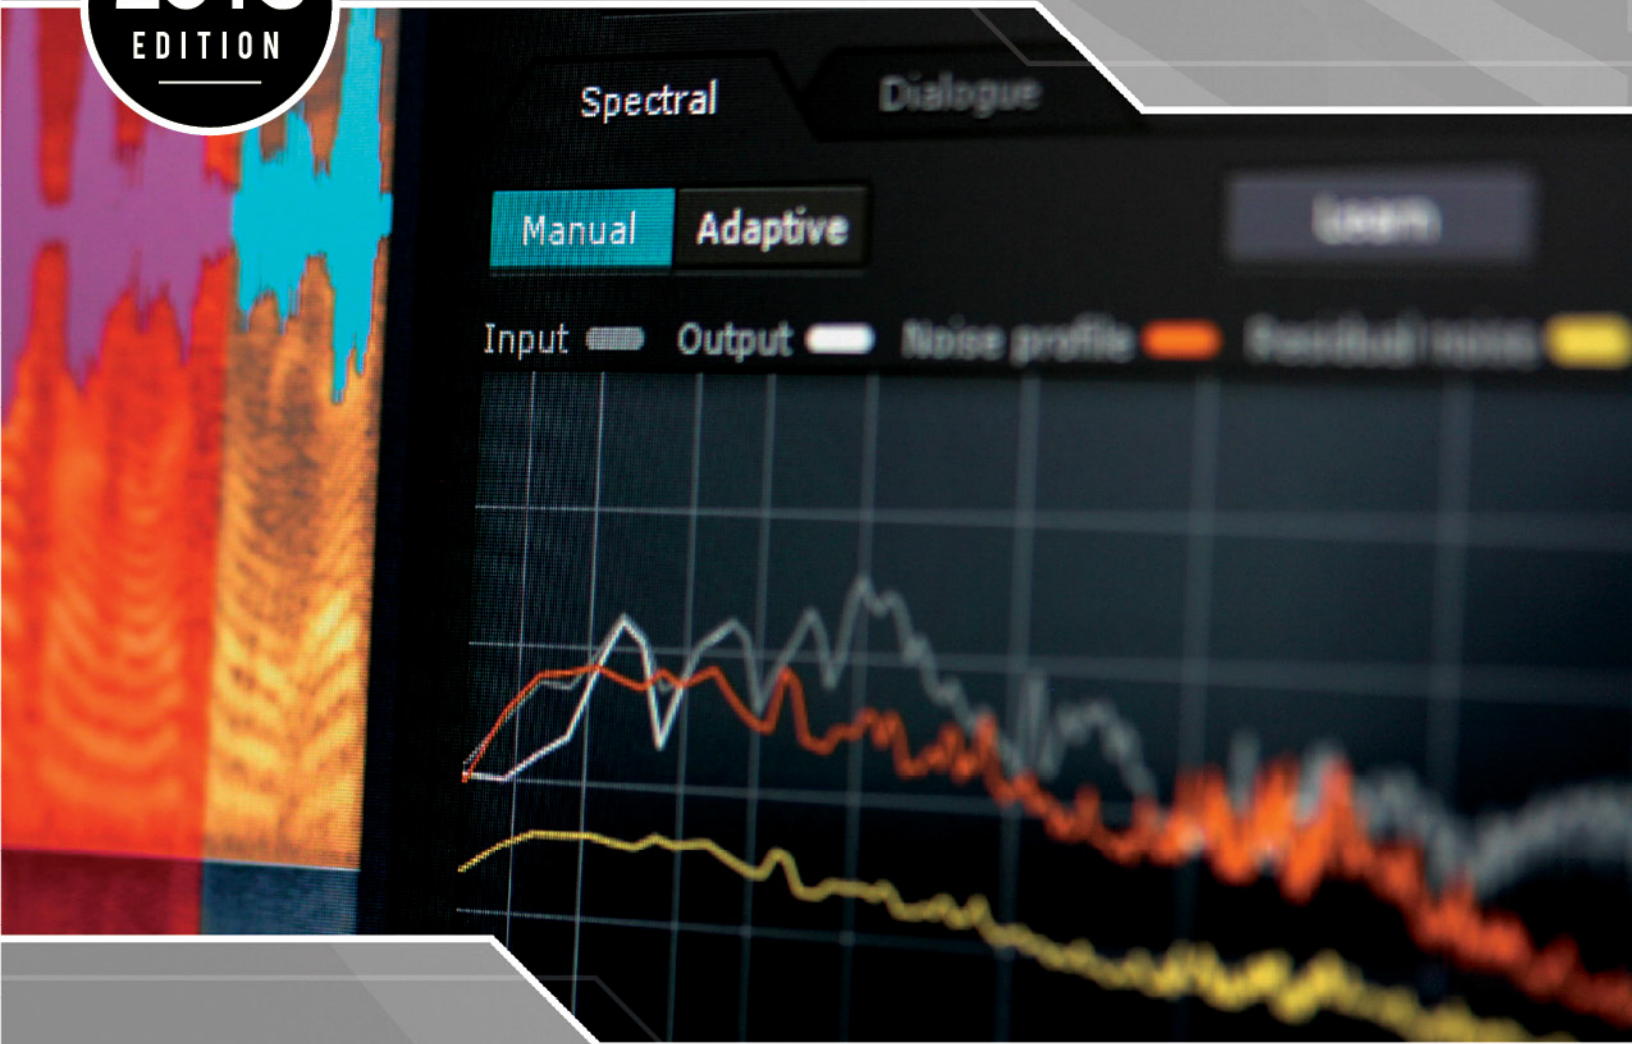


AUDIO REPAIR

WITH RX[®] 3

2013
EDITION



TOOLS, TIPS,
and TECHNIQUES



TABLE OF CONTENTS

| | |
|--|-----------|
| Intended Audience for this Guide | 5 |
| About the 2013 Edition..... | 5 |
| Additional Resources | 5 |
| About iZotope | 5 |
| 1: INTRODUCTION | 6 |
| 2: WHAT IS AUDIO REPAIR AND RESTORATION? | 7 |
| 3: AUDIO REPAIR AND RESTORATION BASICS | 8 |
| Defining the Outcome..... | 8 |
| Repair and Restoration Tools | 8 |
| General Recommendations when Repairing or Restoring Audio..... | 10 |
| Restoration Quick Tips | 10 |
| The Tools of the Trade | 11 |
| 4: UNDERSTANDING SPECTROGRAMS / IDENTIFYING AUDIO PROBLEMS .. | 12 |
| What's the Goal of Using a Spectrogram?..... | 12 |
| Principles of Spectrogram Displays..... | 13 |
| Spectrogram Types..... | 16 |
| Using a Spectrogram to Identify Audio Problems | 17 |
| 5: WHAT IS RX 3? | 23 |
| What Can RX 3 Do?..... | 23 |
| Where Can I Use RX 3? | 23 |
| 6: DENOISING | 24 |
| What's the Goal of Denoising?..... | 24 |
| Principles of Denoising | 24 |
| Using the Denoiser and Hum Removal in RX 3..... | 26 |
| Broadband Noise Reduction | 26 |
| Hum and Tonal Noise Reduction..... | 31 |

| | |
|---|-----------|
| Tonal Noise | 39 |
| Dialogue Denoising | 42 |
| General Denoising Tips | 46 |
| 7: REMOVING INTERMITTENT NOISES AND GAPS | 47 |
| What's the Goal of Audio Repair? | 47 |
| Principles of Audio Repair..... | 47 |
| Using the Spectral Repair in RX 3..... | 51 |
| Intermittent Noises..... | 54 |
| Audio Dropouts..... | 58 |
| General Audio Repair Tips..... | 62 |
| 8: REMOVING CLICKS AND POPS | 63 |
| What's the Goal of Removing Clicks and Pops?..... | 63 |
| Principles of Removing Clicks and Pops..... | 64 |
| Using Declick & Decrackle in RX 3..... | 65 |
| Analog Clicks..... | 67 |
| Digital Clicks | 68 |
| General Tips for Removing Clicks and Pops..... | 69 |
| 9: REMOVING CLIPPING | 70 |
| What's the Goal of Removing Clipping? | 70 |
| Principles of Removing Clipping..... | 70 |
| Using Declip in RX 3 | 72 |
| General Tips for Removing Clipping | 75 |
| 10: REMOVING REVERB | 76 |
| What's the Goal of Removing Reverb?..... | 76 |
| Principles of Removing Reverb..... | 77 |
| Using Dereverb in RX 3 | 78 |
| General Tips for Removing Reverb..... | 80 |
| 11: EXPORTING AND DELIVERING AUDIO | 81 |
| What's the Goal of Exporting and Delivering Audio?..... | 81 |

| | |
|---|------------|
| Principles of Exporting and Delivering Audio | 81 |
| Exporting and Delivering Audio in RX 3 | 81 |
| GENERAL TIPS FOR EXPORTING AND DELIVERING AUDIO..... | 83 |
| 12: SUMMARY | 87 |
| About the Authors | 88 |
| APPENDIX A: GETTING SET UP TO REPAIR AND RESTORE AUDIO | 89 |
| Equipment | 89 |
| Monitoring Audio | 90 |
| Transferring Media to the Computer | 91 |
| Recording from Analog Formats | 91 |
| APPENDIX B: GENERAL RX 3 TOOLS | 94 |
| Comparing Settings | 94 |
| Batch Processing..... | 95 |
| Waveform Statistics..... | 96 |
| EQ | 96 |
| Presets..... | 97 |
| Keyboard Shortcuts | 97 |
| APPENDIX C: REPAIRING EXAMPLE AUDIO | 98 |
| Example 1: Removing Broadband Noise from a Concert Recording | 98 |
| Example 2: Restoring an Historical Speech: Making Voice More Intelligible..... | 100 |
| Example 3: Cleaning up a Phone Interview with Declick and Spectral Repair | 102 |
| Example 4: Removing Clicks and Pops from a Concert on Record | 105 |
| Example 5: Removing Clipping from a Phone Interview..... | 107 |
| Example 6: Removing Guitar String Squeaks with Spectral Repair | 108 |
| APPENDIX D: TIPS FROM THE PROS | 109 |

INTENDED AUDIENCE FOR THIS GUIDE

If you don't know anything about audio repair and restoration, this guide is a great place to start. It will help you develop practical and effective methods to remove noise and fix audio problems. Sure, we think you should use iZotope RX 3® (www.izotope.com/rx) to repair and/or restore your audio... you certainly could. But we've learned so much from the audio community over the past 10+ years that we're happy to give something back in return: a guide that's useful for anyone that wants to learn more about audio repair and restoration. As a result, this guide can be freely copied or distributed for non-commercial purposes.

If you own RX 3 but don't fully understand audio repair and restoration, this guide will help you better understand the powerful tools at your disposal. Each chapter demonstrates many useful concepts that you can apply to your next piece of problematic audio.

If you own RX 3 and already know the basics of audio repair and restoration, this guide will show you new tricks and techniques that are possible in RX 3's modules. Just read through and say, "Yeah, I knew that" when appropriate for the other parts.

You can also follow along by downloading the free 10-day trial of RX 3 at www.izotope.com/rx.

ABOUT THE 2013 EDITION

The 2013 edition of this guide has been revised and updated by the experts at iZotope HQ, based on years of research and consultation.

ADDITIONAL RESOURCES

If you're interested in exploring audio repair and restoration in greater depth, there are plenty of resources available. The iZotope YouTube channel (www.youtube.com/izotopeinc) has a number of audio repair- and restoration-centric tutorial videos aimed at anyone from the beginner to the advanced audio expert. Additionally, third party training is available via the iZotope website. Learn more at www.izotope.com/store/tutorials.asp

ABOUT IZOTOPE

iZotope makes innovative products that inspire and enable people to be creative. Based in Cambridge, Massachusetts, iZotope has spent over a decade developing award-winning products and audio technologies for professionals and hobbyists alike. Learn more at www.izotope.com.

1: INTRODUCTION

In today's world, audio and video recordings can be made by almost anyone. From smartphones to sophisticated studios and sound stages, large amounts of media content are created daily. And as our world becomes increasingly and audibly congested, the rate of ruined recordings is rising in tandem. Human error, unexpected electrical or mechanical interference and unwelcome intrusions from aircraft, cell phones, pets, people and Mother Nature regularly impact even the most seasoned professional. Many times, it just isn't possible to record that "perfect take" over again, particularly when editing coverage of a live event, or working to meet budgets and deadlines.

Whether you're a professional or just getting started, you've probably ran into difficult audio situations like those mentioned above, and wished you could correct what seemed like an impossible dilemma. We hope this guide will help you understand the basics (and beyond) of audio repair and restoration and enable you to fix previously unusable audio.

2: WHAT IS AUDIO REPAIR AND RESTORATION?

When you hear the words “repair” and “restoration,” you might be inclined to think of dusty vaults filled with aging master tapes and records. Some audio repair and restoration projects may indeed involve taking old recordings and reviving them, but the methods used are useful for a wider range of scenarios. Every time you record audio—whether at home, in the studio, or on location—there’s always the chance of encountering unexpected and unwelcome audio “guests.”

“Audio repair and restoration” is a phrase used to describe the various processes and techniques one can use to remove noise and other imperfections from sound recordings. When used correctly, these techniques can alleviate problems including:

- Ambient background noise
- Tape hiss
- Electronic interference such as hum and buzz
- Sudden background noises (coughs, ringing cell phones, etc.)
- Clicks and pops from older vinyl, shellac or phonograph recordings
- Clipping in both the analog and digital domain

3: AUDIO REPAIR AND RESTORATION BASICS

DEFINING THE OUTCOME

The goal of good audio repair and restoration is to render the best possible sonic result with the least audible human intrusion. In essence, your intervention in the original recording should be transparent and not introduce new artifacts that distract the listener. Sometimes it's possible to solve an audio problem entirely, and other times it's about finding the right balance between reducing the problem and preserving the original audio.

It's useful to remember that no recording is truly perfect, and any statement as such is purely subjective.

In the early 1980's for instance, when CD players became readily available and affordable, commercial tape vaults and archives turned into gold mines. But strangely, during the early days of CD restoration, whole armies of purists became outraged when record companies sought to change the character of certain, revered recordings (no matter how old) from their original, boxy-quality and dull, wooly ambience. Today, we have fresher ears that seem to be more objective.

Whatever your tastes, and even as times and nostalgic aesthetics change, the basic intention of restoration should remain the same: render the best possible sound with the least obvious interference.

REPAIR AND RESTORATION TOOLS

Repairing and restoring audio typically involves working with the following types of processors:

- Denoisers are used to reduce and remove steady state background noise. "Steady state" means slowly changing noise. It might include constant ambient noise or tape hiss (referred to as "broad-band" or "noisy" noise), or electrical buzz and hum (referred to as "tonal" noise because it typically exhibits recognizable pitches or harmonics). Denoisers can be spectral or multiband, software or hardware (such as iZotope ANR-B), and are sometimes designed for a specific use case, such as vocals.

RX 3 TIP • RX 3 includes a single band Denoiser with independent control for both tonal and noisy problems (it includes an envelope that may be used to shape the noise reduction curve). RX 3 Advanced also includes a multiband Dialogue Denoiser, which is more suited for real-time noise reduction on spoken word or sung vocals within a mix.

- Declickers are used to reduce and remove intrusive clicks and pops. These can be caused by anything from dust and scratches on an old record, a CD skipping on playback, or even mouth clicks and lip smacks from a voiceover.
- Decracklers are closely related to Declickers, but are optimized to help reduce and remove a more continuous, quieter stream of clicks that blend together to cause what the human ear perceives as a general crackle.

TIP • Using a Decrackler before using a Denoiser is often a very effective way of dealing with surface noise recorded from vinyl or shellac records.

- Declippers are used to repair digital and analog clipping artifacts. These artifacts occur when overloading an A/D converter or over-saturating magnetic tape.
- Visual Editing Tools vary by manufacturer, but the basic premise combines visual representations of audio, via a waveform or a spectrogram, with tools allowing you to select and edit certain audio events rather than the entire file.

TIP • Chapter 4, Understanding Spectrograms / Identifying Audio Problems, will help outline the fundamentals of working with audio in this new, visual way, before we dive in to the various tools in depth.

- Dereverbs are a new, cutting edge technology, and are designed to remove or reduce reverberations from audio. They are particularly useful for dialogue editing and ADR matching, and allow the engineer to remove unwanted or distracting reverberations from dialogue recordings.

RX 3 TIP • RX 3 Advanced has a proprietary Dereverb tool developed by iZotope.

With all of these tools available, you might wonder where to begin. There isn't a single "correct" order in which to use them—it all depends on the audio material you're restoring. Always begin with the most obvious or obnoxious audio problem that you can hear and identify. Then, depending on the audio, it may make sense to perform some processing tasks before others. For example, a loud hum, a heavy crackle, or severe clipping might at first prevent you from hearing and dealing with additional audio problems. Peeling away that first layer may make the next step more obvious to you. Don't be afraid to try out different combinations of the tools to get the result you want.

GENERAL RECOMMENDATIONS WHEN REPAIRING OR RESTORING AUDIO

We highly recommend that you educate yourself about the function of individual tools in your toolbox. Dedicated audio repair and restoration tools can do a fantastic and fairly autonomous job, but learning when, where, and to what degree of strength to use a specific tool can yield better, faster, and more transparent results.

It's also beneficial to establish your aim before setting out. Do you want to treat each file with individual care and attention, or would you rather define the most appropriate settings with which to batch process hundreds of files quickly?

RESTORATION QUICK TIPS

1. Back up your work. Always make a backup of the original audio file before you begin attempting to restore it. Depending on the tool, some edits become permanent once the file is saved, so it's always advisable to maintain a prior backup.

RX 3 TIP • RX 3 allows you to save your work and unlimited undo history as an RX 3 document, which can prevent losing or overwriting work.

2. Keep the ears rested and the mind open. While doing audio restoration work, you'll likely spend a lot of time focusing on subtle details. Taking breaks will help you return with a fresh mind and see and hear the bigger picture.
3. Make multiple versions. Sometimes it helps to try doing the same audio repair more than once with different settings and then compare the results.

RX 3 TIP • RX 3 has a great Compare Settings tool that helps A/B results as you go. Also, you may come back to a version you tried a few days earlier when you were tired, and now find it sounding worse than ever. This happens to all of us! See suggestion #2.

4. Keep detailed notes. This is invaluable, particularly when there are so many different methods for dealing with different audio problems. For forensics work, documentation is often a required deliverable.

RX 3 TIP • Using RX 3 document files and saving module-specific presets can save you the trouble of writing out all of the parameters on a recall sheet as you would in the analog domain.

5. Back up your work. The first and last rule of any audio editing project! You never know when a hard drive, backup device or original master might fail. Again, always back up your work!

THE TOOLS OF THE TRADE

In the following chapters, we'll briefly examine some of the essential audio repair and restoration tools. This isn't meant to be a comprehensive guide, but will hopefully give you some focused thoughts about the tools and their uses.

Every noise is different, so a certain element of trial and error will always remain. As you continue to learn by experimenting with different tools and processes, you'll be able to refine and optimize your techniques over time.

4: UNDERSTANDING SPECTROGRAMS / IDENTIFYING AUDIO PROBLEMS

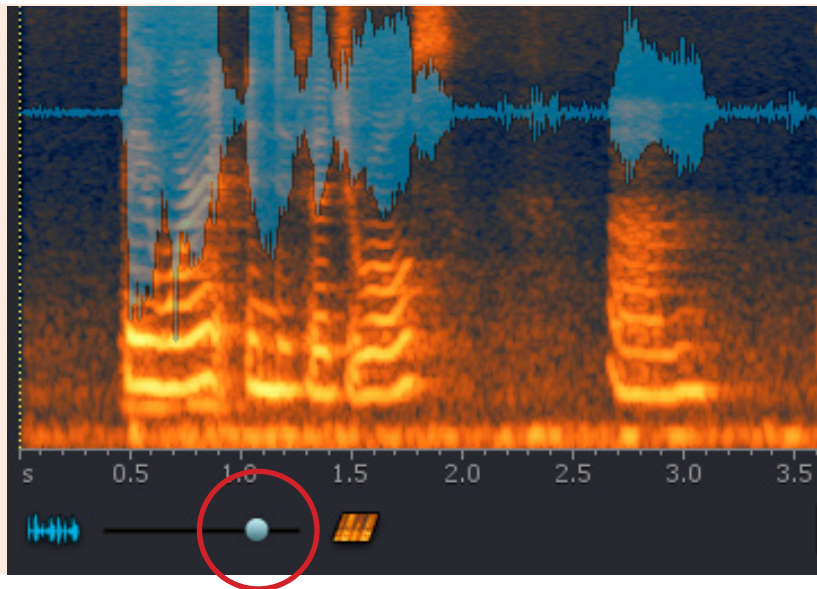
As with medical diagnostics, the key to successful audio restoration lies in your ability to correctly analyze the subject's condition. This can be a life-long, never-ending quest—constantly honing the ear to distinguish the noises and audio events that need to be corrected.

To get started, it's important to identify the problems with your file and identify which tool(s) will give you the results you want. Let's briefly look at how to examine your audio using the spectrogram and waveform display tools, then consider how to identify audio problems using these displays.

WHAT'S THE GOAL OF USING A SPECTROGRAM?

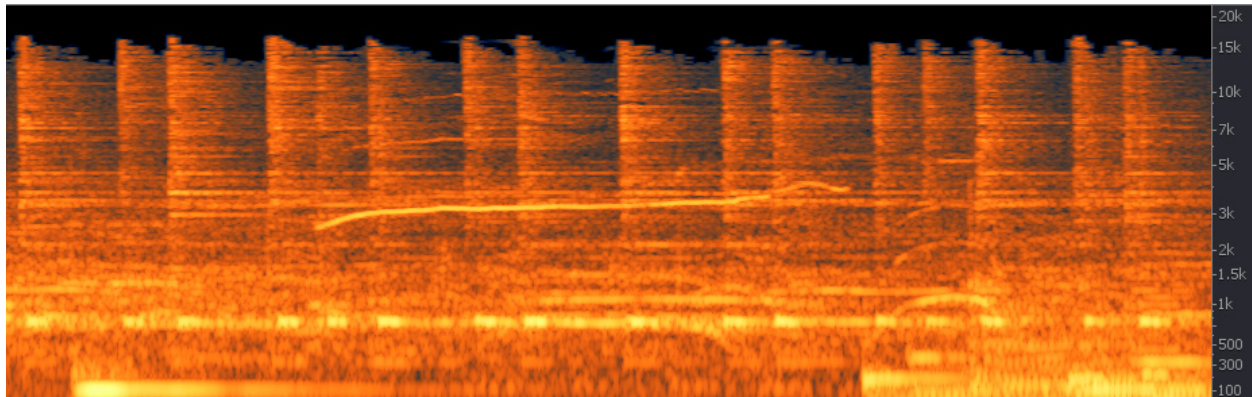
The aim of any good visualization tool for audio repair and restoration is to provide you with more information about an audible problem. This not only helps inform your editing decisions, but, in the case of a spectrogram display, can provide new, exciting ways to edit audio—especially when used in tandem with a waveform display.

RX 3 TIP • Use this special slider in RX 3 to blend between detailed spectrogram and waveform views.



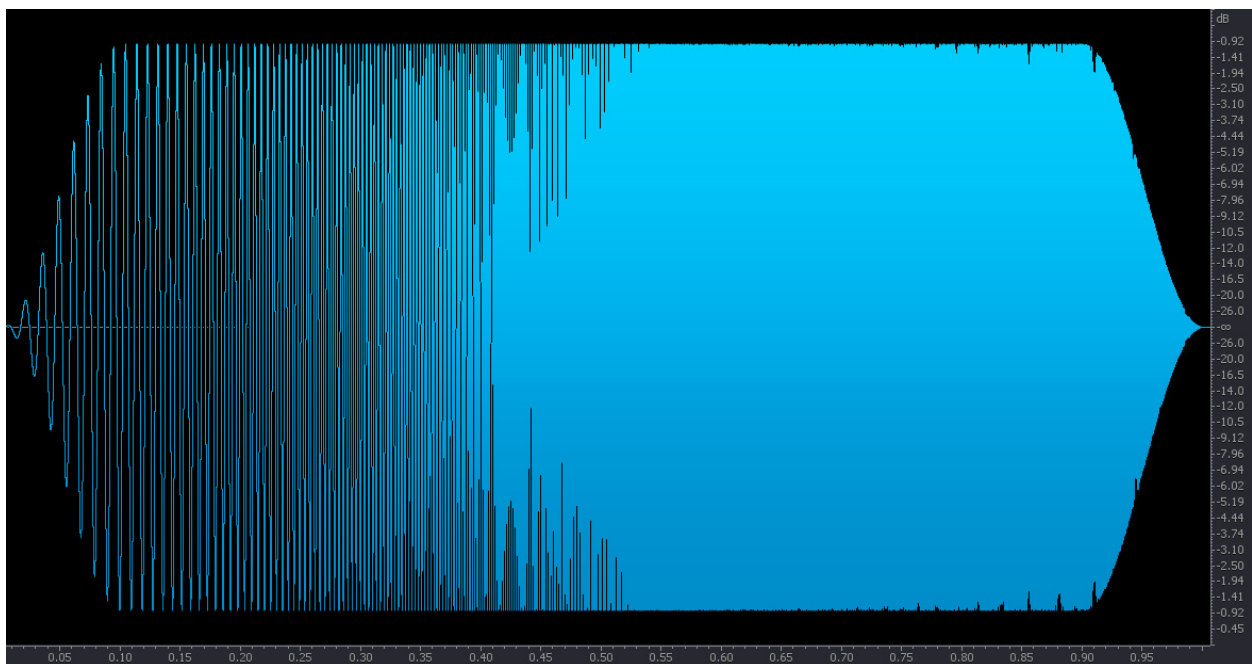
PRINCIPLES OF SPECTROGRAM DISPLAYS

So what's a spectrogram? A spectrogram is a very detailed, accurate image of your audio, displayed in either 2D or 3D. Audio is shown on a graph according to time and frequency, with brightness or height (3D) indicating amplitude. Whereas a waveform shows how your signal's amplitude changes over time, the spectrogram shows this change for every frequency component in the signal.



