is recommended that limiting to the final output level is done at the delivery sample rate. For example, when delivering a Red Book CD master, run the Oxford Limiter at 44.1 kHz.

If this is not possible,

- Leave some headroom before resampling. To do so, simply reduce the Oxford Limiter's Output Trim to at least -0.1 dB.
- Use a resampler which can normalise down to 0 dBFS automatically rather than clipping.

## 12 Specifications

### 12.1 Latency


Round-trip latency in samples:

	44.1/48 kHz	88.2/96 kHz	176.4/192 kHz
Native	141	227	415
AAX DSP	150	237	425

### 12.2 Pro Tools | HDX – Instances per chip

	44.1/48 kHz	88.2/96 kHz	176.4/192 kHz
Mono (full)	6	4	1
Mono (no recon)	15	7	3
Stereo (full)	5	2	1
Stereo (no recon)	11	5	2

## 13 Preset Manager Toolbar



Sonnox Oxford plug-ins come equipped with their own onboard Preset Manager, which is displayed at the top of the plug-in window. The reasoning behind this is to allow increased portability of your presets across all the host applications, while also providing a consistent and versatile interface. While most host platforms allow creation and loading of presets, those host-created preset files are not portable between different host applications. With the Oxford plug-ins' Preset Manager, you can create a named preset in one host application and load it when using an alternative application.

The Sonnox Preset Manager is fully described in a companion document — Sonnox Toolbar and Preset Manager User Guide — available for download at [www.sonnox.com/docs](http://www.sonnox.com/docs)

## 14 Supported Platforms

- Avid Pro Tools (AAX Native 32/64-bit)
- VST hosts (32/64-bit)
- AU hosts (32/64-bit)
- Mac OSX 10.7 or higher
- Windows 7 or higher (32/64-bit)



## 15 System Requirements

For latest System requirements, please visit [www.sonnox.com](http://www.sonnox.com).

### All versions

- Free iLok account
- Appropriate product licence
- iLok2

### Pro Tools

- Approved Avid CPU and hardware configuration
- Pro Tools 10.3.5 (Native or HDX), or higher

## 16 Copyright and Acknowledgements

Trademarks and content copyright © 2007-present Sonnox® Ltd. All rights reserved.

Sonnox® and the five dots logo are registered trademarks of Sonnox Ltd.

This product is manufactured and supplied by Sonnox Ltd. This product is protected by one or more European and/or US patents.

DIGIDESIGN, AVID and PRO TOOLS are trademarks or registered trademarks of Avid Technology Inc.

VST is a trademark of Steinberg AG.

All other product and Company names are trademarks or registered trademarks of their respective holders.