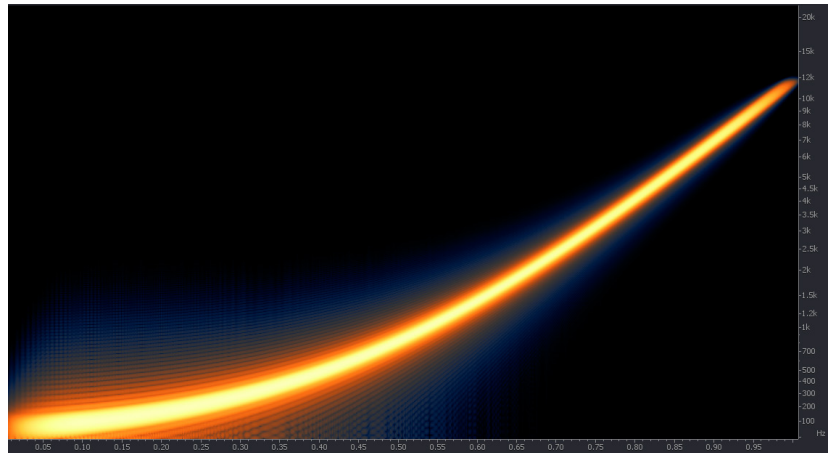
If you're used to using the waveform display, it may take a while to get your head around this unique way to "see" the audio. As a start, let's look at a few simple pieces of audio.

Here's a picture of a sine wave moving up in pitch from 60 to 12,000 Hz as seen using a waveform view.

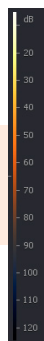
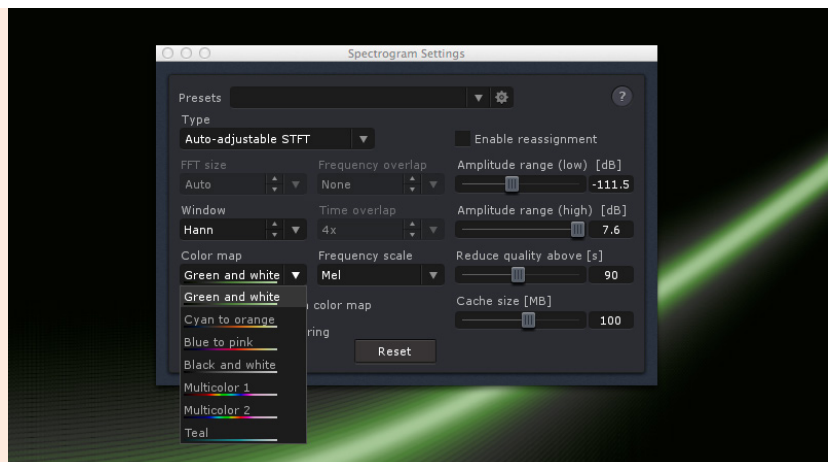


One thing you'll notice when looking at the waveform display is that it's good at showing audio amplitude, but less effective at showing what's happening at different frequencies. For example, we can easily see here that the sine wave is the same loudness for the entire duration of the file. However, we can't tell much about how the pitch or frequency changes over time. Now let's look at this same audio file using a spectrogram.

Now it's very obvious that the pitch of the audio is moving up! The horizontal axis shows time, just like the waveform display. But now, the vertical axis shows us frequency in Hz—the pitch of the event that's happening. We can see how loud events are by how bright the image is. The black background is silence, while the bright orange curve is the sine wave moving up in pitch.

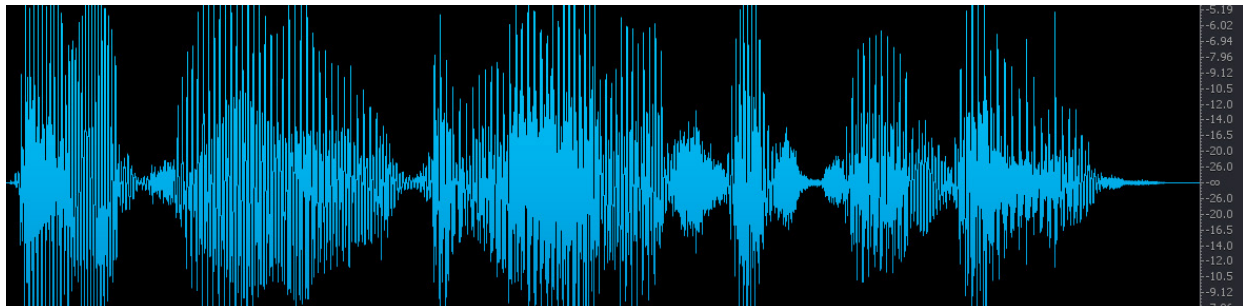


RX 3 TIP • The blue/orange color map is optimized for easy sound identification, and is the default spectrogram color map in RX 3. In the spectrogram settings, you may change the color map based on personal preference.

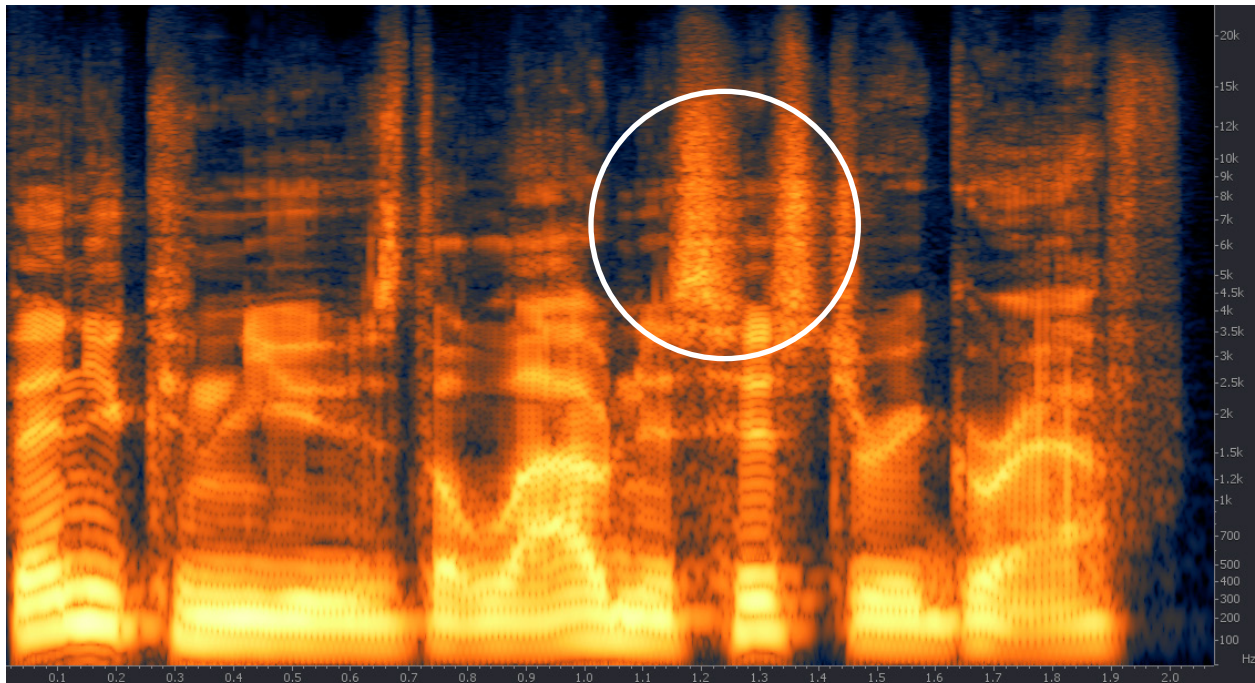


RX 3 TIP • Using this color map slider, you can increase the brightness of the Spectrogram to make certain things easier to see.

Now let's look at something more complex: the human voice. Here's a short, spoken phrase as seen in with a waveform display:



What we're seeing here is the amplitude of the spoken words over time. If we switch to the spectrogram view, we'll see many things we can't see in the waveform view:



The human voice is much more complex than it might seem from looking at the waveform view. Each word is made up of a fundamental frequency (at the bottom of the spectrogram), harmonics that extend above that frequency, sibilance ("S" sounds) that begin or end words, and more. And of course, you can now see more clearly the noise that is surrounding the voice.

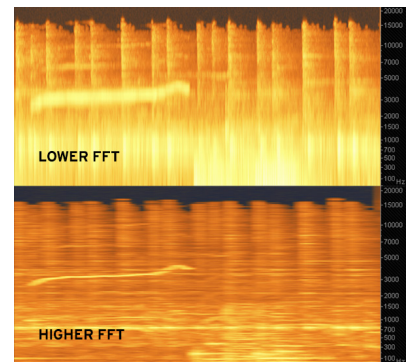
This is why having a detailed spectrogram display is so important to doing audio restoration. It helps you clearly see the problems that you're trying to fix.

SPECTROGRAM TYPES

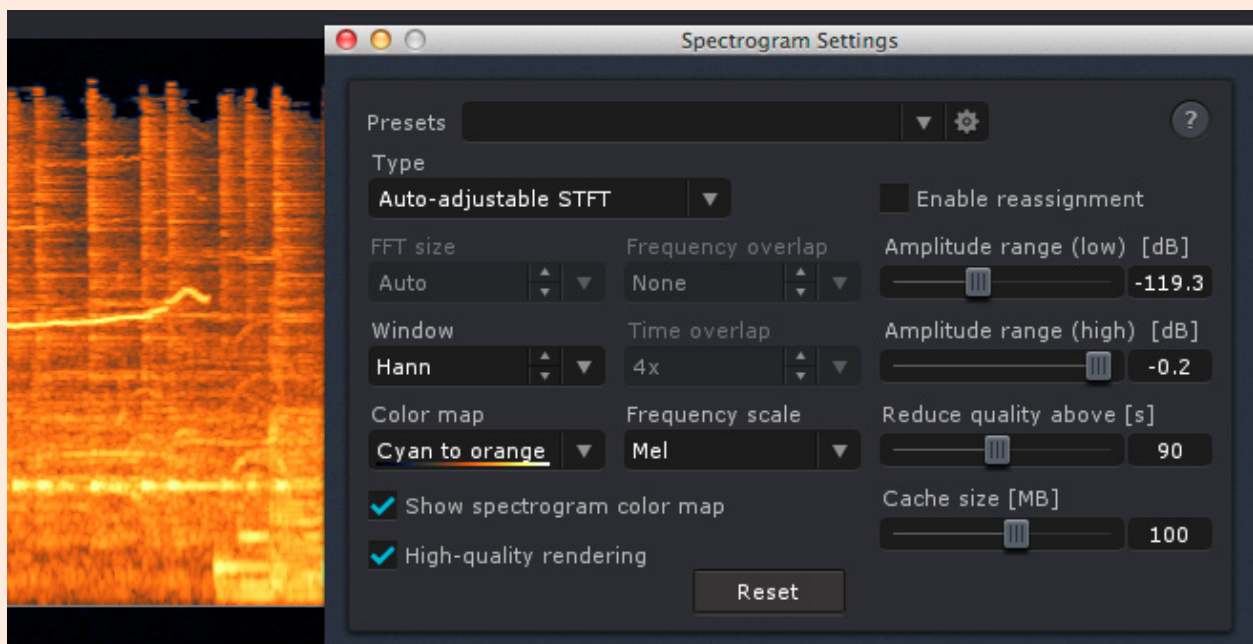
Not all spectrograms are created equal. An algorithm known as the “Fast Fourier Transform”, or FFT for short, is used to compute this visual display. Many products that feature a spectrogram display allow you to adjust the size of the FFT, but what does this mean for audio repair and restoration? Changing the FFT size will change the way the algorithm computes the spectrogram, causing it to look different. Depending on the type of audio you’re working with and visualizing, this may help. As a rule, higher FFT sizes more accurately display low frequency information, while lower FFT sizes more accurately display transient information.

If you’re trying to identify a plosive, mic handling noise, or other muddy low-frequency information, a higher FFT size in your spectrogram settings will help. If you’re trying to identify a high frequency event, or working with a transient signal (such as a percussion or drum loop) choose a lower FFT size.

The following image is of a drum loop in a live concert setting, with a member of the audience whistling. You can see how the different FFT sizes affect the way we see high vs. low frequencies, as well as transients vs. sustained notes.



RX 3 TIP • RX 3 has an Auto-adjustable STFT mode that automatically selects FFT size according to your zoom ratio, so it's always easy to see everything that's going on.

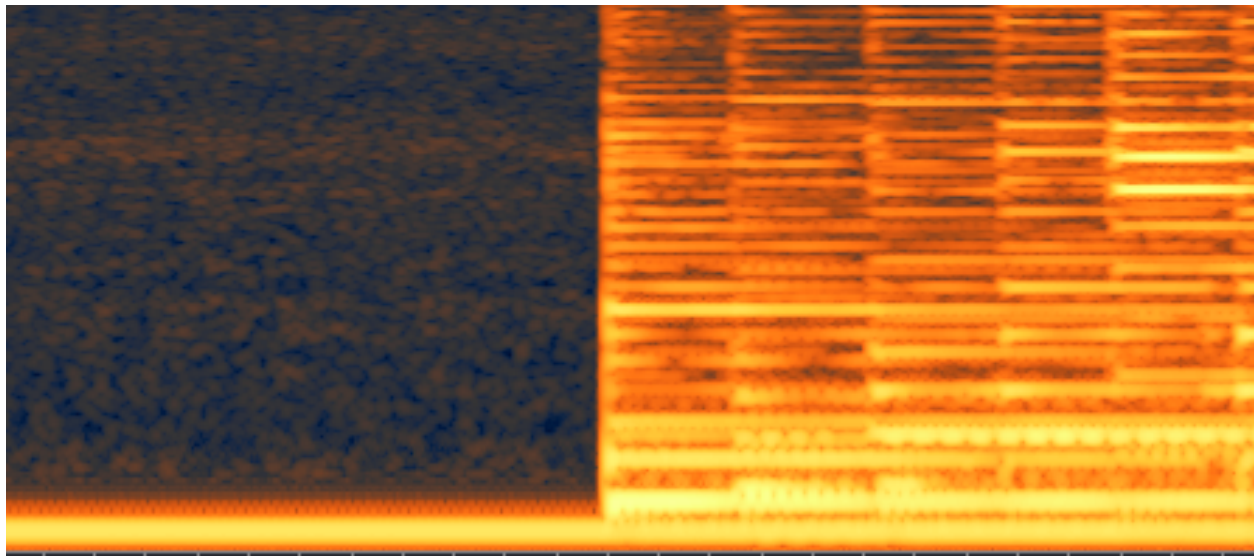


Now we'll move on to some specific examples of how to visually identify noise and other audio problems. Learning to identify these problems by sight will greatly help you—it means you'll be able to use any software that includes spectrogram technology.

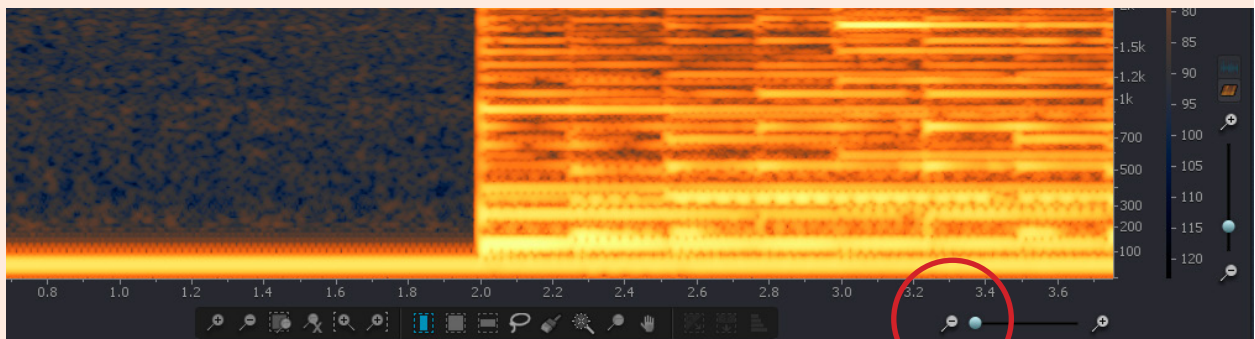
USING A SPECTROGRAM TO IDENTIFY AUDIO PROBLEMS

Hum

Hum is usually the result of electrical noise somewhere in the recorded signal chain. It's normally heard as a low-frequency tone based at either 50Hz or 60Hz depending on whether the recording was made in North America or Europe. If you zoom in to the low frequencies, you'll be able to see hum as a series of horizontal lines, usually with a bright line at 50Hz or 60Hz and several less intense lines above it at harmonics. See below:



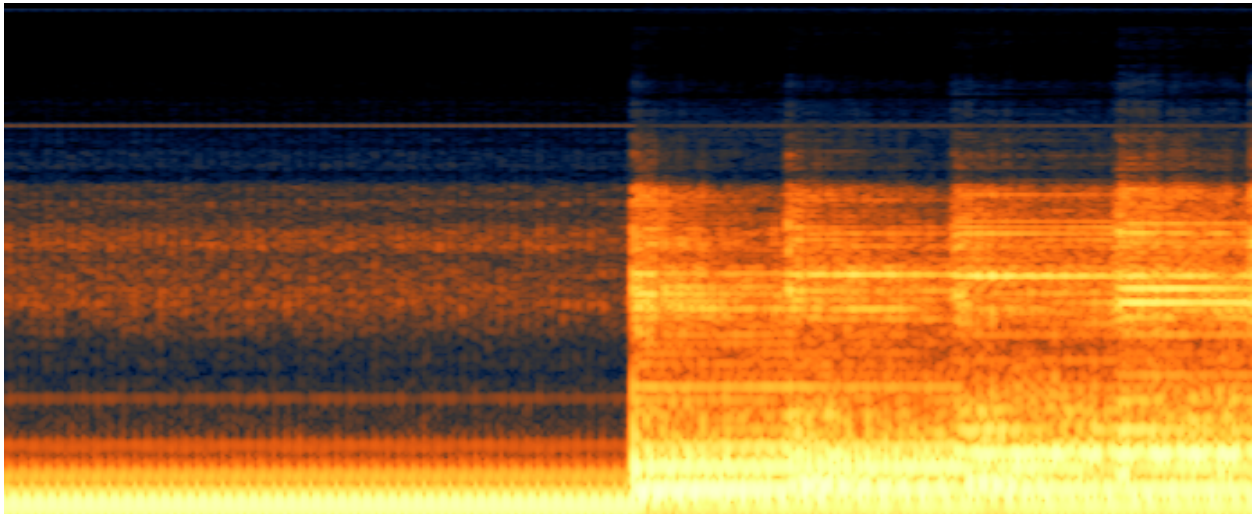
RX 3 TIP • To zoom in on the spectrogram with RX 3, you can scroll the mouse wheel, or use the magnifying glass tool.



TIP • In situations like this example, which only exhibits a few harmonics, a hum removal tool is an ideal choice for getting the job done.

Buzz

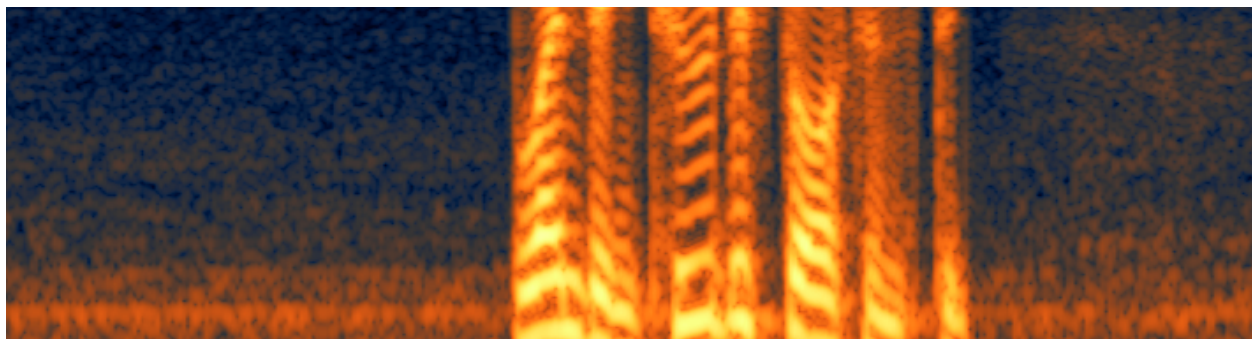
In some cases, electrical noise will extend up to higher frequencies and manifest itself as a background buzz. See the example below:



TIP • Hum removal tools usually focus on low-end hum, so when the harmonics extend to higher frequencies, a denoiser is more effective at removing the problem.

Hiss and other Broadband Noise

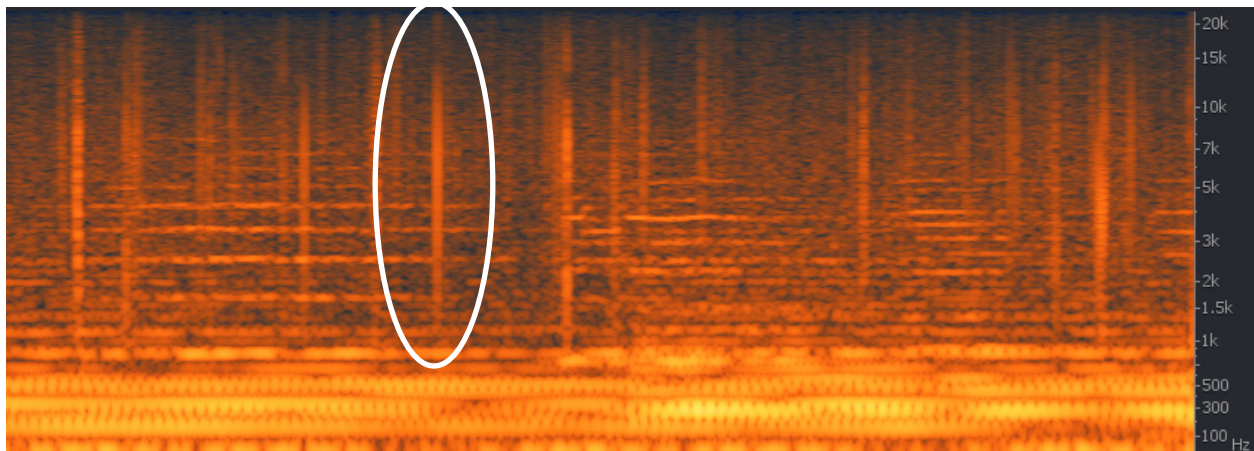
Unlike hum and buzz, broadband noise is spread throughout the frequency spectrum and isn't concentrated at specific frequencies. Tape hiss and noise from fans and air conditioners are good examples of broadband noise. In a spectrogram display, broadband noise usually appears as speckles that surround the program material. See image below:



TIP • Denoisers are very effective at dealing with this type of broadband noise.

Clicks, Pops and Other Short Impulse Noises

Clicks and pops are common on recordings made from vinyl—but can also be introduced by digital errors, including recording into a DAW with improper buffer settings, or making a bad audio edit that missed a zero crossing. Even mouth noises such as tongue clicks and lip smacks fall into the clicks category. These short impulse noises appear in a spectrogram as vertical lines. The louder the click or pop, the brighter the line will appear. The example below shows clicks and pops appearing in an audio recording transferred from vinyl:



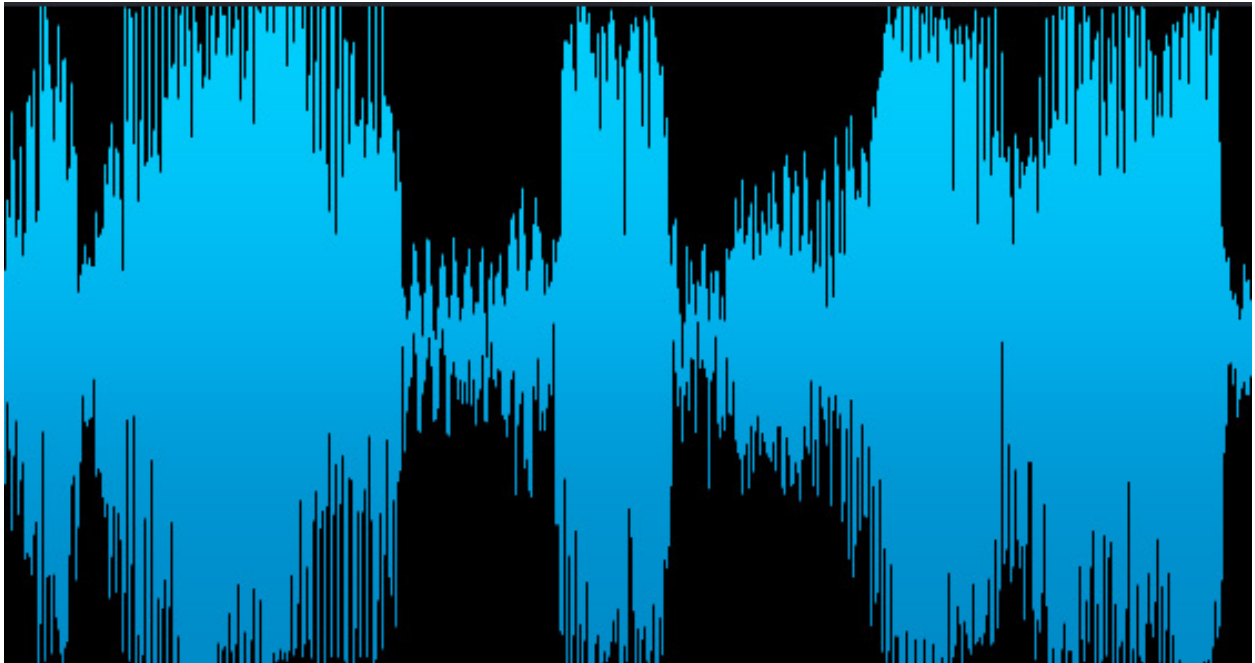
TIP • Declicking tools can recognize, isolate, then reduce and remove clicks such as these.

Clipping

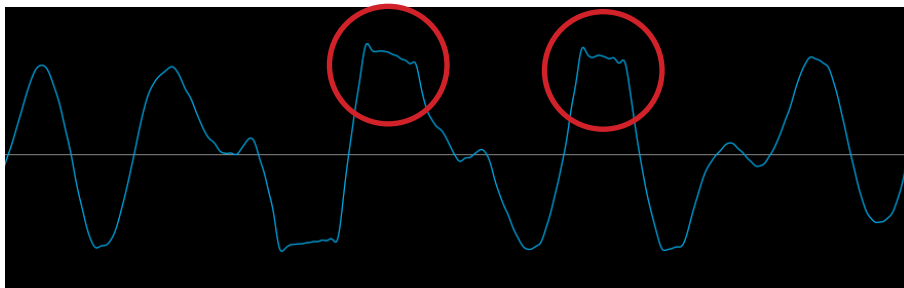
Clipping is an all-too-common problem. It can occur when a loud signal distorts on input to a sound card/converter, mixing console, field recorder or other sound capture device. A spectrogram is not particularly useful for identifying clipped audio—for this you'll want to work with a waveform display.

RX 3 TIP • Move the slider below the spectrogram to the left to superimpose the waveform display.

As you'll see in the image below, the clipping appears as “squared off” sections of the waveform.



Many software programs allow you to zoom in on a waveform and see in detail where the waveform has been truncated.

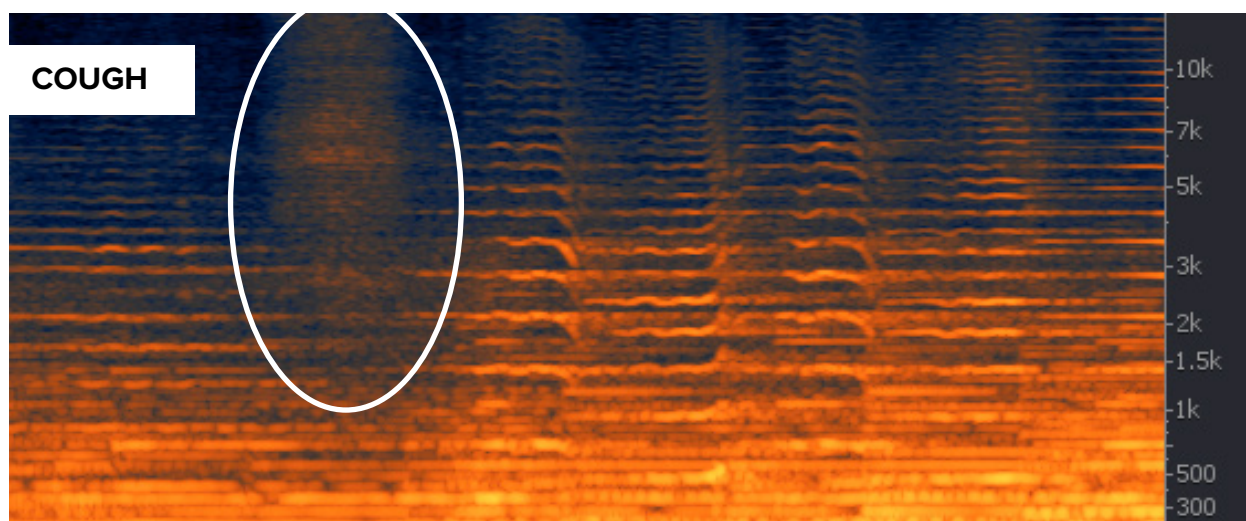
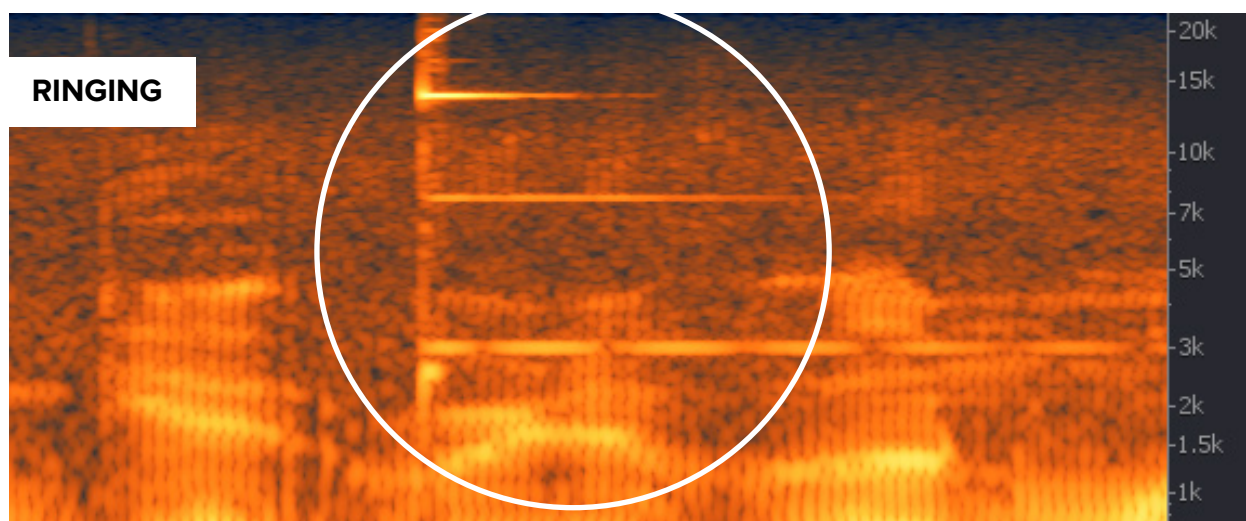


TIP • Declicking tools can intelligently redraw the waveform to where it might naturally have been if the signal hadn't clipped.

TIP • Sometimes, heavily limited audio will also appear “squared off” when zoomed out, but this doesn't necessarily mean it will sound as heavily distorted. You can zoom in to see if individual waveform tops are clipped.

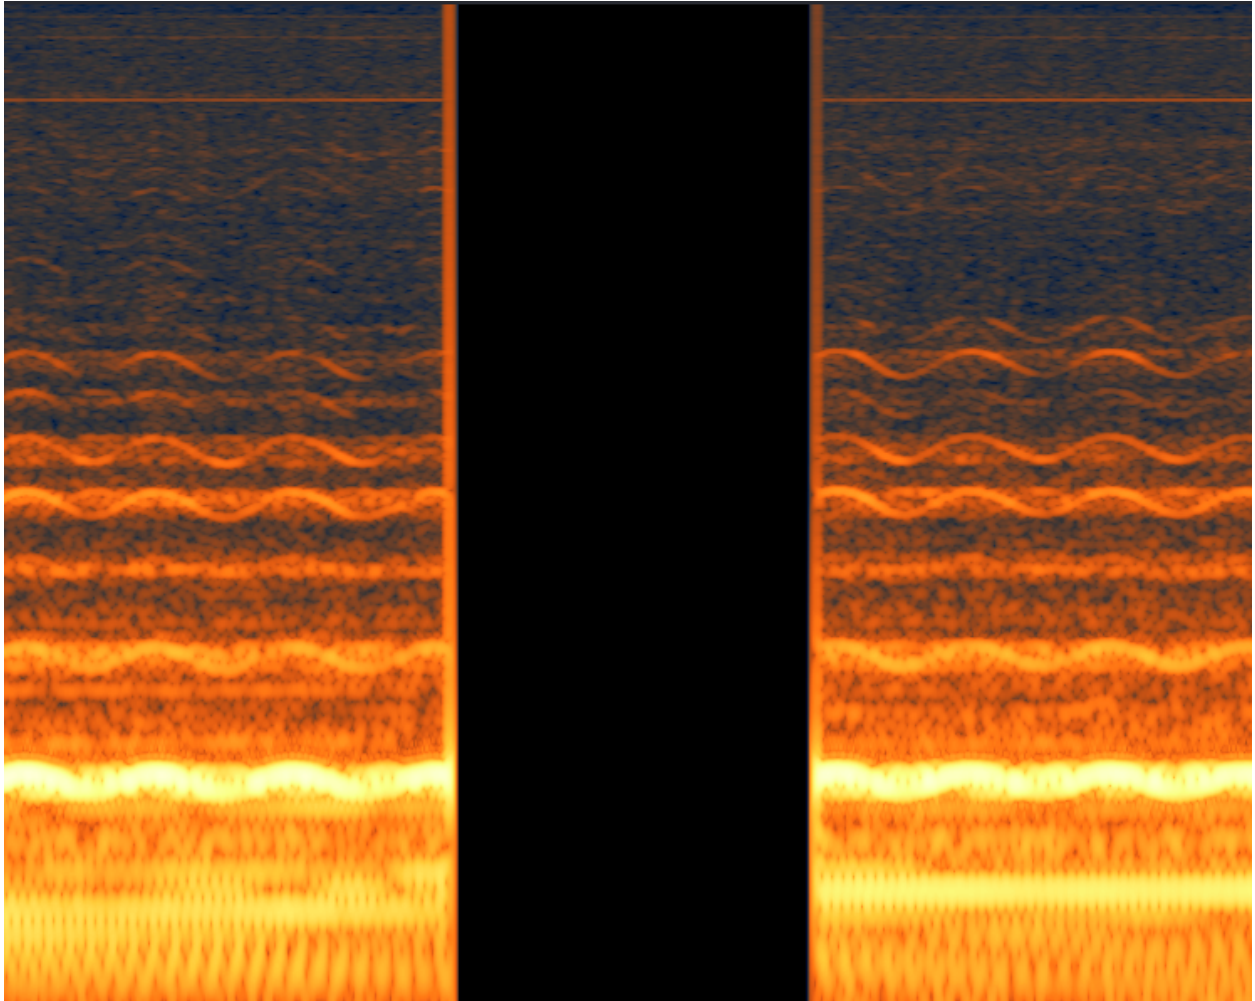
Intermittent Noises

Intermittent noises are different than hiss and hum—they may appear infrequently and may not be consistent in pitch or duration. Common examples include coughs, sneezes, footsteps, car horns, ringing cell phones, etc. The images below represent two different examples of these noises:



Gaps and Drop Outs

Sometimes a recording may have short sections of missing or corrupted audio. These are usually very obvious to both the eye and the ear! See below:



TIP • Noises and dropouts like the example above are often unpredictable, and usually need to be removed or patched manually using a visual/spectral editing tool.

RX 3 TIP • The Spectral Repair module in RX 3 is one of the most revolutionary and effective ways of repairing noises and dropouts.

5: WHAT IS RX 3?

As we explore the various ways one can approach audio repair and restoration, we'll be using iZotope RX 3 to demonstrate key concepts and showcase some unique tips and tricks.

RX 3 is a complete audio repair and restoration suite, consisting of both a standalone audio editor and software plug-ins for use within a DAW. RX 3 offers advanced visual representations of your audio and innovative new processing technology that removes noise and repairs audio. Its unique workflow features are designed to help you get great-sounding, efficient results even when working on challenging projects.

WHAT CAN RX 3 DO?

RX 3 combines Photoshop™-style selection tools with a visual spectrogram display and advanced audio processing, enabling you to perform powerful audio edits. For example:

- Reduce both ambient and signal-based background noises such as hiss, hum and buzz—without sacrificing the fidelity and clarity of the original audio
- Isolate and then reduce or accentuate specific audio events, such as sudden background noises
- Replace damaged or missing sections of an audio file with seamless, natural-sounding patches
- Eliminate pops, clicks, and mouth noises without audible artifacts
- Repair audible analog and digital clipping with perfect precision

WHERE CAN I USE RX 3?

As a standalone software editor and plug-in suite, RX 3 is ideal for performing both real-time processing and offline editing in various audio workflow situations. Such use cases include post-production mixing for broadcast (TV, radio, web), audio mixing and mastering, audio restoration and archiving, digital audio forensics and any other scenario that demands flawless audio delivered in a timely manner.

6: DENOISING

“Denoising” in the context of audio typically refers to the reduction or removal of steady state background noise. As explained in Chapter 3, steady state noise might include constant ambient noise, tape hiss or electrical buzz and hum.

WHAT’S THE GOAL OF DENOISING?

There are many situations where removing steady state background noise can greatly improve the quality of the material.

Removing steady state noise can improve the overall listening experience, increase intelligibility in dialogue tracks, remove interferences like electrical hum from a musical performance, and make it easier to combine different audio files into one scenario when mixing sound for picture.

However, there are other considerations too. If the noise carries a lot of low frequency information (often described as rumble, hum etc.), reducing this can improve headroom and open up a musical or post-production mix to greater dynamic range and mixing possibilities.

PRINCIPLES OF DENOISING

Denoisers are one of the most commonly used tools in audio repair and restoration. They work by first taking a sample of the noise’s frequency spectrum, either manually or automatically, referred to as a “noise profile.” This noise profile is then used to distinguish between the desirable (and usually much more variable) audio, and the steady, undesirable background noise. As the algorithm identifies and separates the noise from the desired signal, the level of noise is then intelligently suppressed. The user is able to control the level of noise reduction applied, as well as certain parameters affecting how the noise reduction works.

TIP • These additional parameters will vary by product and by manufacturer. Try downloading a trial or listening to demonstrations online to find the best fit for you! For instance, more information on iZotope RX 3 can be found at izotope.com/rx, or via our YouTube channel, youtube.com/izotopeinc

There are a number of denoising tools available, both software (such as iZotope RX 3’s Denoiser and Dialogue Denoiser) and hardware (such as iZotope’s ANR-B).

When denoising, there are some universal home truths that apply regardless of the type of background noise or the tool you’re using:

1. Reduce > Remove. Often, if the goal is to improve the listening experience or increase headroom, the best result isn't necessarily obtained by attempting to remove the noise entirely. Rather, reducing the noise to an acceptable or indistinguishable level often yields better results.
2. Doing multiple, gentle processes will often yield a more natural result than one harsh process of denoising. Though you do need to relearn the new noise before each processing pass, this softly-softly approach allows the Denoising tool to more easily reduce noise without affecting the desirable signal.
3. You can often reduce tonal noise more heavily, and with fewer artifacts, than you can reduce broadband noise. Identify the type of noise(s) present in your signal, and keep this in mind when treating them.

RX 3 TIP • The iZotope RX 3 Denoiser provides separate control over the tonal and noisy components in your noise profile. They are linked by default, which usually delivers the best results. There are use cases for delinking, however—like when trying to remove low-end amp buzz without reducing the high-end noisy shine the amp gives the guitar signal.



4. Applying different levels of noise reduction to different areas of the frequency spectrum can help reduce artifacts. For example, try applying more gentle processing to the high-mids than you would to the low-end, which usually contains less important audio information.

TIP • Multiband denoisers, or denoisers with enveloping capabilities, are designed to accommodate this workflow. RX 3 Advanced makes both options available should you need them.

USING THE DENOISER AND HUM REMOVAL IN RX 3

There are some slight differences in the tools you may use to treat tonal versus broadband noise, which we'll now investigate.

BROADBAND NOISE REDUCTION

In Chapter 4, we learned how to use a spectrogram to visually understand what's happening in our audio.

Broadband noise appears as random speckles or a haze over the audio file—similar to the static you see on your TV when there is no signal.

Below are screenshots of a noisy voiceover sound recording:

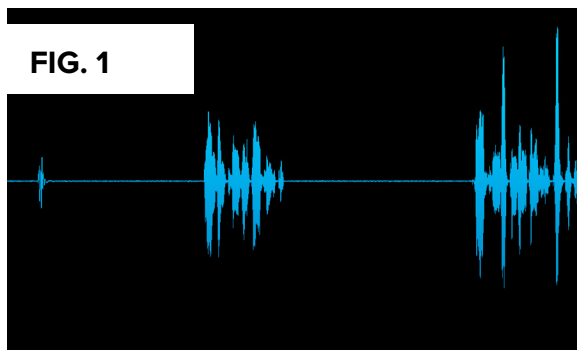


FIG. 1

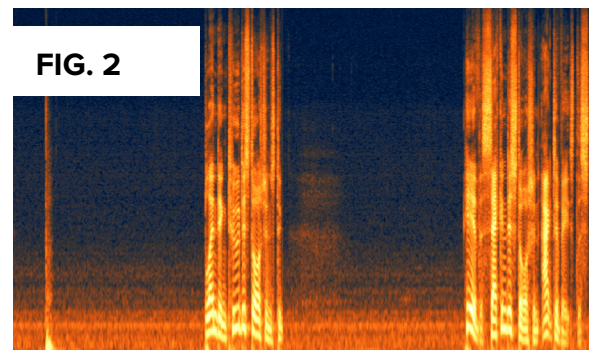


FIG. 2

Fig. 1 demonstrates the audio file via a waveform view, and we can see the blocks of noise between the transient audio events that make up our voiceover.

Fig. 2 demonstrates the same audio file via a spectrogram view, and there are visible speckles. It's important to use a spectrogram when working with noise. As we see in this view, the block of noise on the left is a better point to learn the noise profile from than the block of noise on the right, which has some other audio elements mixed in that could confuse the noise learning process.

Step 1:

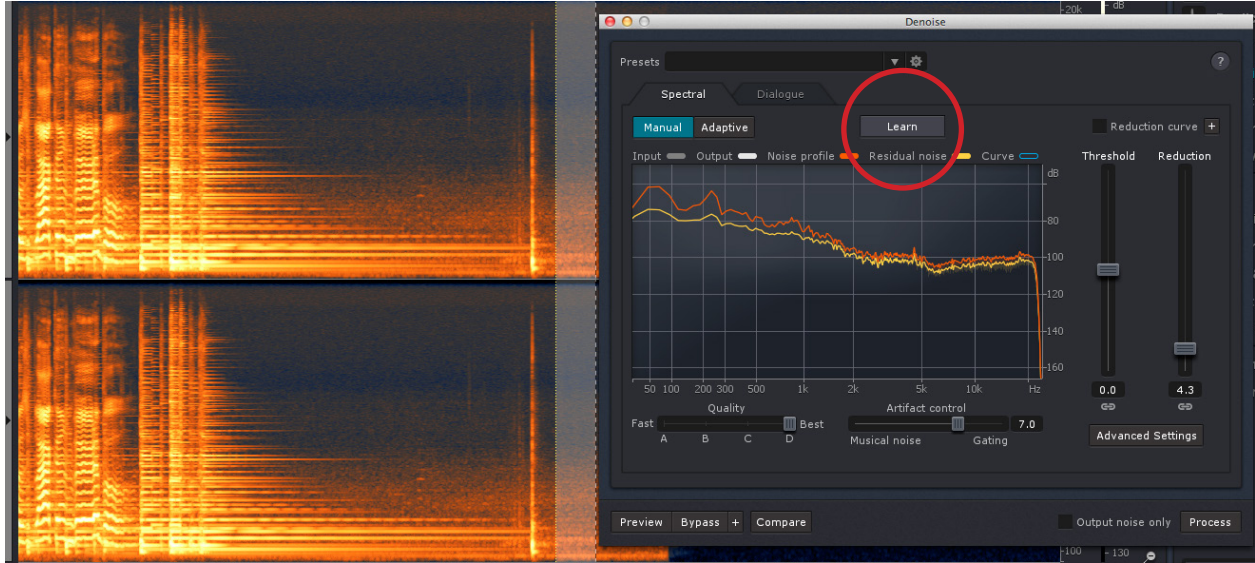
For broadband noise problems, always use a Denoiser. The first step is always to create a noise profile. This may be done manually by isolating the noise, or automatically. A good denoiser offers both manual and automatic modes.



Automatic modes—such as Adaptive in the RX 3 Advanced Spectral Denoiser—listen to the audio and learn the noise profile on your behalf. These modes are useful when there is no easy-to-identify area of noise from which to learn, or for when the noise profile changes slowly over time.

Some dedicated Dialogue Denoisers (such as RX 3 Advanced's aptly named Dialogue Denoiser) typically work in automatic mode, and may sound better when used on dialogue. For now, we'll focus on manual learning, as it's the best solution for a wider variety of audio, including music. We'll also be using the Spectral Denoiser in RX 3.

Select an area of the signal that contains nothing but noise, and click Learn (or the equivalent in your Denoising software). As a general rule, a selection of 100ms should be long enough, but the longer your selection of noise, the more accurately the computer can reduce the noise. We recommend a selection of at least 1s, preferably 2–4s or longer if possible.



Step 2:

Begin to adjust the level of noise reduction and listen closely to the changes that occur in the audio. Some people prefer to slowly increase the amount of noise reduction until it sounds right, and stop once they start hearing artifacts. Other engineers prefer to go hard, and then dial it back until it sounds appropriate. Do whatever works best for you!



RX 3 TIP • Use the Output Noise Only feature in RX 3's Spectral Denoiser to listen to just the noise. If you hear any desirable signal bleeding through, like low-level bursts of dialogue or music, you should reduce the amount of noise reduction.



Levels of between -5 and -15dB of noise reduction are common. If you find yourself approaching -15dB of noise reduction and you aren't satisfied, consider doing a gentler pass of around -7dB. Relearn the noise profile, and perform a second gentle pass of -8dB. You'll end up with around the same level of noise reduction and hopefully less artifacts.

Play close attention to your original audio. The peaks of your waveforms should still be intact, and it should sound untouched. If you hear dulling or artifacts, it's a sign you've gone too far.

RX 3 TIP • Algorithm D in RX 3 utilizes advanced high-frequency synthesis to avoid dulling. In RX 3 Advanced, you can access additional advanced-level controls to further tweak this functionality.

Step 3:

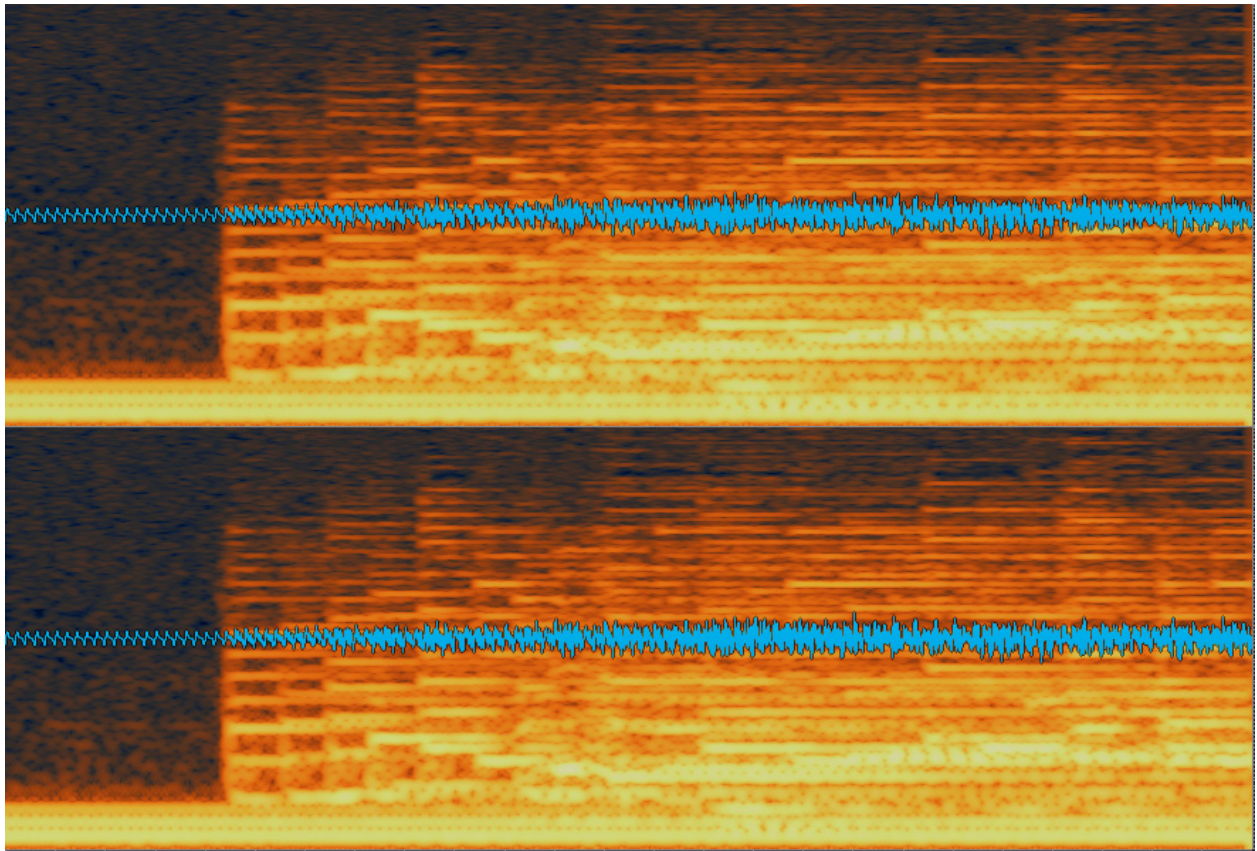
Once your noise reduction is beginning to sound effective, direct your attention to any smoothing filters your tool may offer, such as the RX 3 Spectral Denoiser's Artifact Control slider. Adjusting these smoothing tools can help eliminate any artifacts and preserve the main goal of improving the listening experience. We certainly don't want to make it worse by taking out some noise but adding in artifacts!

HUM AND TONAL NOISE REDUCTION

Hum and tonal noise are usually visible as bright, steady, horizontal lines. Hum and tonal noise look different from each other, and there are different methods available for treating them. Let's begin with hum:

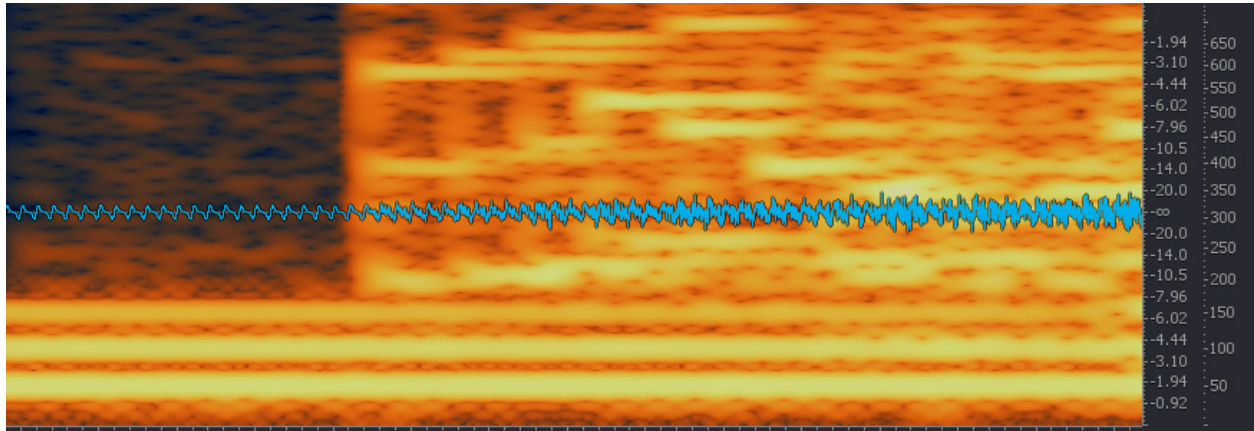
Hum

Below is a screenshot of a recording with hum.

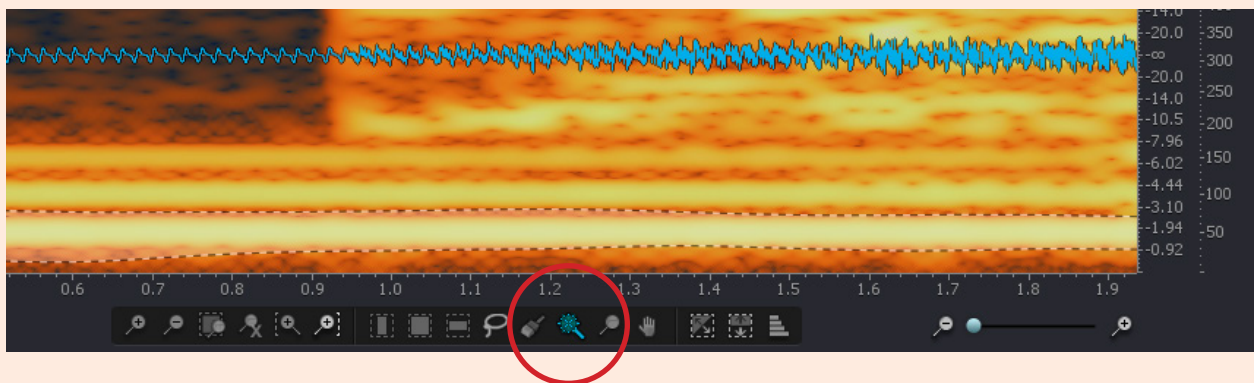


TIP • In order to see which frequencies make up the tonal noise, you must use a spectrogram to view the audio. This can be much more accurate and easy to read than a spectrum analyzer.

It's usually easy to identify the fundamental base frequency, as it's usually the most visible. Often, the fundamental frequency will be 50Hz (Europe) or 60Hz (North America) due to those regions' difference in electrical delivery. This audible hum originates from bad ground connections or inducted power sources such as AC mains or transformers.



RX 3 TIP • The Magic Wand tool in RX 3 can automatically select the fundamental frequency, and double clicking will automatically select all the additional harmonics.

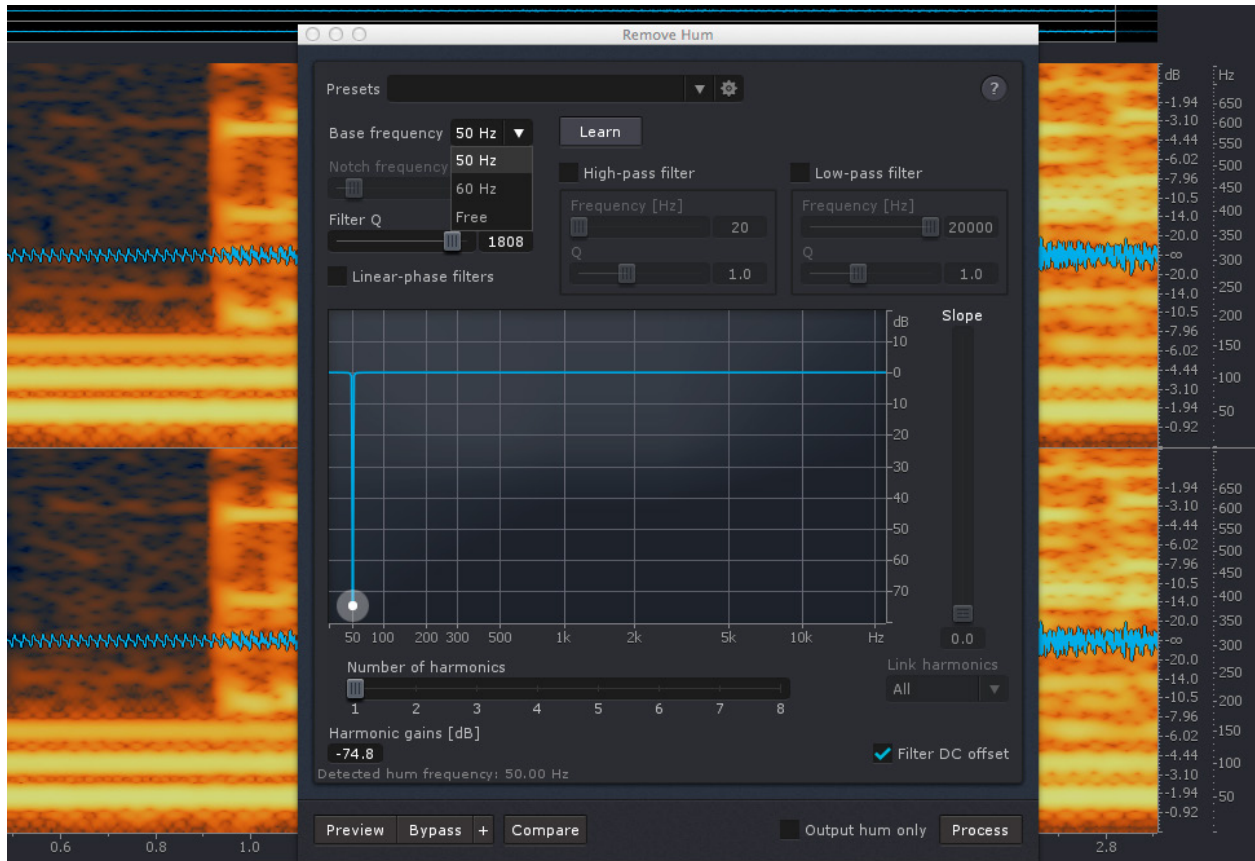


We'll start by using the Hum Removal tool. Hum removal tools, such as the one in RX 3, are extremely precise filters, designed to notch out very specific frequencies. For basic hum with anywhere from 2 to 7 harmonics, this is an effective tool.

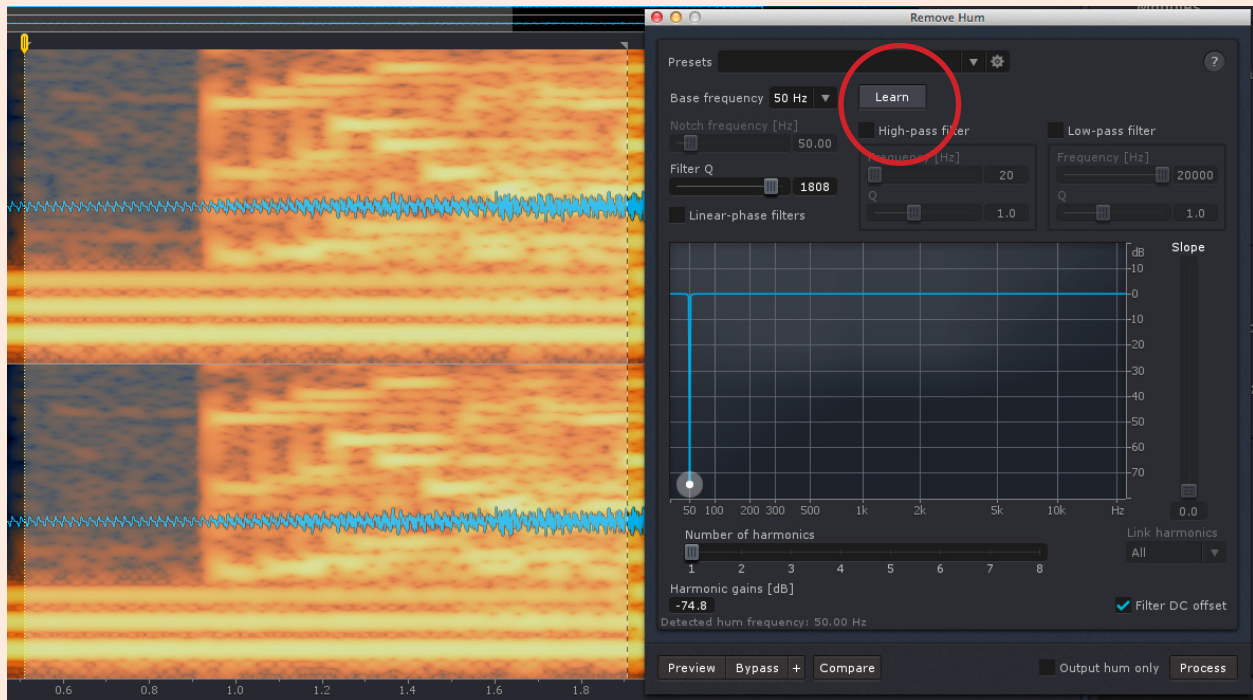


Step 1:

To begin, identify the frequency of your hum. Your ears and the spectrum analyzer module may help. Set the base frequency to 50Hz or 60Hz, depending on where the audio was recorded. If the hum isn't located at 50Hz or 60Hz, you may use Free mode, which unlocks the filters and allows you to set the base frequency yourself.

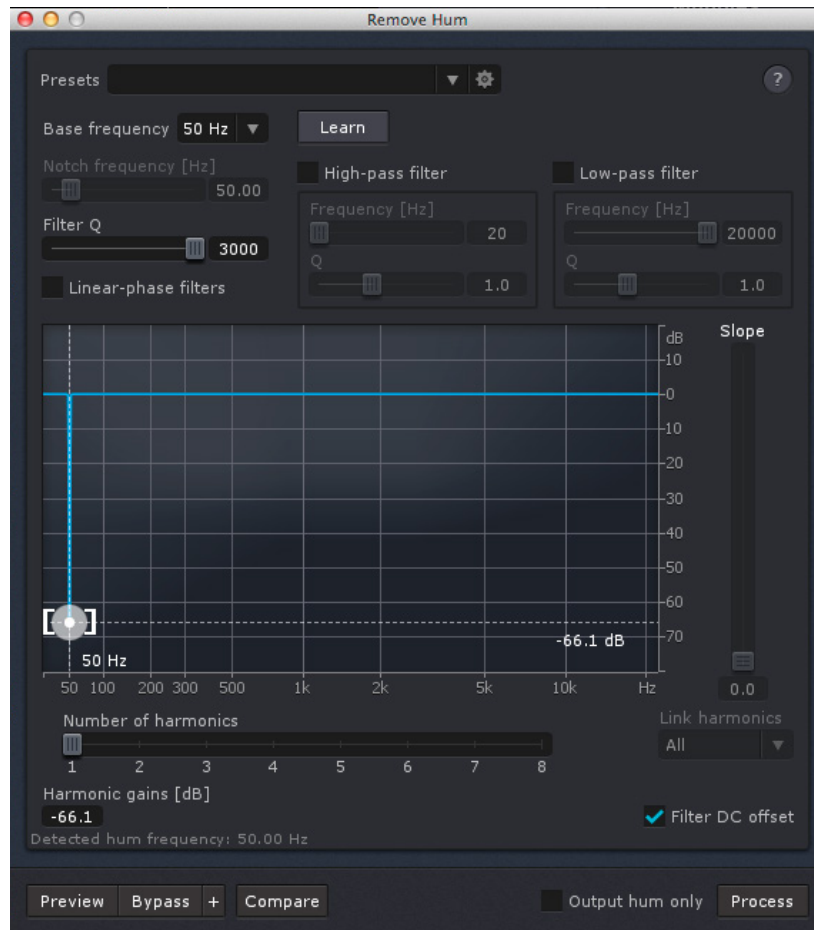


RX 3 TIP • Use the time selection tool to select an area of hum and click Learn. This automatically sets the filters to the correct frequencies based on your selection.

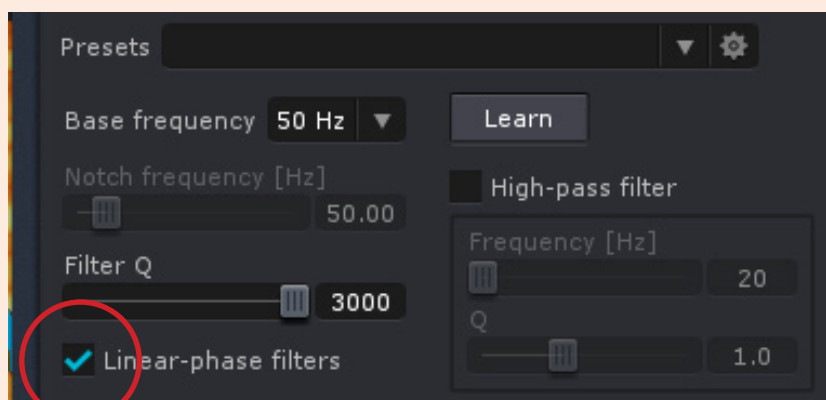


Step 2:

Next, adjust how aggressively Hum Removal attacks the primary frequency by pulling down the first frequency node. The deeper the cut, the more hum will be removed—but you may also adversely affect wanted audio in that frequency range, so tread carefully.

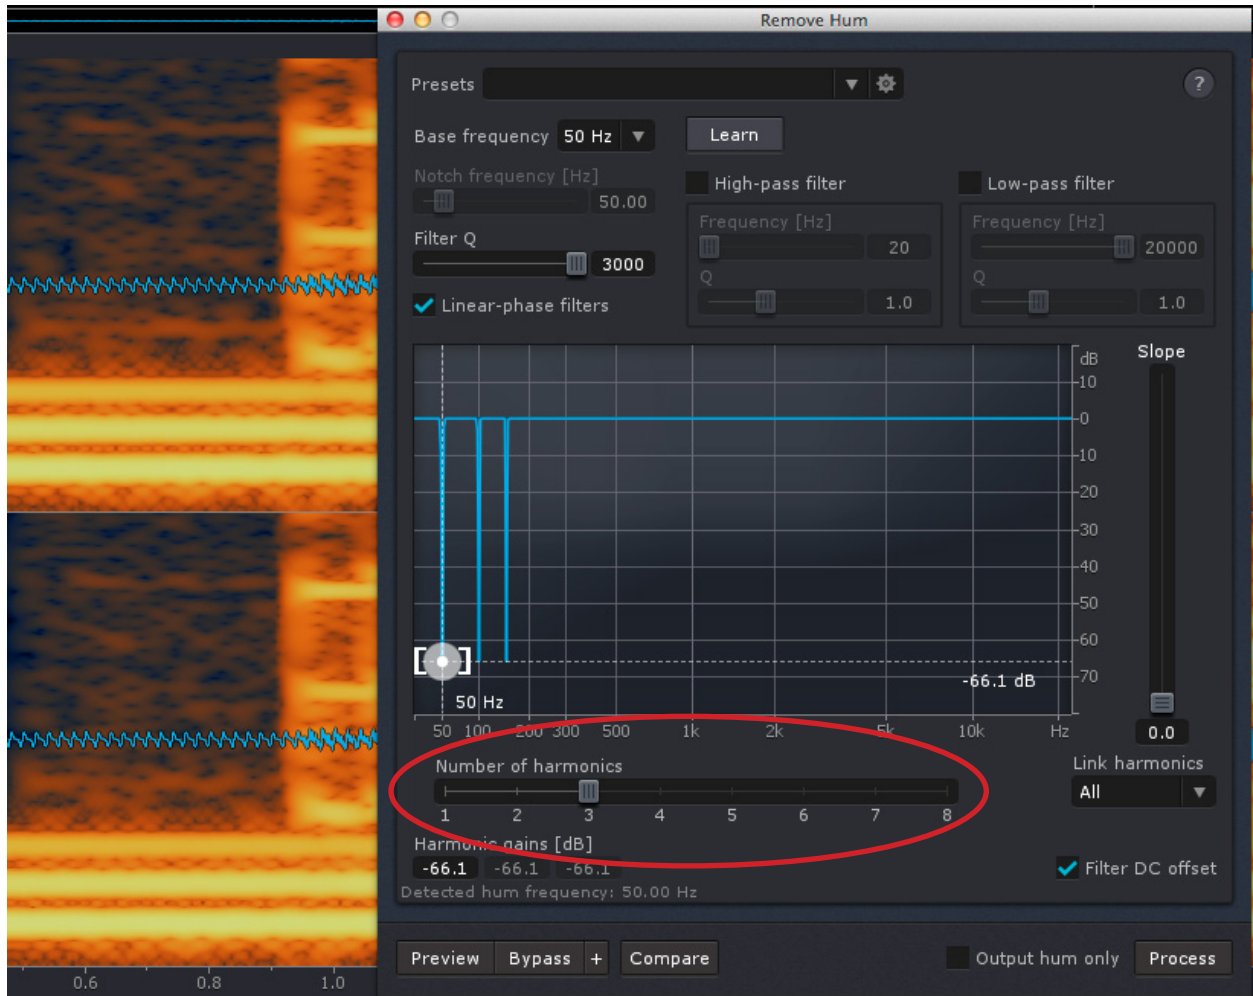


RX 3 TIP • Try selecting the Linear-phase filters check box when you're making very deep filter cuts in full frequency range content. This helps preserve surrounding audio due to the filter's exceptional steepness.



Step 3:

If you can see and hear additional harmonics, start reducing the gain of these as well.




RX 3 TIP • Use the Number of harmonics control to select up to 7 harmonics above the primary frequency. The spectrogram display makes it easy to identify the number of hum harmonics in your project.

Step 4:

As you start reducing the additional harmonics, it's a good idea to ensure you aren't removing any desirable audio. A good hum removal tool allows you to adjust the "Q", or width value of the filters. Generally speaking, narrower filters will provide a better result, as they can more accurately notch out frequencies without affecting the surrounding audio.

RX 3 TIP • Use the Output hum only control to isolate and hear the audio that will be attenuated. If you hear any desirable audio creeping in, narrow the filter Q.

 Output hum only

RX 3 TIP • Filter DC offset removes the DC (direct current) offset caused by the imbalance that sometimes occurs in A/D converters. DC offset is exhibited by the waveform appearing above or below the zero line, and is undesirable because it prevents you from achieving maximum audio levels. DC offset may also cause audible artifacts depending how the signal is being processed.

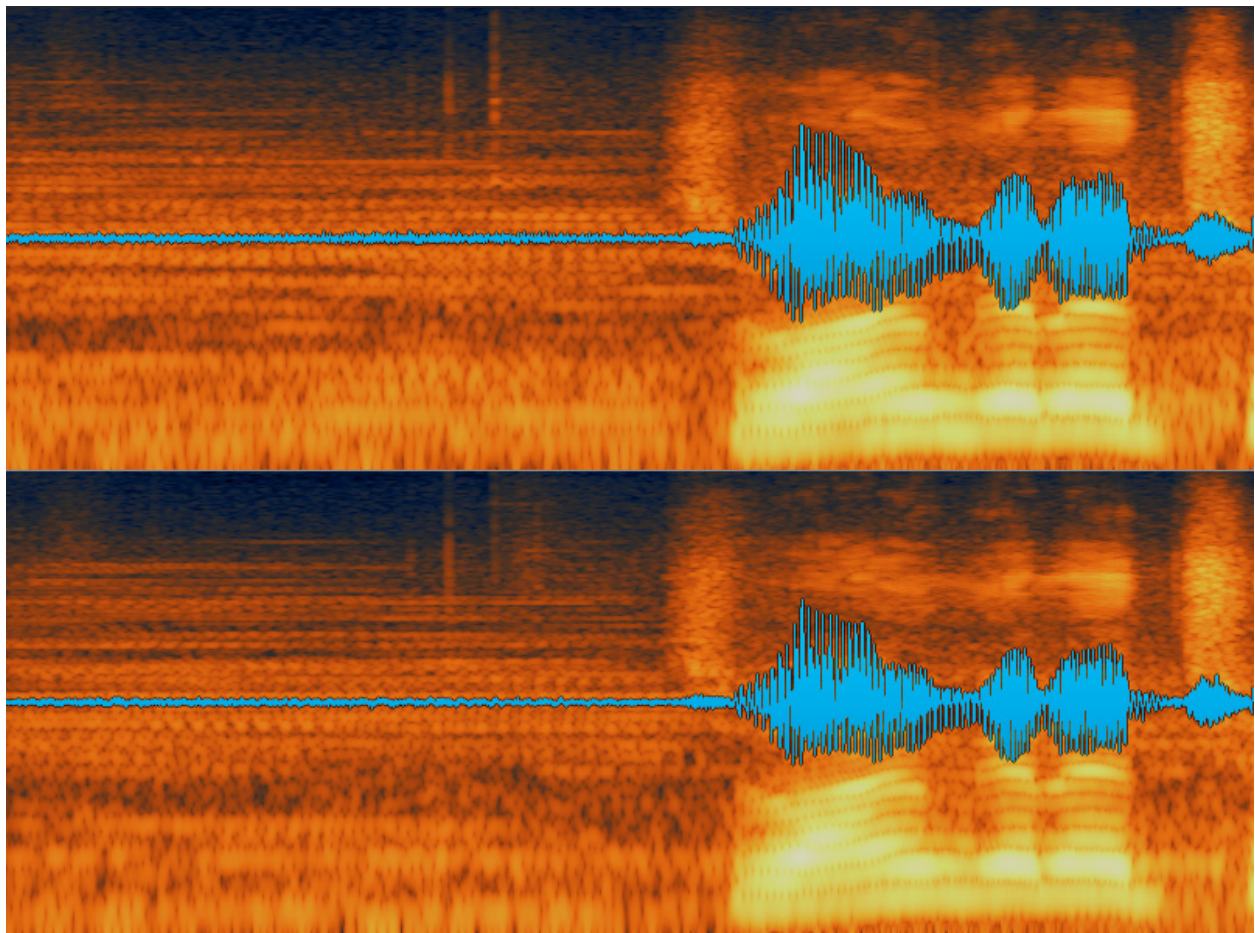
 Filter DC offset

TONAL NOISE

Tonal noise, such as buzz from fluorescent lighting, is more difficult to trace back to a primary frequency and may extend up into high frequencies that are out of the scope of what RX 3's Hum Removal module can handle.

Tonal noise is different from the broadband noise discussed previously in that it tends to be concentrated at certain frequencies. For this, we can use the RX 3 Denoiser module, which gives us separate control over the tonal and broadband noise elements in a noise profile.

Below is a screenshot of a recording with tonal noise. As you can see, there are a lot of harmonics present:

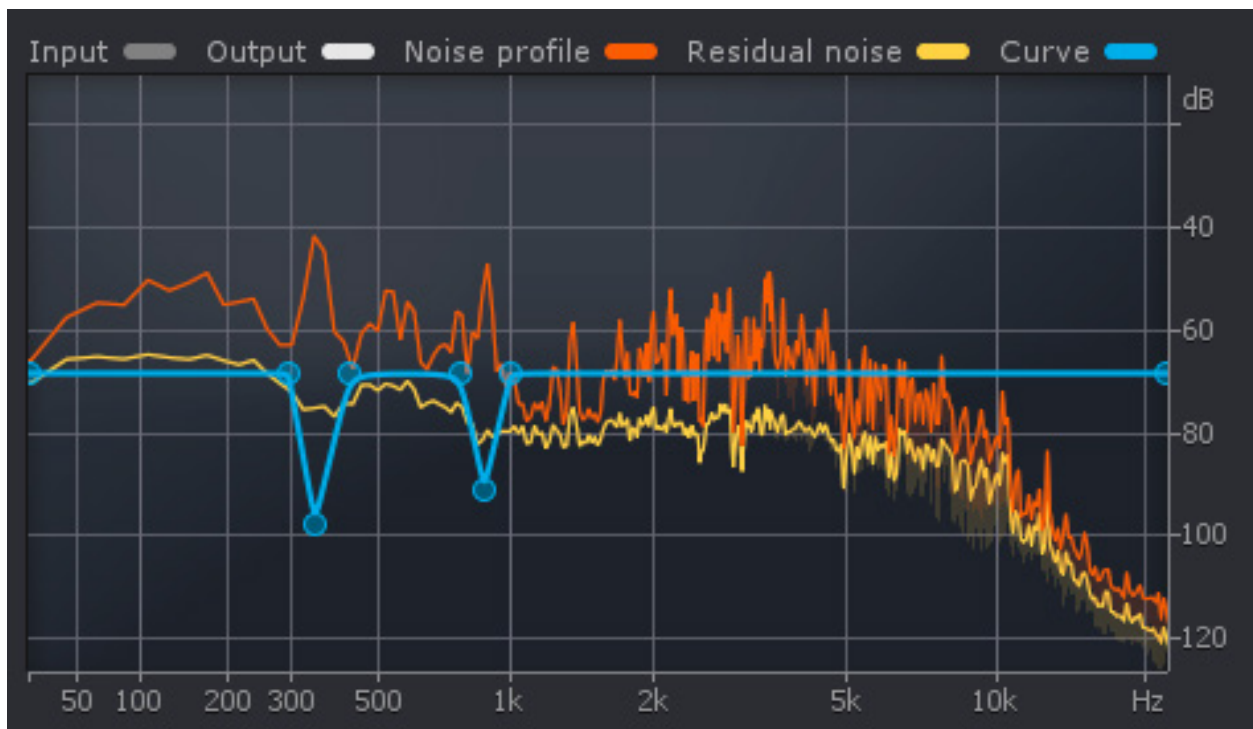


To remove Tonal Noise using the Denoiser, follow the steps for dealing with broadband noise outlined earlier in this Chapter, but with the following additions:

- When adjusting the level of noise reduction in Step 2, de-link the sliders by clicking on the Chain-link icon. Then, increase the amount of tonal noise reduction using the now de-linked Tonal slider. This will start to clamp down more heavily on the tonal elements of your noise profile.



- Make sure the Reduction curve is showing on the noise profile graph. Click the toggle box to show it. Now you may use this reduction curve as an envelope to notch certain areas of the noise profile to be affected more heavily. Where you see tonal peaks, you can direct the Denoiser to be more aggressive in those areas.



DIALOGUE DENOISING

Recorded dialogue may have broadband or tonal noise problems, so why not just treat dialogue with the aforementioned methods? Well, you can—and would likely get a good result!

However, as mentioned in Chapter 3, some denoisers are designed specifically for use on dialogue and vocals. Because of the unique nature of voice recordings—plus our innate ability to hear the minute subtleties of the human voice—regular denoisers that might sound fantastic on music or other audio may not deliver the best possible results on dialogue and vocals. Dialogue-specific denoisers use custom algorithms better suited for vocals.

RX 3 TIP • iZotope RX 3 Advanced includes a dialogue denoiser. We'll take a look at using it shortly.

Another reason for using a dialogue-specific denoiser has to do with workflow. In post-production projects, there might be tens or hundreds of dialogue regions in any given session. It's much quicker and more efficient to use a denoising plug-in in real-time than to destructively edit and then reimport so many files. Dialogue-specific denoisers are low latency and can run in real-time with minimal impact on system resources.

Step 1:

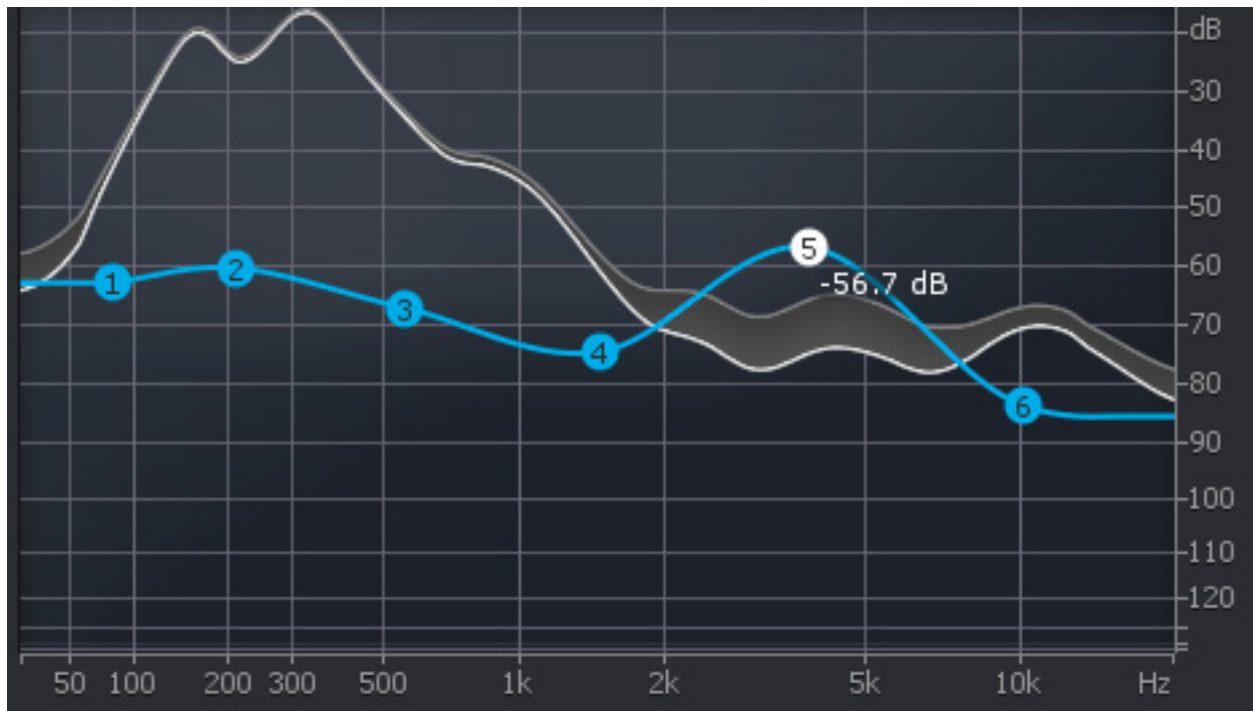


You can use the Dialogue mode in the RX 3 standalone application's Denoiser module, or use the dedicated plug-in in your DAW.

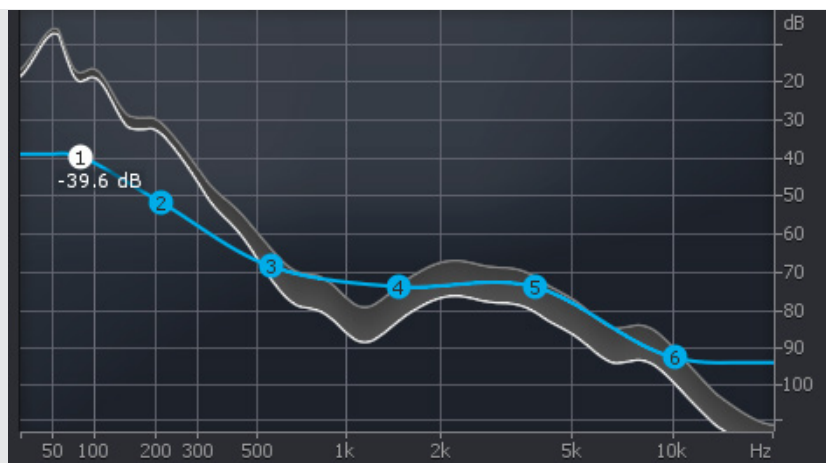
Insert your dialogue-specific denoiser on your audio track, and allow the audio to begin playing. iZotope's Dialogue Denoiser has two modes, Manual and Auto. Make sure you start with Auto.

A good dialogue-specific denoiser will offer multiband functionality, allowing you to quickly adjust the amount of noise reduction applied to different areas of the frequency spectrum. In the Dialogue Denoiser Manual mode, you can adjust these multibands.

This is useful if your audio has a lot of hiss or low-end rumble, for instance.



TIP • Apply gentler processing to the high-mid frequencies, as this tends to be where vocal artifacts are more audible.



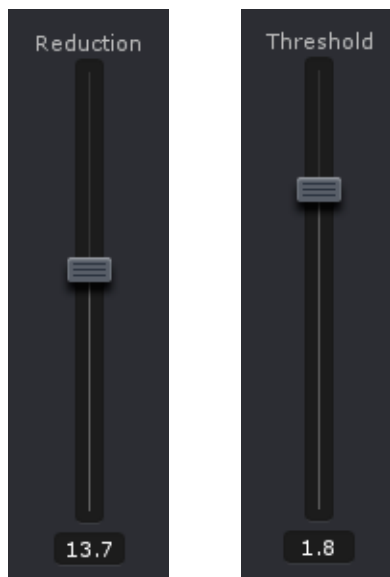
Step 2:

Increase the amount of reduction until you hear the noise begin to disappear. Listen closely, as you'll want to set the amount of reduction slightly below the level at which it begins to affect the voice.

TIP • Remember that a higher value for the denoiser's reduction parameter doesn't mean more noise, it means more reduction (and therefore less noise).

If your denoiser has a threshold control (as the Spectral and Dialogue Denoisers in RX 3 do), adjusting it can help remove more noise with fewer artifacts.

A higher threshold value will reduce more noise, but suppress low-level signal components, so if you go too far, you may start to hear the ends of words being truncated slightly. A lower threshold value helps preserve these low-level signal components, but if you go too low, you may cause noise modulation. It's all about finding the right balance!



GENERAL DENOISING TIPS

1. Softly, softly. Be gentle and do multiple passes if necessary. This will often lead to a better sonic result than one harsh denoising pass.
2. Reduce with caution. The amount of reduction is the most important control, so listen very closely as you adjust it. It's possible to over-compress or over-saturate an audio signal using conventional mixing tools, and it's also possible to overuse a denoising tool.
3. Know your noise. Using the tips outlined above, identify the components of your noisy signal. Broadband noise and tonal noise require different approaches, so treat them accordingly.
4. Learning is good. If time permits and the noise profile is constant throughout, using Manual mode on a carefully learned noise profile can produce better results than Auto.

7: REMOVING INTERMITTENT NOISES AND GAPS

Intermittent noises can include a wide range of intrusive sounds that don't fall into the broadband or tonal noise categories. These can include a cell phone ringing at a violin recital, a door hinge squeaking during an interview, or traffic noise interrupting the dialogue being recorded for a film.

Gaps, dropouts and short sections of corrupted audio are also common audio problems. These can be caused by everything from a loose audio cable to digital errors.

These audio problems tend to occur suddenly and last for a short duration—that's what differentiates them from the conventional noises we dealt with in Chapter 6. We'll refer to the process of repairing these intermittent problems as audio repair.

WHAT'S THE GOAL OF AUDIO REPAIR?

Repairing audio involves carefully patching over troublesome areas and performing precise audio edits—all without causing audible artifacts that listeners may detect.

Audio repair might be performed for an obvious benefit, such as salvaging a live recording that suffers from an audio dropout or sudden and distracting audio event. However, the principles of audio repair can also be used for more subtle edits and enhancements. For example, removing springtime birdsongs from an audio recording intended to be set in the depths of winter, eliminating vocal plosives, and smoothing abrupt transitions in certain frequency areas are all ways in which audio can be polished in the post-production stage.

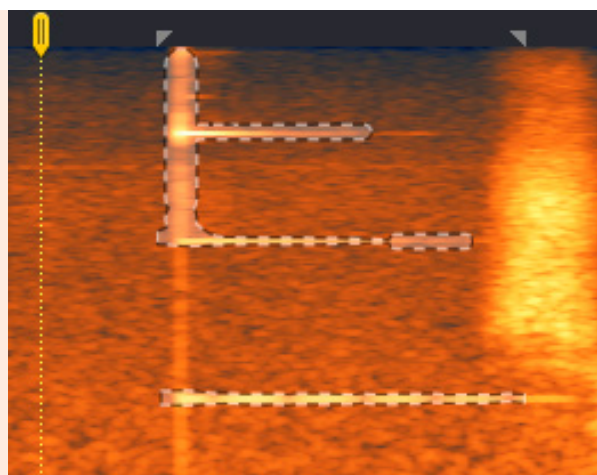
PRINCIPLES OF AUDIO REPAIR

No matter what software you choose, the audio repair process will require using some sort of visual editing tool(s).

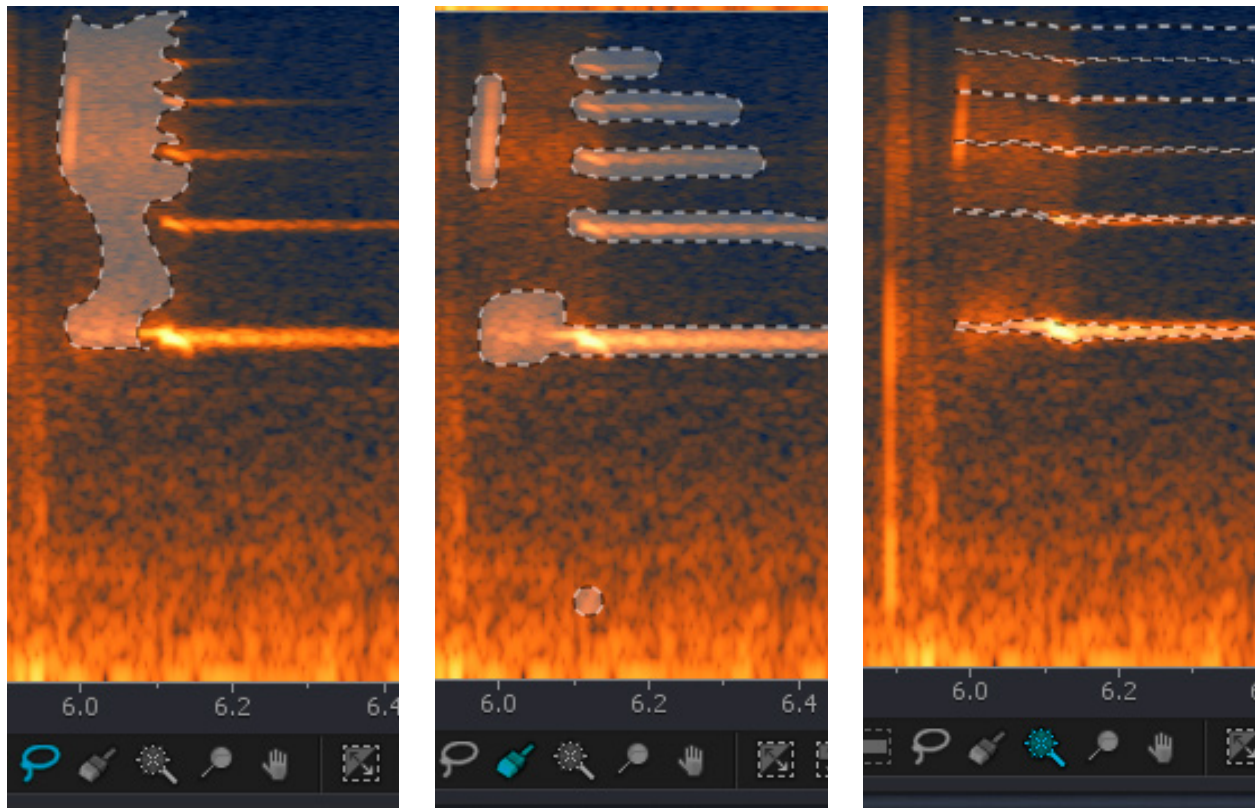
Perhaps the most important is the spectrogram display, which is used to make precise selections of identifiable audio events. These precise selections can then be patched or repaired. If you didn't know how to read a spectrogram display when you started reading this guide, you probably now have a better sense of what they can be used for.

The tools that allow you to interact with the spectrogram are also important. Good audio repair and restoration software should include a number of selection tools, which you can use to draw or highlight specific audio events and frequencies that you observe on a spectrogram.

RX 3 TIP • In addition to the common horizontal and vertical time/frequency selection tools, RX 3 includes several advanced selection tools, such as Lasso, Brush and Magic Wand. These provide much more control over selecting sudden audio events that change frequency and move about.



Terms such as brush and lasso are common across visual editing platforms such as Adobe Photoshop™. Here's what they mean in the context of audio repair:



- **Lasso:** A lasso tool lets you use your mouse to outline a freeform selection of an image.
- **Brush:** A brush tool lets you use your mouse to outline a freeform selection with a defined brush size. The brush size is usually adjustable.
- **Magic Wand:** A magic wand tool lets you automatically and intelligently select a specific audio event within a spectrogram (or certain pixels making up part of an image).

Once you've highlighted certain events, the final step to audio repair will be audio processing.

There are several things that make intermittent noises especially hard to fix:

- They can be wildly unpredictable in frequency and timing.
- Unlike broadband noise, hum, clicks and crackles, noises like this can't be removed with an automated process and can be time consuming to fix.
- Most traditional audio editing tools cannot effectively remove them without leaving many artifacts or damaged audio.
- There's a lot of diversity between the audio repair methods and tools available, so it's worth checking out audio examples and downloading demos to find the right solution for your needs.

USING THE SPECTRAL REPAIR IN RX 3

In this Chapter, we'll explore using the RX 3 Spectral Repair module to remove intermittent noises and fill in unwanted gaps. We'll begin with these algorithms as they work in particularly unique ways. Please note that other audio repair and restoration solutions may not include the following functionality, or might achieve the intended results differently.

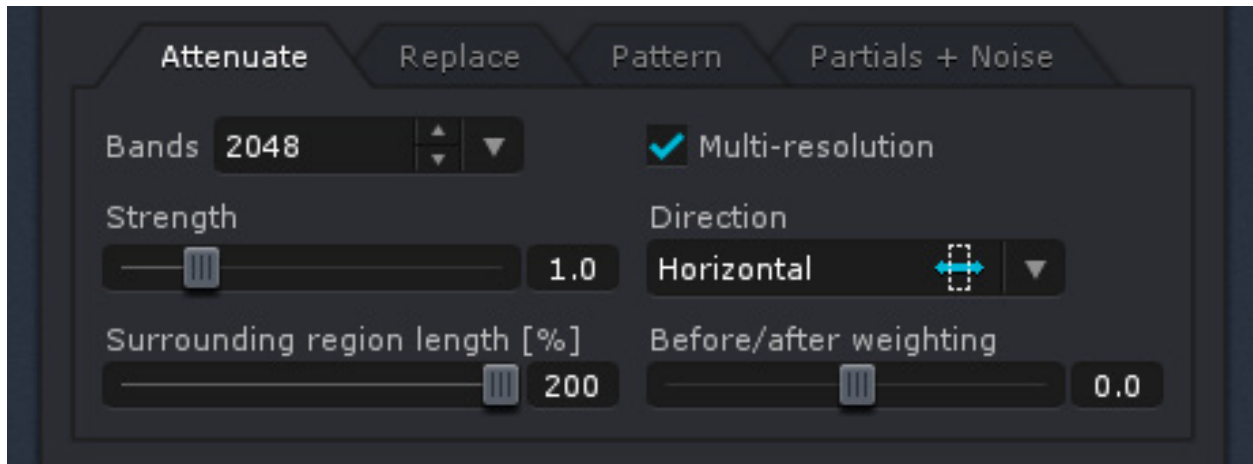
The process of using Spectral Repair requires making precise selections using either the basic or advanced tools. If you're using Spectral Repair as a plug-in inside a DAW, note that it includes a built in spectral editor window that works in a similar way to RX 3's interface.

RX 3 TIP • The Spectral Repair plug-in only works with DAWs that support a technology called offline processing. Pro Tools is one such example.

TIP • We strongly recommend that you try the Spectral Repair examples included in the Appendix, as well as the audio demo files that accompany this guide.

Spectral Repair offers four different algorithms that you can use for audio repair:

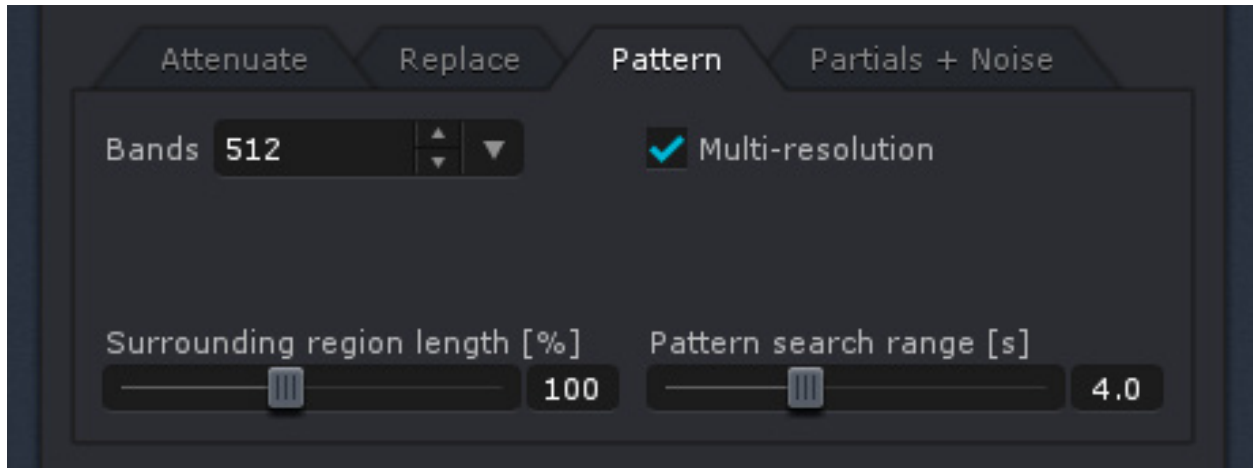
1. Attenuate is an intelligent gain adjustment. It can be used to push unwanted audio events into the background—particularly useful if the unwanted audio events don't completely obscure the desired signal.



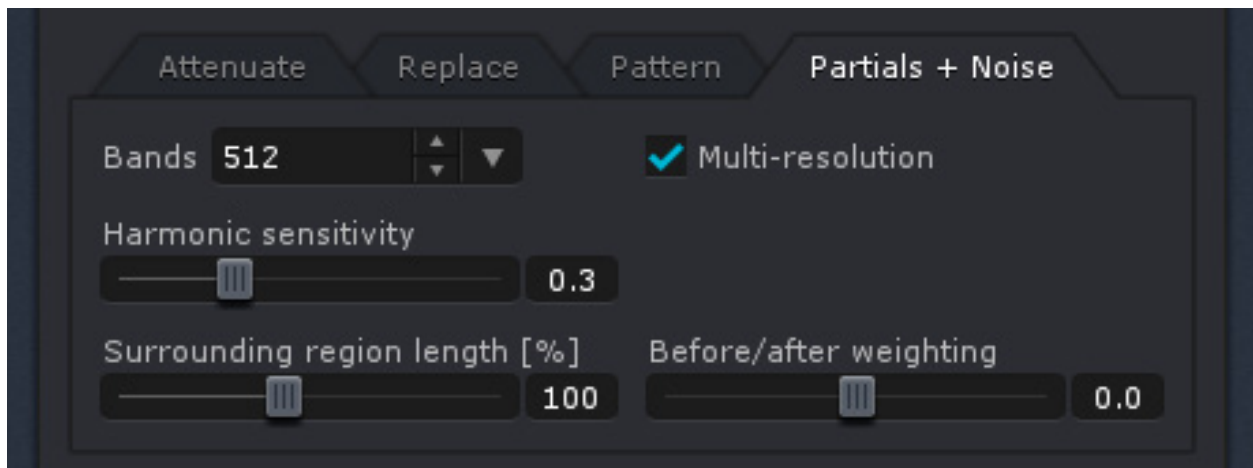
2. Replace is used to replace damaged audio, including entire gaps and dropouts. It can resynthesize audio using the audio information surrounding the damaged area.



3. Pattern is suited to patching badly damaged audio that contains repeated components, such as instrumental vibrato. It's an intelligent copy and paste function that incorporates advanced blending techniques.



4. Partials+Noise is a more advanced version of Replace. It's more effective on heavily harmonic content, and focuses on detecting and resynthesizing harmonics.

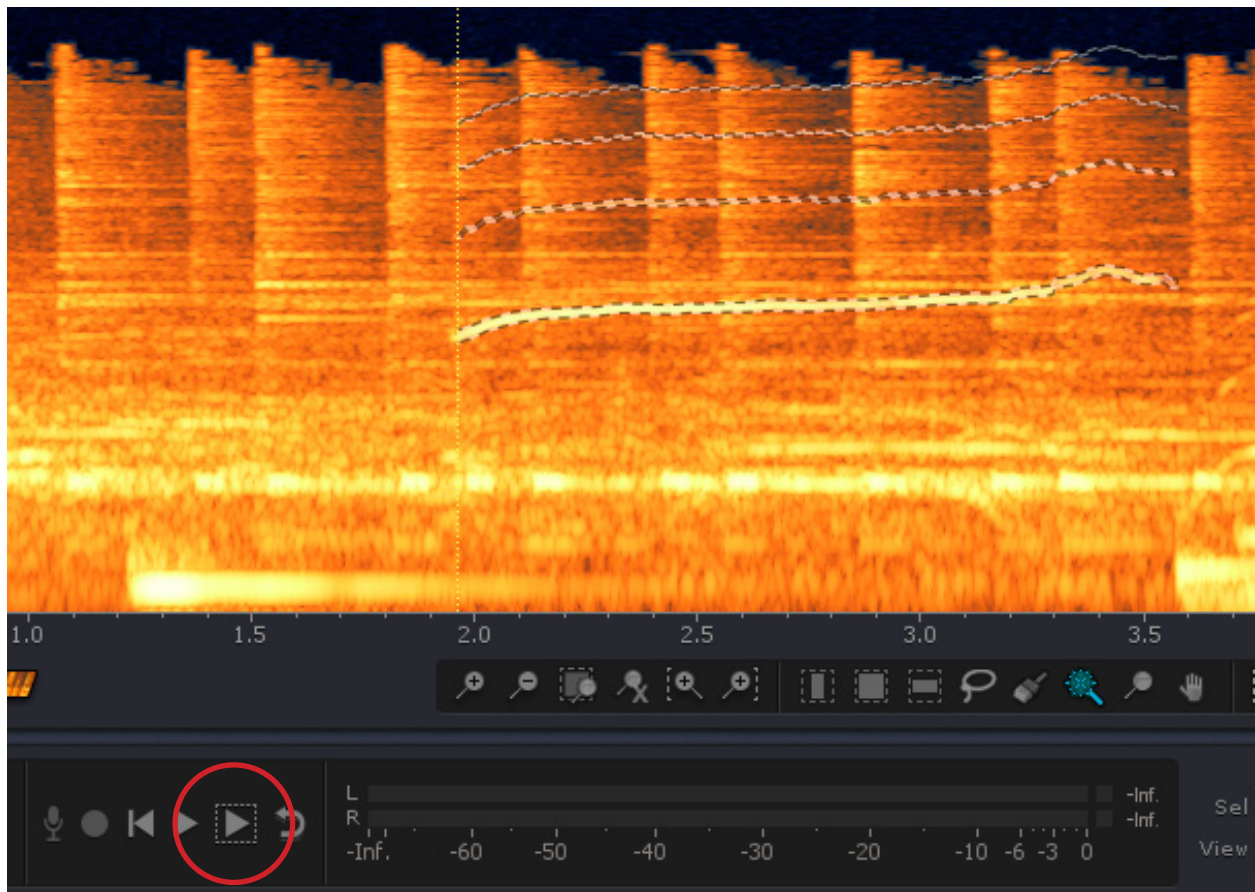


INTERMITTENT NOISES

Step 1:

Begin with identifying unwanted noise in the spectrogram. See Chapter 4 for more details on using the spectrogram to hone in on different types of problematic audio.

Using the selection tools, isolate the noise as precisely as possible. You can either draw a freehand selection around the edges of the unwanted audio, or use the Magic Wand tool in RX 3 to make the selection automatically.



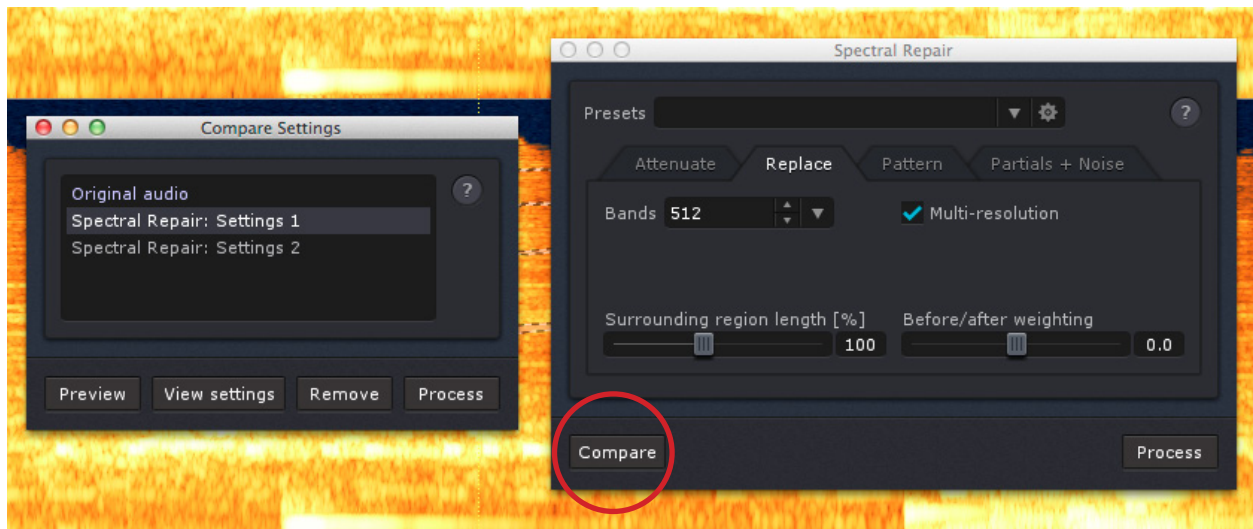
RX 3 TIP • Use the Play Frequency Selection tool to play back the selection of audio in the spectrogram. This helps to determine whether you've correctly set the boundaries for the selection.

Step 2:

Once you've isolated the unwanted noise, open Spectral Repair and ask yourself the following questions:

- Do you want to push the noise gently into the background? If so, use Attenuate to process the selection.
- Do you want to completely remove the noise? If so, use Replace or Partial+Noise to process the selection.

As a general rule, Replace is a good starting point for removing sounds entirely. Partial+Noise might be a better choice if the surrounding audio is heavily harmonic—like a chair squeak from a live orchestral recording.



RX 3 TIP • Using the Compare function will let you quickly A/B the results of your chosen modes and settings whilst experimenting with audio treatments.

Step 3:

Now that you've identified, isolated, and treated unwanted noise, it's time to play back your audio. Listen out for any changes that may have been made to the desirable audio, such as the introduction of artifacts.

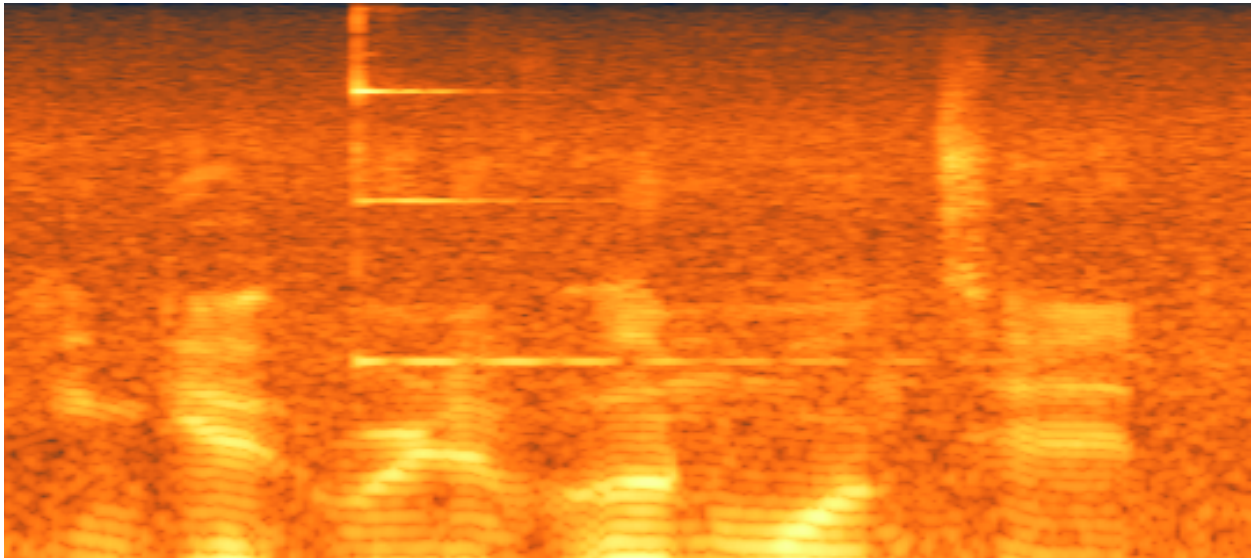
Remember that listeners will only hear what you've left in the mix, and not what you've taken out. If you fail to identify and treat leftover artifacts, they may distract the listener. It's important to make an edit sound as seamless as possible!

Step 4:

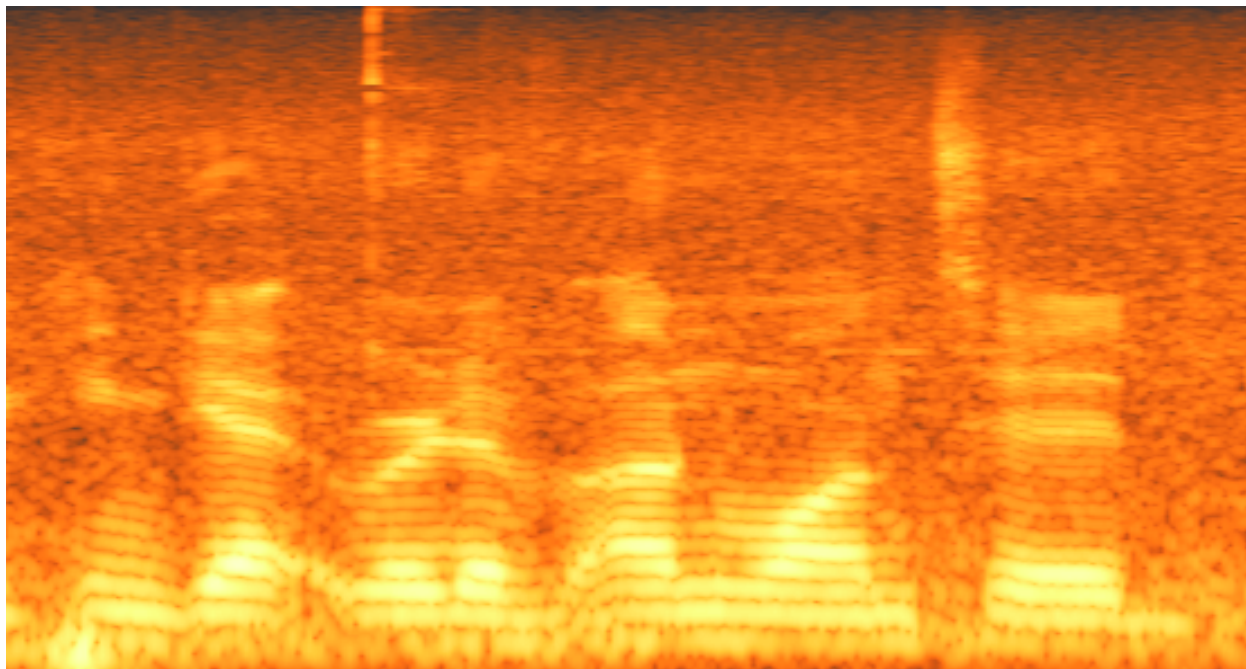
Removing the noise can usually be done in one pass.

In other cases, you might need to do a couple of patch repairs to remove the problem entirely.

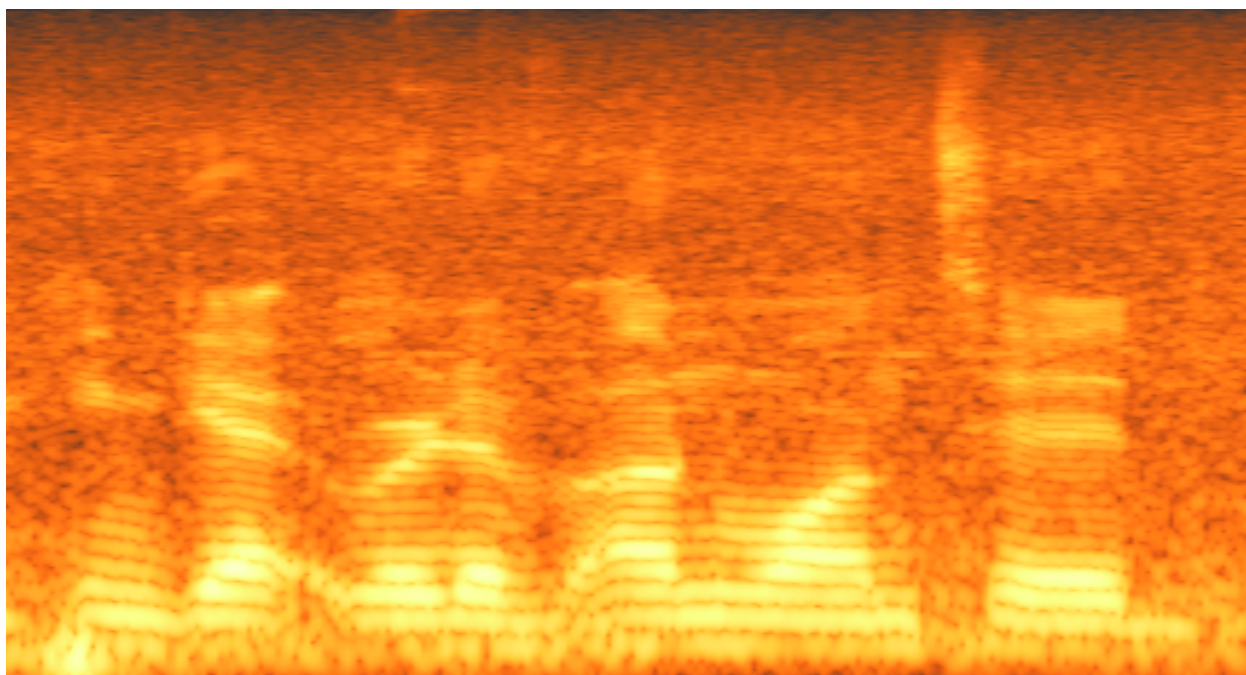
In this image, the unwanted noise is a bicycle bell. You can use one pass to remove the ring of the bell, which is visible as a collection of horizontal frequencies.



In this second pass, you can select and patch the initial attack of the bell, which is a more mechanical sound.



Understanding the different components that make up unwanted noise can help you identify how best to treat it, and whether one or several passes would be most effective. This process is similar to how you identified different types of noise and the ways to treat it in Chapter 6: Denoising.



AUDIO DROPOUTS

Step 1:

Audio dropouts are easily spotted in either a waveform or a spectrogram view.

Working with a spectrogram view does make the repair process easier.

Firstly, you'll want to make sure you have the gap highlighted. A full-bandwidth selection tool is most preferred for this, and you should highlight a very small amount of audio to the left and the right of the gap.

TIP • It's best to select a small amount of audio to the left and right of the gap as clicks will be present. This is a result of the waveform being interrupted between zero crossings. This technique also helps hone in on material that will later be replaced.

Step 2:

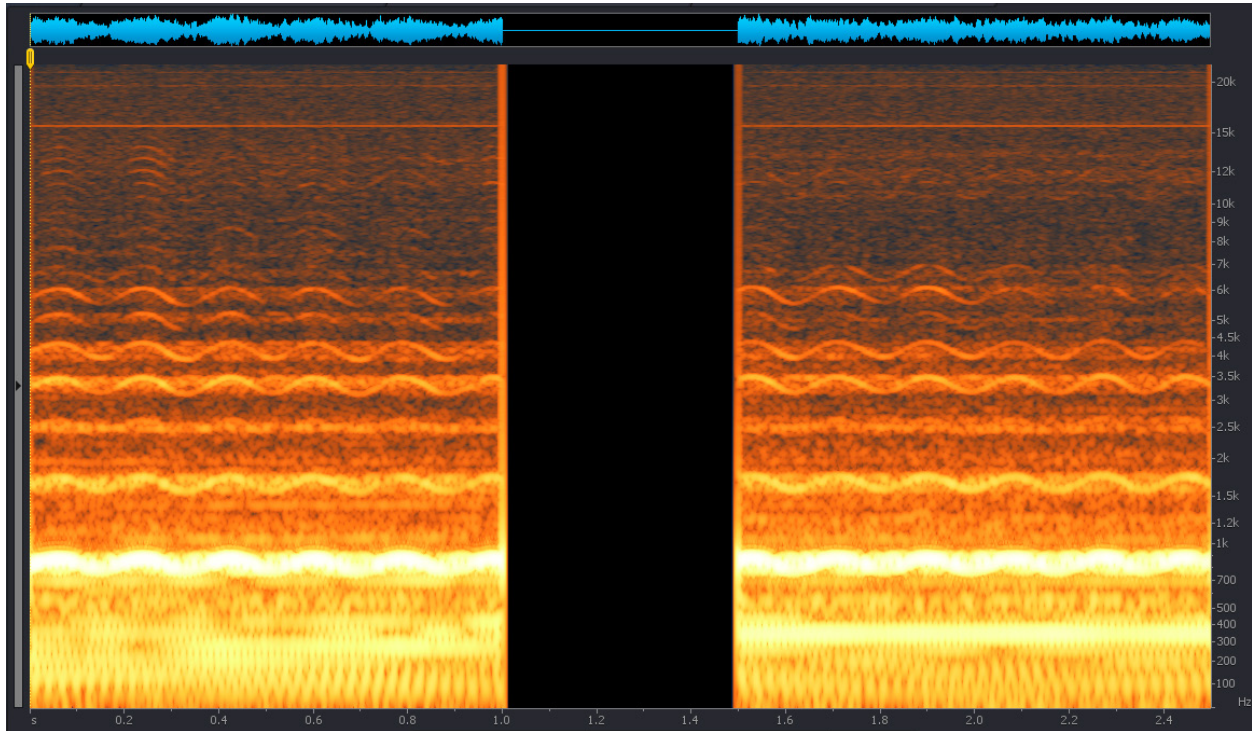
Once you have the audio gap highlighted, ask yourself these questions before processing:

- Is the audio fairly steady, with low harmonic content? If so, use Replace to process the gap.
- Does the audio have a repeating pattern, such as a vibrato? If so, use Pattern to process the gap.
- Does the audio seem suited for treatment with Replace, but is much more harmonically intense (like music)? If so, use Partials+Noise to process the gap. Replace always connects harmonics from two sides of the gap horizontally, while Partials+Noise can connect frequency-varying tones, more common in music.

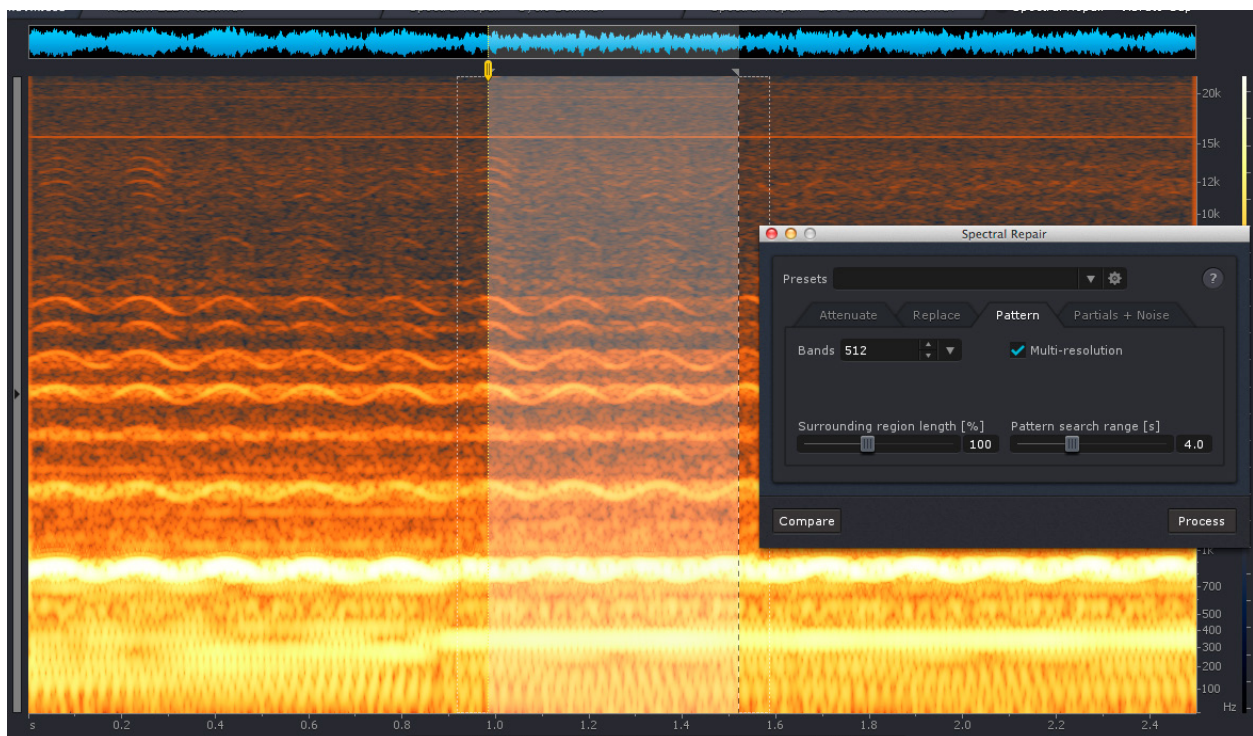
Step 3:

As mentioned above, listening back is important—particularly when performing repair on full bandwidth audio dropout.

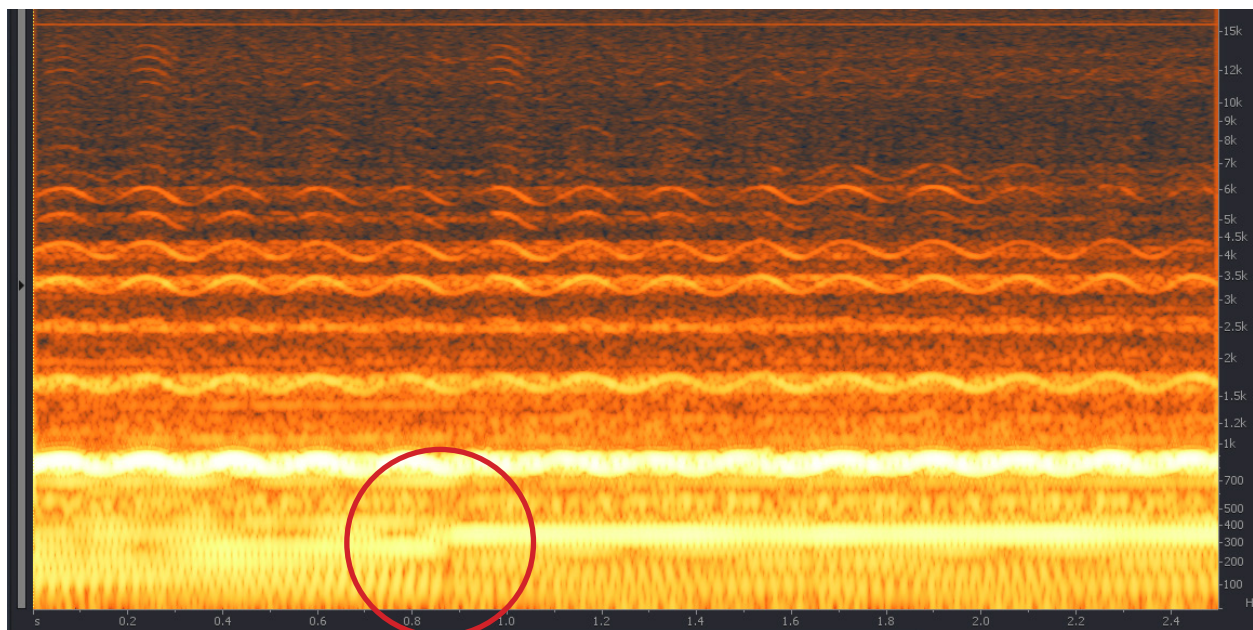
Repairing the audio dropout may work the first time, but sometimes you'll need to do a couple of passes.



For instance, in the above image, we have an opera singer and an orchestra. Pattern was used to replace the dropout, resulting in this:

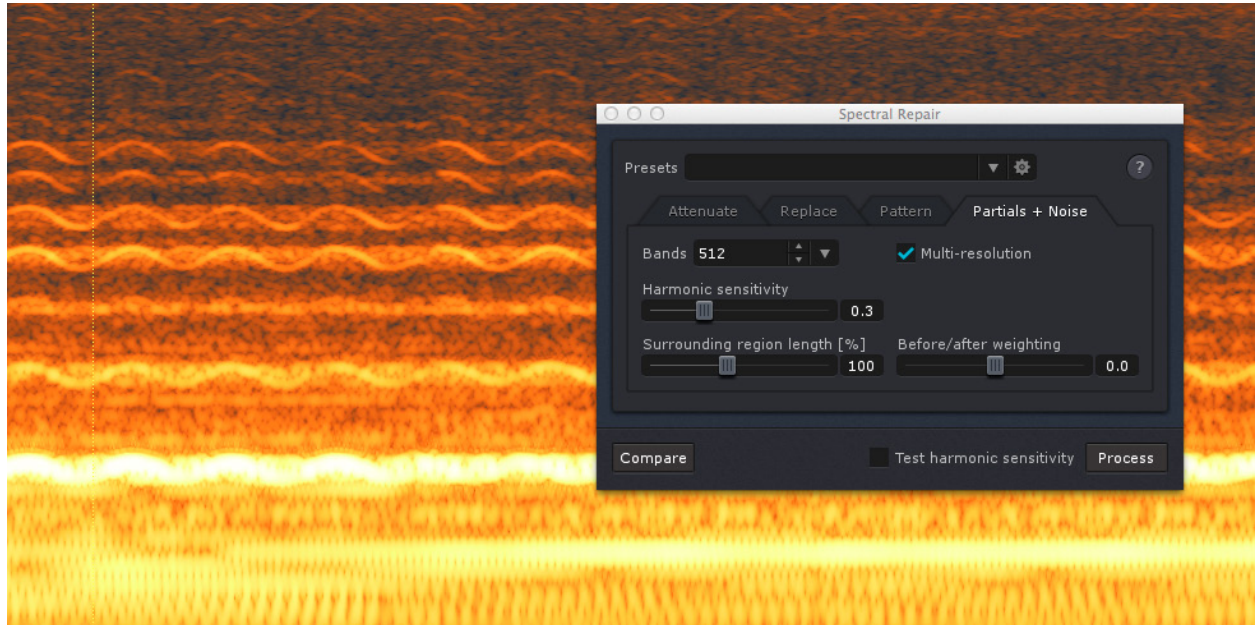


Although this may sound good, we can see that the slowly evolving orchestral chord has been repaired abruptly.

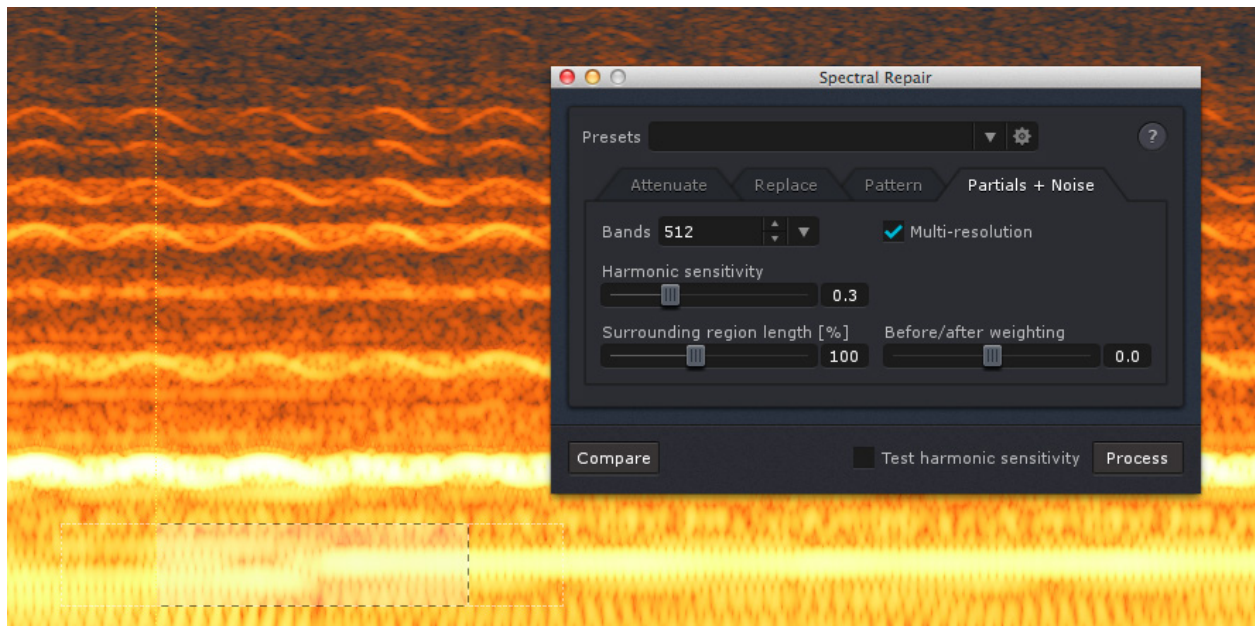


In situations where a whole signal does not transition smoothly, it's advisable to perform a second pass.

Target the specific frequency area and process it with the appropriate algorithm. In this case, you should use Partials+Noise to replace the harmonic content.



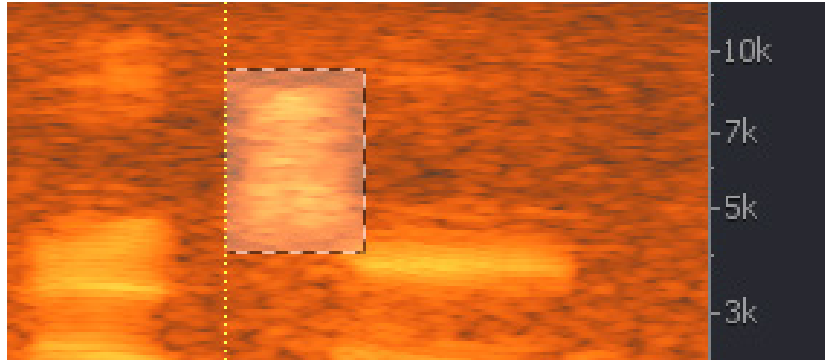
In this image, the audio dropout has been repaired and some additional steps have been taken to ensure a seamless edit.



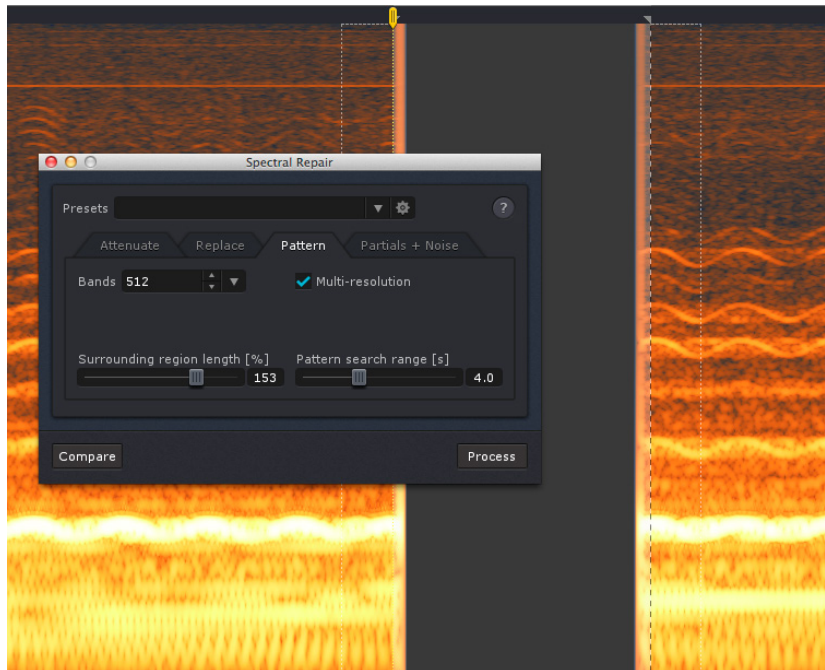
GENERAL AUDIO REPAIR TIPS

1. Trial and error is good. Although you'll likely get great results the first time around, you'll become more effective at using Spectral Repair the more you experiment and work with it.

2. Provide more information. If you fail to get a perfect result the first time, simple tricks like extending the surrounding region length parameter may help.



3. Look around you. If a noise only takes up a certain area of the frequency spectrum (between 5–8 kHz, for instance), you might get a more transparent result performing area-specific spectral repair, rather than full bandwidth repair. This is because the audio surrounding the area being patched (anything below 5kHz and above 8kHz in the above case) is left unaffected and will help smooth over the transition.



4. Keep it simple. For example, you might find that plosives are better treated using the Gain module by simply turning the gain down to zero. But most times, Spectral Repair is the answer—its default parameters are well designed and are effective without too much fiddling!

8: REMOVING CLICKS AND POPS

Clicks and pops can occur at almost any stage of the recording process. They can be caused by surface noise from mechanical media (especially discs), static electricity, power lines, cell phones, mouth noises, inadvertent physical contact with a microphone and bad audio connector cables. Occasionally, digital errors will result in unexpected clicks too.

WHAT'S THE GOAL OF REMOVING CLICKS AND POPS?

In the case of old analog playback mediums such as vinyl or shellac records, clicks and pops are understood to be commonplace. It's extremely difficult to obtain a perfect analog to digital transfer of an old record without including clicks and pops. In this example, the goal is to remove the extraneous clicks and pops that might distract the listener, and still retain the true character of the original audio.

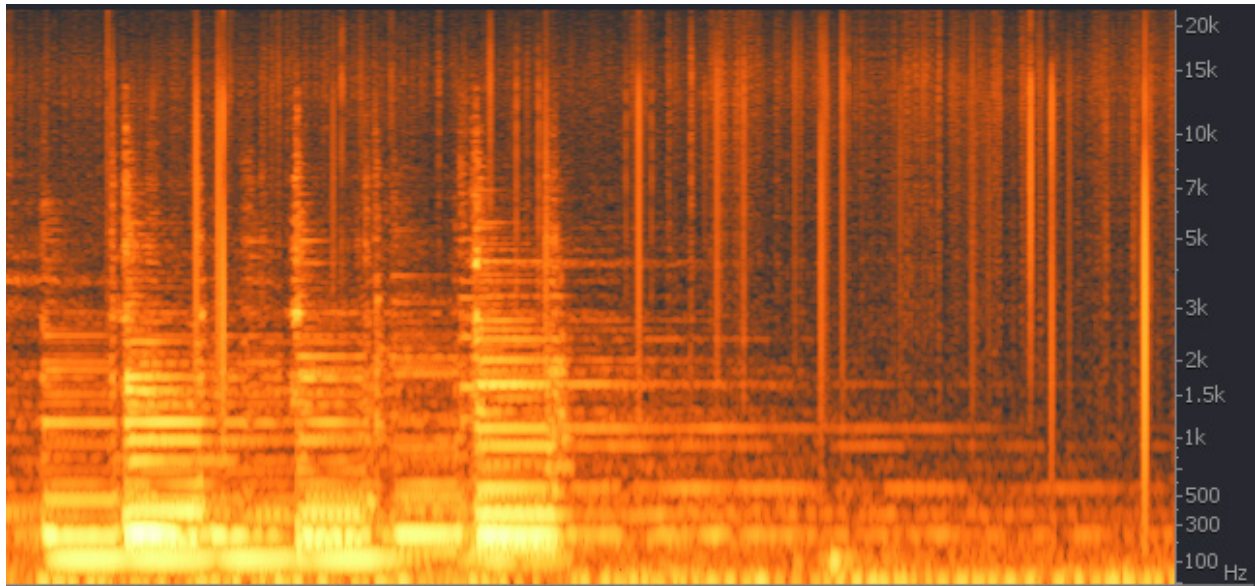
TIP • Many are nostalgic about the sound of vinyl. For this reason, there exists dedicated software that makes digital audio sound as if it were playing from an old record. iZotope Vinyl helps achieve this sound, and was iZotope's first ever plug-in—released in 2001. It's available for free from the iZotope website, at www.izotope.com/vinyl

Clicks and pops that are caused by digital errors or interferences have the potential to render audio useless. Listeners are far less accepting of clicks and pops caused by such errors, as they not only sound different, but are not expected in commonly pristine digital audio.

In the case of mouth noises, removing clicks helps improve the general sonics of a vocal recording. Professional vocal talents are often highly skilled at controlling mouth noises, but clicks from the mouth are usually a perennial problem in dialogue editing.

PRINCIPLES OF REMOVING CLICKS AND POPS

It's not possible to remove clicks and pops effectively with conventional tools,, nor is it possible by notching with a precise EQ. This is because clicks usually cover a wide frequency bandwidth and are extremely momentary.

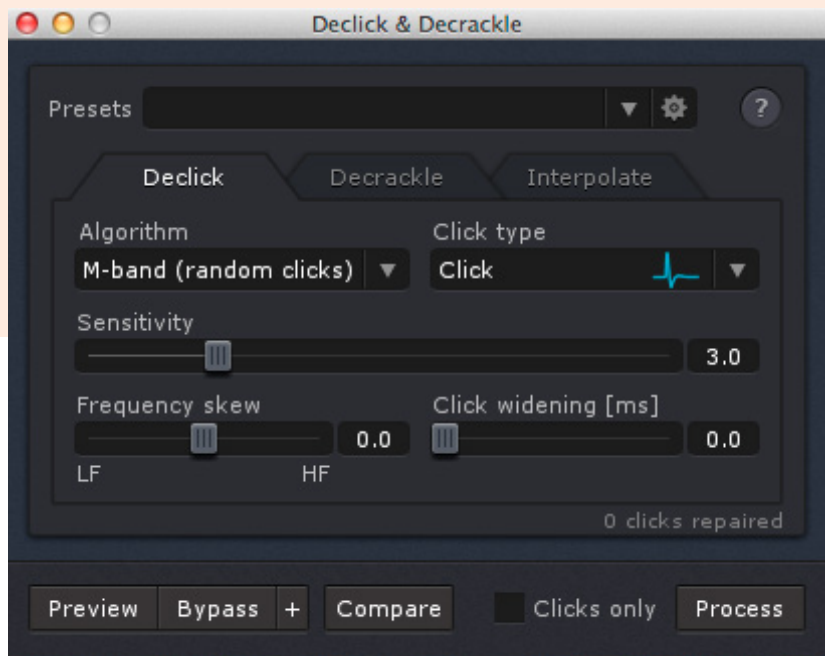


Declicking tools are designed to identify and recognize the sonic signature of a click, enabling the user to then attenuate or remove it entirely, either one by one, or by processing the entire audio file at once. Some declickers may have a single mode, but a good declicker should have several different modes optimized for dealing with the variety of clicks mentioned above.

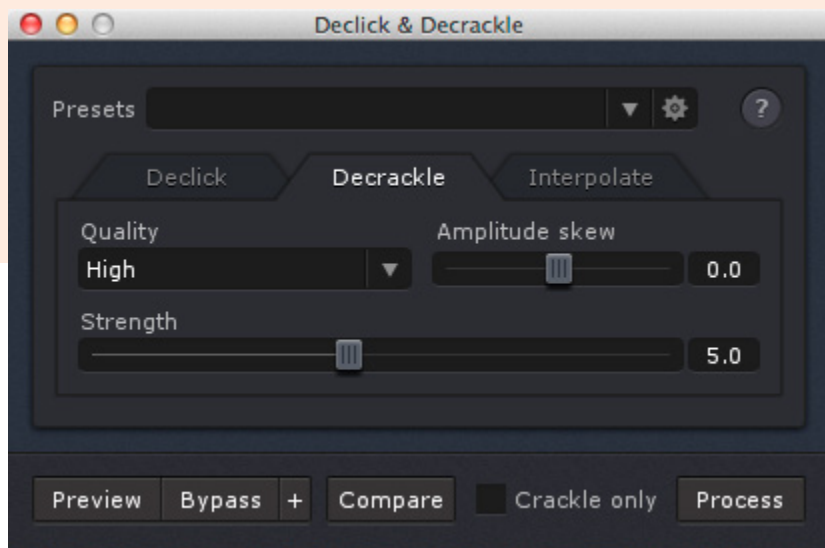
USING DECLICK & DECRACKLE IN RX 3

Chapter 7 covers using Spectral Repair to eliminate certain clicks and pops, but for more troublesome distractions, an even better tool is available for use. The Declick & Decrackle module in RX 3 repairs and reduces clicks, pops and other impulse noises within the waveform. It has three modes: Declick, Decrackle, and Interpolate.

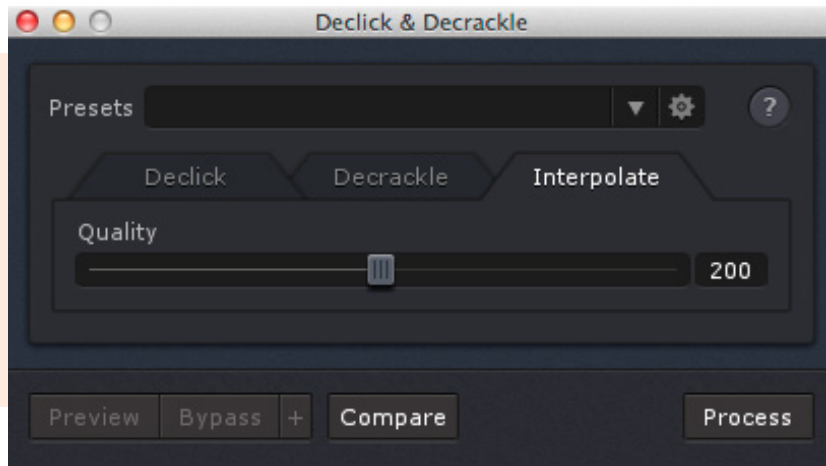
RX 3 TIP • Declick automatically reduces the vast majority of clicks. There are several options in the RX 3 Declicker that combat these, and we'll explore them all shortly.



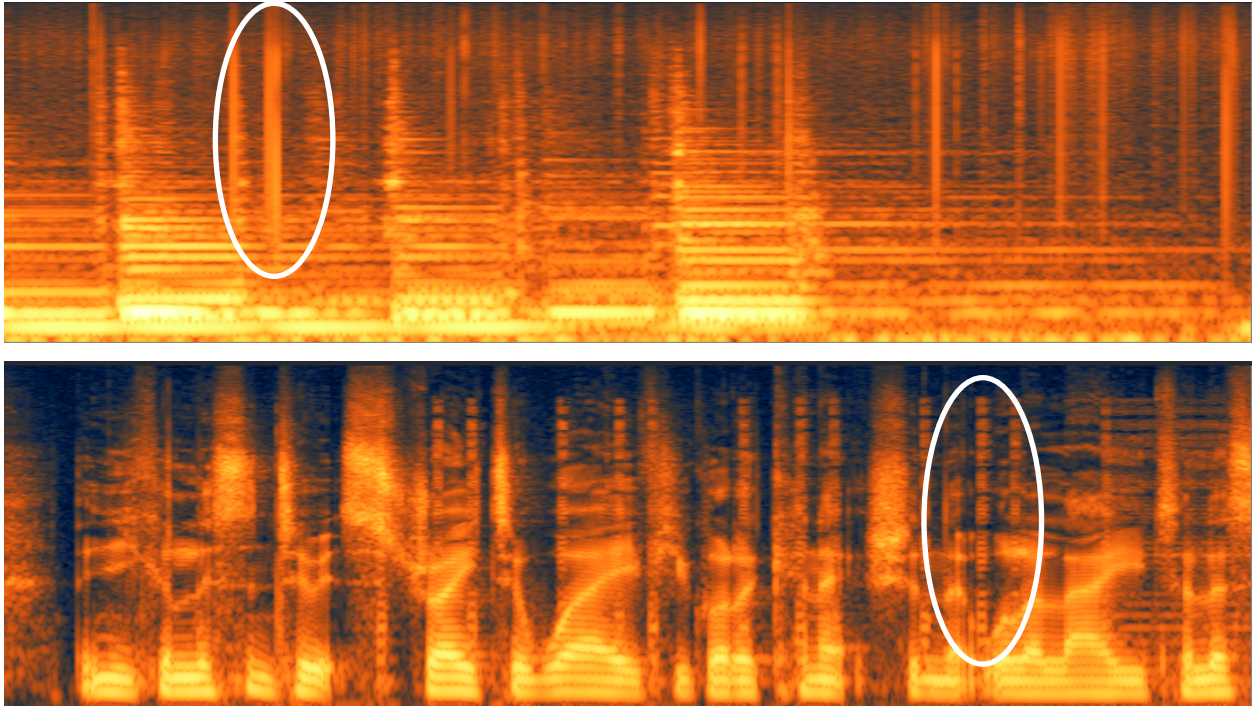
RX 3 TIP • Decrackle removes repeated, quieter clicks that blend together to form what we perceive as crackle.



RX 3 TIP • Interpolate performs one-off manual click repairs that Declick is unable to treat. Interpolate can also be used to fix very slight audio dropouts.

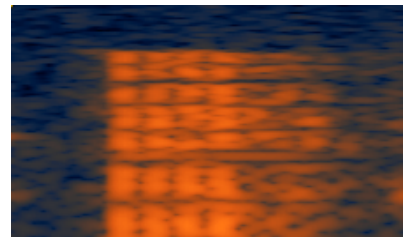


Clicks occurring in the analog domain, whether caused by the surface of a mechanical playback medium or a mouth, are sonically and visually different to digital clicks.



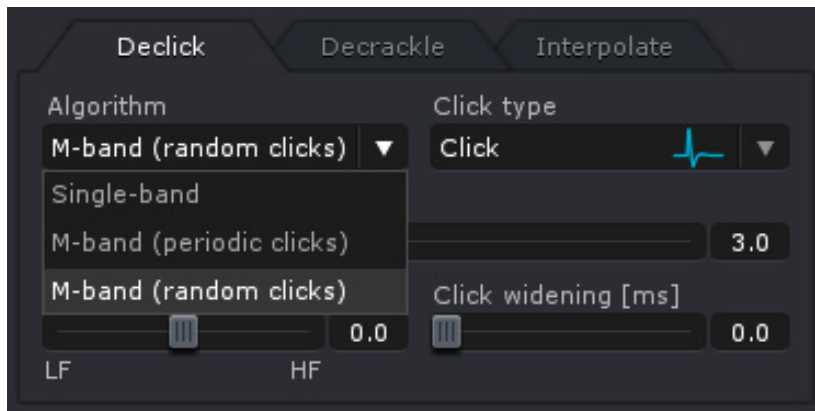
The first image demonstrates a click that occurs during playback of a vinyl record. It's a solid frequency event on the spectrogram, and can be considered random.

The second image demonstrates a click caused by cell phone interference. When zoomed out, it looks similar, but zoomed in, we can see that the single click is actually a very short, periodic signal. The RX 3 Declicker features both a Random and a Periodic mode that you can use to tackle these problems respectively.



ANALOG CLICKS

Step 1:

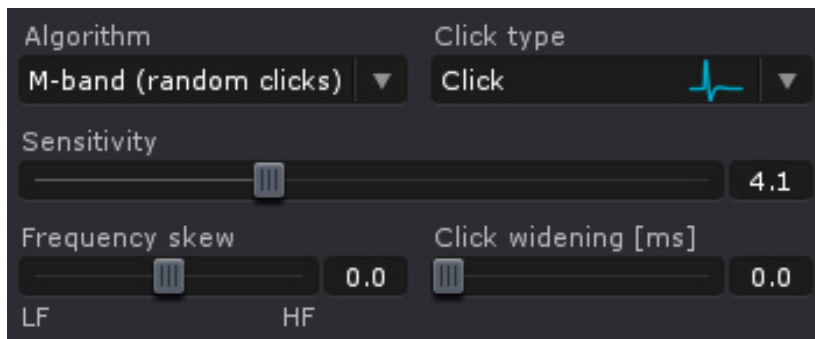


Open Declick & Decrackle. Select the Declick tab, and then select M-band (random clicks) mode.

“Click” will be the click type selected by default. If you believe that the click is being caused by a low-end thump, you should select Thump as your click mode instead. The third option, Discontinuity, is ideally suited to clicks caused by bad audio cuts that have missed a zero crossing.

Step 2:

Before adjusting the strength beyond the default settings, click Preview to hear the result of the processing on your audio file. Many declickers will offer an appropriate level as a starting point.



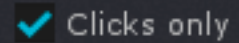
As the audio plays back, you can adjust the strength of the declicker up or down. Higher values will reduce more clicks.

If you're working with single clicks on a one by one basis, process each and listen back regularly. You can always undo any processes that don't work effectively.

Step 3:

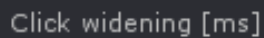
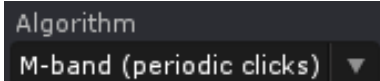
After choosing an acceptable level of click reduction, listen carefully to ensure you aren't unintentionally harming any transients. Using a declicking tool on its highest settings can audibly soften transients. If you can hear that the transients are losing their edge, turn the strength of the declicker down a little.

RX 3 TIP • Use the Clicks Only feature in RX 3 to listen out for the isolated clicks. If you hear desirable audio creeping in, your settings are likely too harsh.



DIGITAL CLICKS

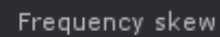
Repeat the same steps used to treat the analog clicks, but instead use M-band (periodic clicks) mode.



0.8

RX 3 TIP • Use Click widening to expand the click size if your clicks are not being fully recognized and removed by the Declicker.

RX 3 TIP • Using the Frequency skew can achieve better results if a click is focused specifically towards the high or low end. Generally, however, this is best left in the middle.



3.0

LF

HF

GENERAL TIPS FOR REMOVING CLICKS AND POPS

1. Two is better than one. As with denoising, there are scenarios where two processes of any declicker are better than one, particularly if you're dealing with an old recording full of unwanted clicks. Performing one process of declicking will remove the most obvious clicks, allowing the second process of declicking to remove the quieter ones.
2. When is a click not a click? The simple answer is: when it's a transient. Overly harsh declick settings can harm the attack of transients within your audio, particularly instruments such as brass, acoustic guitar and percussion. It's important to listen out carefully for this, and dial the declicker back if it's occurring.
3. Choose your battles. It's easy to process an entire file with only one problematic click. But there's no need to risk the fidelity of good audio if it only contains a small number of clicks. For infrequent click problems, treat them individually before moving on.

9: REMOVING CLIPPING

Live concerts and on-location interviews frequently become victims of signal overload or clipping. This can happen in both the analog and digital domain, as well as during the A/D process. This is often the result of time constraints in setting up and sound checking. In music recording, overly enthusiastic singers and drummers are frequently the worst offenders!

WHAT'S THE GOAL OF REMOVING CLIPPING?

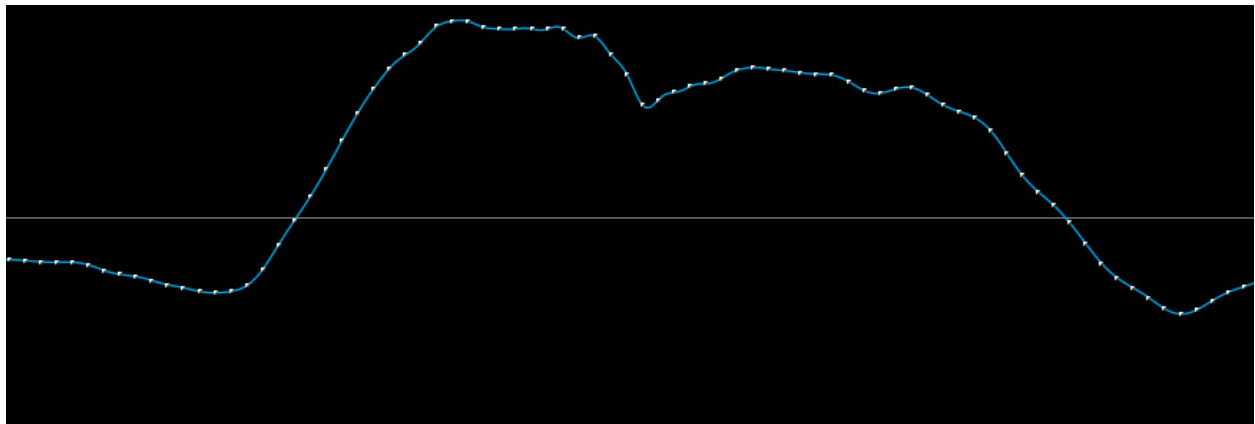
The goal of removing clipping is to repair clipped sections while allowing the original audio to sound as natural as possible.

Although a certain level of residual noise is often deemed acceptable by the human ear, clipped audio causes actual destruction of the audio—something the human ear tends to be less forgiving with.

Some distortion cannot be fixed. For example, old records that have been over played may have groove wall distortion from record wear—something that is virtually impossible to fix.

PRINCIPLES OF REMOVING CLIPPING

Declicking tools can help treat both analog and digital clipping or overload distortion.



Zoom in on a clipped waveform to clearly see where the audio has been truncated. The squaring off of the waveform occurs when there isn't enough dynamic range available to express the amplitude of the audio signal.

It's technically possible to obtain moderately good results by using the mouse to slowly redraw each sample—restoring it to what it might have been prior to clipping. Declipping tools, however, use more advanced techniques to intelligently redraw the waveform. This saves the audio engineer lots of time, and removes manual work.

In RX 3, you can use the Declip module. It uses advanced interpolation techniques to intelligently rebuild the peaks of clipped audio.

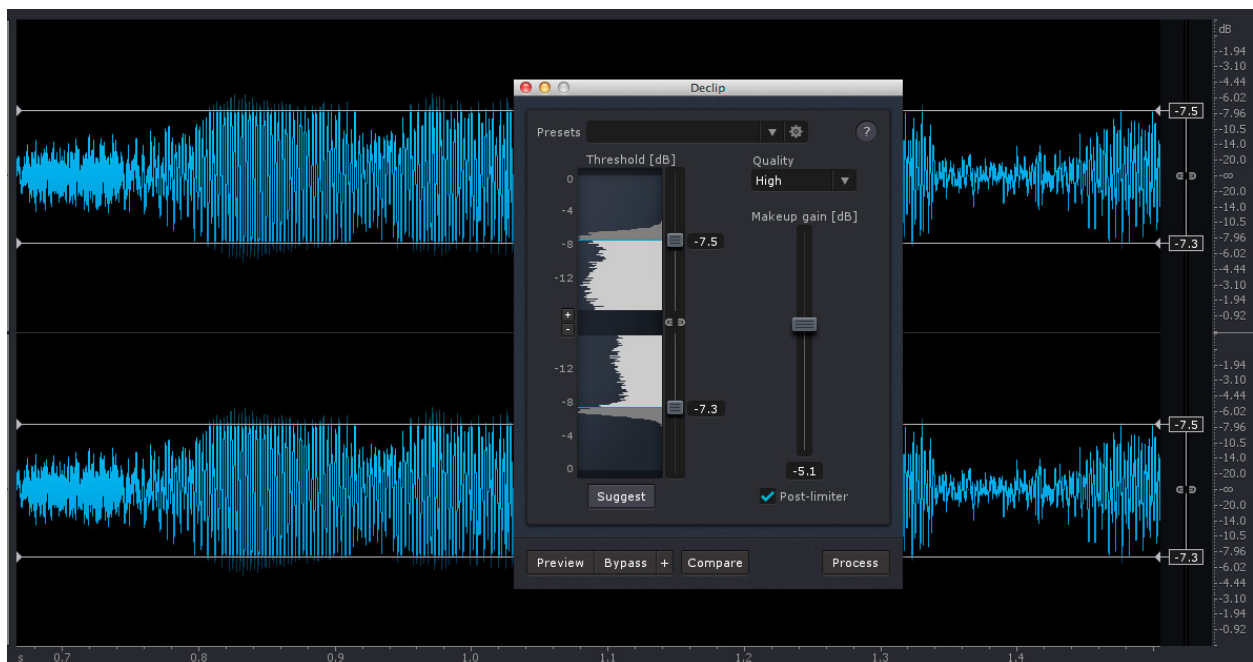
USING DECLIP IN RX 3



Step 1:

If you can actually see where the clipping occurs, Declip is both easier to use and more likely to succeed. Once you've identified exactly where the clipping occurs, you can open Declip in the RX 3 standalone application and set a suitable threshold. There are two ways you can do this:

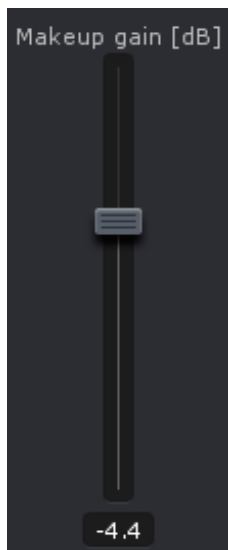
1. Zoom in on the waveform and drag the threshold down until it sits just below the truncation.
2. Adjust the threshold so that it sits just below the clipping as displayed on the histogram display—this automatically computes for you, based on the selected audio.



RX 3 TIP • The height of the histogram represents the high and low levels of the audio, whilst the width of the bar indicates the most frequent levels. Here, a white line that covers the full width indicates that clipping is present.

RX 3 TIP • Clicking on the Suggest button will cause Declip to automatically place the threshold. This may save you some time and manual work.


Step 2:



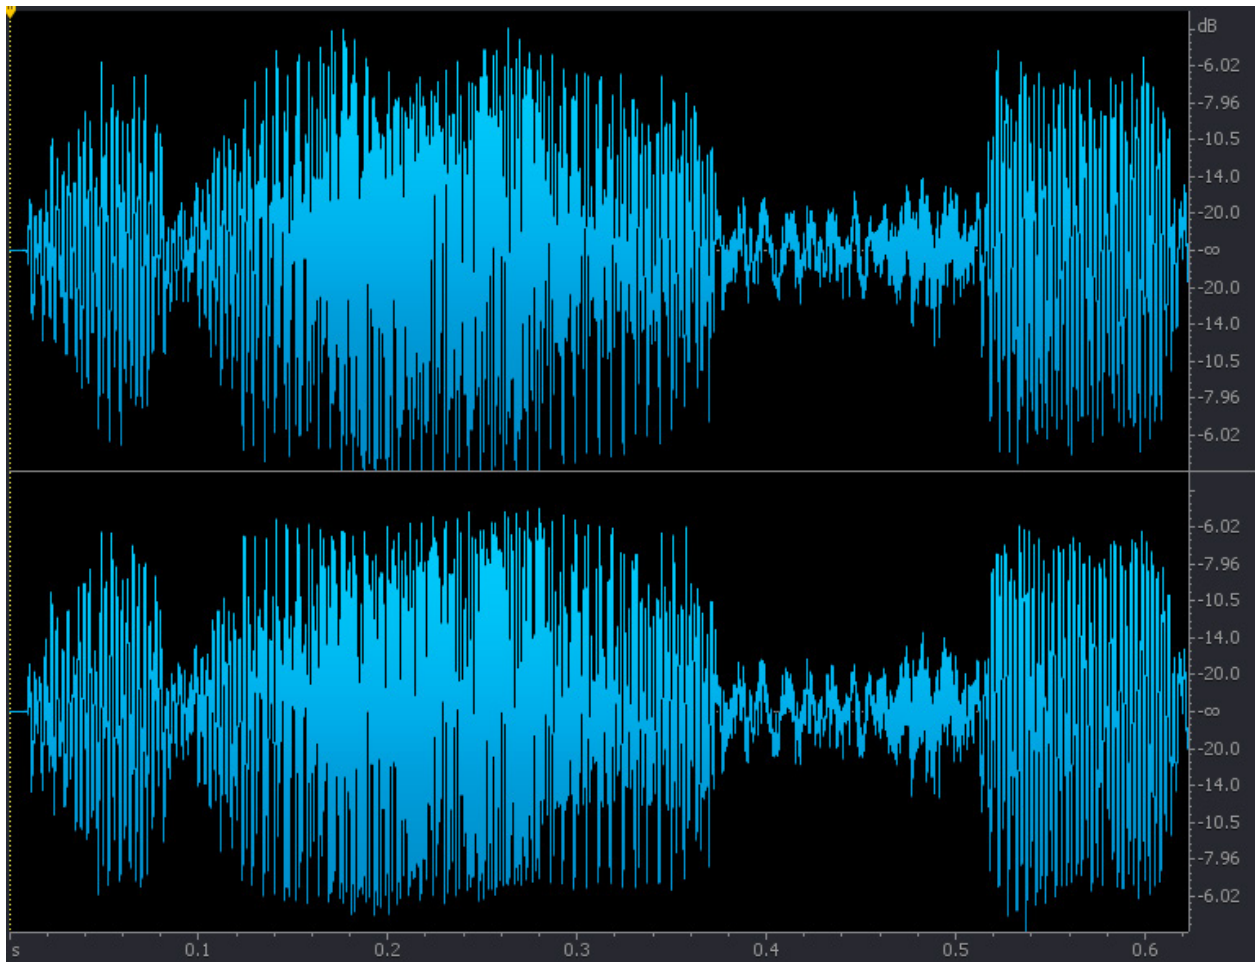
Declipping tools redraw a truncated waveform. Therefore, if the waveform is clipped at 0dB, redrawing it will cause the waveform to go above 0dB. Declipping tools utilize makeup gain to counteract this, which allows you to recover the natural sound without creating additional clipping. This facility accommodates the increased peak levels caused by the declipping process.

Generally, makeup gain of around -3dB to -6dB is a good amount to begin with. Try that, and click process.

RX 3 TIP • Using the Post-limiter in the Declip module will automatically limit the audio to avoid creating additional clipping. This may be necessary in some cases to avoid a drastic reduction in level, but you shouldn't solely rely on this. You can get a more natural sound with the right makeup gain settings applied.

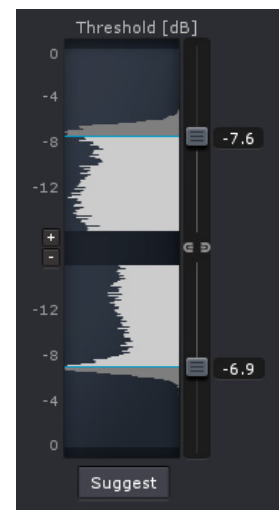
 Post-limiter

Step 3:



Occasionally the audio may have what's called asymmetric clipping—a scenario where clipping occurs at a different level on one side of the waveform.

When confronted with this scenario, click the linked icon to delink the two thresholds. This will allow you to set different threshold values for the positive and negative sides of the waveform, which offers the best sounding results when removing clipping. This feature is exclusive to RX 3 Advanced.



GENERAL TIPS FOR REMOVING CLIPPING

1. If your declipping tool does not work for you, you may be able to use RX 3's Spectral Repair tool to fix short corrupted segments. See Chapter 7 for more details.
2. To avoid volume fluctuations you may want to process the entire audio file rather than specific sections.

10: REMOVING REVERB

Reverb is used to add certain spatial characteristics to audio. It's often used on vocals and instrumentation, but can be used on sound design elements too.

It's a relevant topic to discuss, even for audio repair. Technologies that are able to attenuate reverberations in an audio signal aren't very common, and there are only a few effective tools that do this job well.

RX 3 TIP • RX 3 includes brand new technology created by iZotope that reduces reverb. It's offered as part of the RX 3 Advanced Dereverb module.

WHAT'S THE GOAL OF REMOVING REVERB?

There are two main reasons why an audio engineer might seek to reduce reverb in an audio signal.

Reason #1

The presence of unwanted reverb is a common issue when editing automated dialogue replacement (ADR), matching location recorded dialogue with studio dialogue and mixing dialogue recorded in larger spaces.

Reverb can prevent an audio engineer from achieving a smooth, warm, and present dialogue mix. Particularly if recorded in a large studio or space, reverberations might be so overwhelming that they prevent the audio engineer from adding more reverb as an effect.

In this case, the goal is to remove as much unwanted reverb as possible, without impacting the original audio.

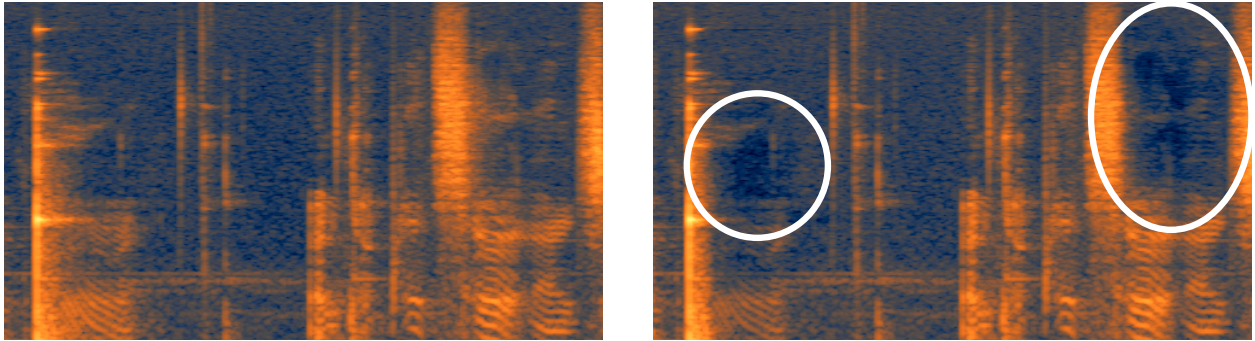
Reason #2

Reducing unwanted reverb before doing heavy denoising is a helpful process when repairing and restoring audio.

An overly enthusiastic denoising process may modulate reverb tails, which results in unwanted artifacts. If you're required to perform heavy denoising, reducing the reverb tails first may lead to a better result.

In this case, the goal is to reduce rather than remove entirely.

PRINCIPLES OF REMOVING REVERB



The above images show a reverberant signal versus the same signal without reverb. The difference between the two is most obvious in the decay that follows the transient, known as the reverb tail.

Technologies that seek to attenuate reverb often look to these tails in order to identify and distinguish the reverb. This is so that attempts can be made at reducing its effect.

Reverbs with short tails or heavy early reflections can be very hard to reduce, as it makes identifying reverb more difficult for algorithms.

It isn't always possible to remove early reflections, and you may have to settle for reverb reduction rather than total removal.

USING DEREVERB IN RX 3

Please note that reverb reduction tools are very few in number, and you may find that the controls mentioned here are specific to RX 3 Dereverb.

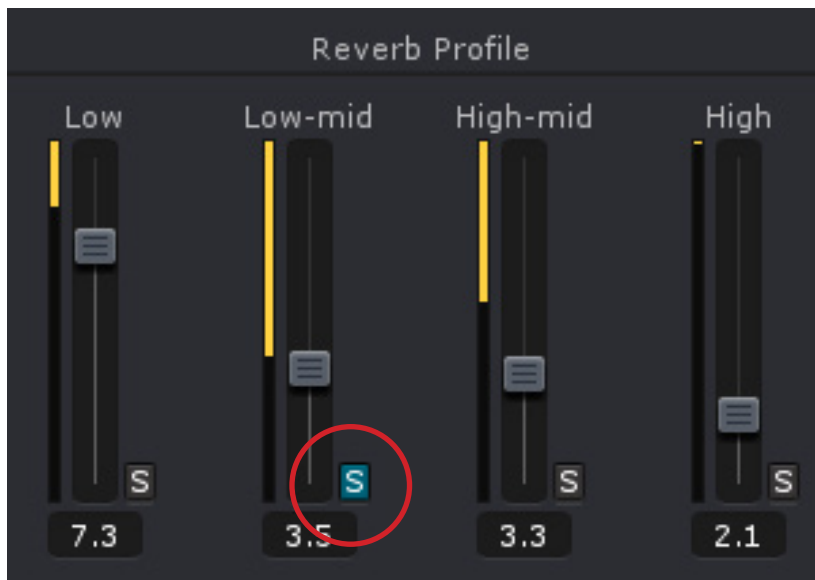
Step 1:

Listen to your audio. Try and discern how long the reverb tail is, and then set the tail length accordingly. Using an incorrect tail length setting can cause more harm than good.



RX 3 TIP • The Learn feature may help here, as this will automatically set the band thresholds for you. For best results, learn from the entire audio file. Dereverb is not like Denoise, thus learning from a selection of 'just reverb', as you would 'just noise', will not help the computer distinguish the reverberant elements of the signal.

Step 2:



RX 3's Dereverb is multiband. There are four sliders allowing you to adjust the amount of reverb reduction in different areas of the frequency spectrum.

If you used the Learn feature in step 1, these sliders may have been initially positioned for you. Step 2 is all about making sure.

Using the Solo function, listen to each band and adjust the slider until you're happy with the amount of reverb being reduced.

Step 3:

Now that you've set the frequency-specific processing details, you can adjust the master Reduction slider to control the overall amount of reverb reduction.

Now you should use the Output Reverb Only function to listen to the signal being removed. As with other repair and restoration tools, it's useful to check what you're actually removing, in case it contains any desirable audio. In this case, you should be hearing the reverb in isolation, and not the direct audio.

RX 3 TIP • Enhance Dry Signal is a useful additional control, designed to help clarify the audio being left behind, and avoid dullness.

GENERAL TIPS FOR REMOVING REVERB

1. Reverb is often centered in the mid range, and focusing your attention on the low and hi-mid multiband sliders will help ensure a good result.
2. Use a gentle pass of 2-4dB of denoising before removing reverb and then doing some heavier denoising. A slight, gentle pass can help remove the top 'silky' layer of noise in the audio signal, allowing reverb removing algorithms a better peek at the reverb underneath. You might find this gets better results even if not explicitly recommended by the reverb tool's manual.

11: EXPORTING AND DELIVERING AUDIO

Once you've completed the repair and restoration process, you'll want to take the necessary steps to export and deliver your work.

This might be as simple as rendering/bouncing your mix from a DAW, or saving an audio file that you can use elsewhere. But sometimes, particularly in the case of audio forensics or audio for TV and broadcast, delivery requirements might be somewhat demanding.

WHAT'S THE GOAL OF EXPORTING AND DELIVERING AUDIO?

Much like the mastering process, the goal is to deliver high-fidelity audio in the correct medium—with the right deliverables attached.

TIP • A deliverable could be something else that's required of the audio engineer, such as the documentation of steps taken (often required in audio forensics), or a proof of loudness compliance (often required for broadcast audio).

PRINCIPLES OF EXPORTING AND DELIVERING AUDIO

It may sound obvious, but knowing what the client requires is perhaps one of the most important principles of exporting and delivering audio.

Establishing the deliverables at the beginning of a project, whether for personal or commercial purposes, can save time and help avoid additional hours of prep work.

EXPORTING AND DELIVERING AUDIO IN RX 3

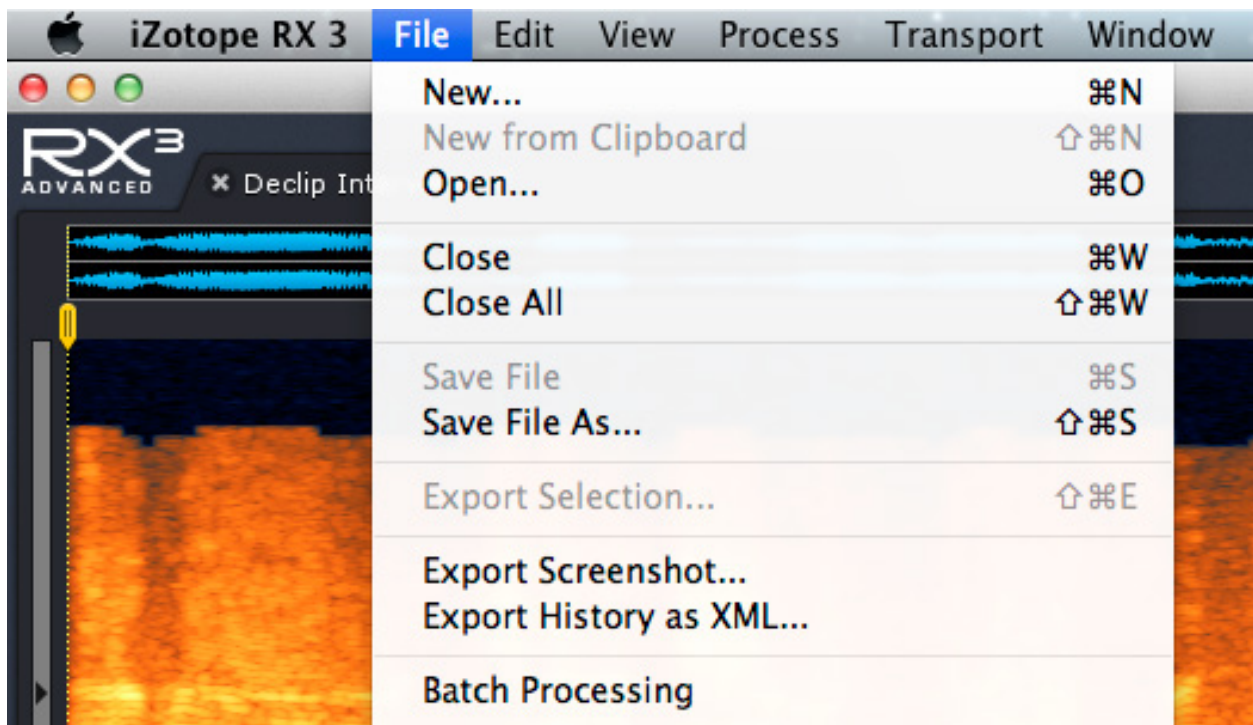
If you've been using the RX 3 plug-ins inside a DAW, you can simply render/bounce your mix session.

RX 3 TIP • If any of the RX 3 plug-ins have been using a lesser-quality setting for latency purposes (such as Low or Medium in Declip), change this to High before bouncing.

When using the RX 3 standalone application, you have several options at your disposal. Although these options are described within the context of RX 3, the principles of maintaining session data and multiple copies of your work still apply across any form of audio repair and restoration work.



Step 1:



Click on File in the upper menu bar.

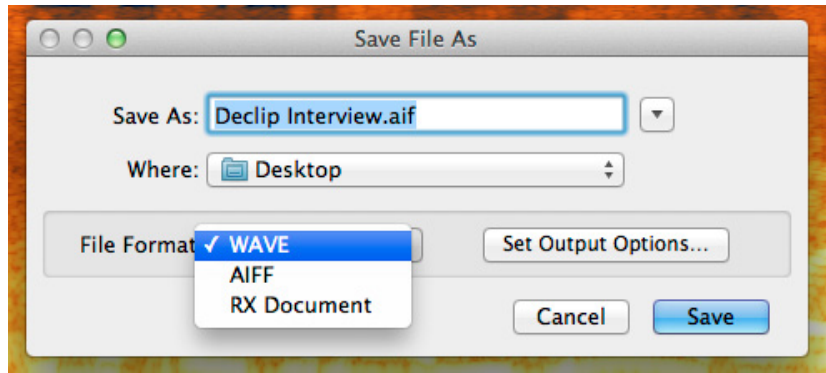
This provides several options.

- Save File As... allows you to save your work as an RX 3 document, which is a session file that saves your processing history for later recall, or as a WAV or AIFF file.
- Export Selection allows you to export the audio you've highlighted, which is particularly useful if you tend to break apart larger projects.

Step 2:

Decide on your output format. The file format options available to you are WAVE, AIFF and RX 3 document.

Should you choose WAVE or AIFF, RX 3 allows you to save in 32-bit floating point. Many other audio editing and repair/restoration suites allow this too.



If you are saving your audio in order to resume work later, you may wish to save it as an RX Document, which retains your edit history.

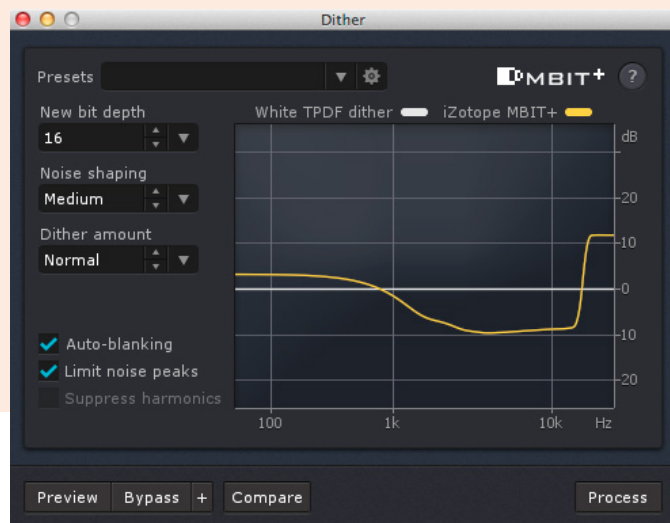
For other contexts, choose 24-bit or 16-bit. 16-bit is preferred if you're delivering media for an audio CD for example.

GENERAL TIPS FOR EXPORTING AND DELIVERING AUDIO

Prior to exporting audio as outlined above, ensure you've at least considered the tips below, as they may be relevant to a specific project or scenario:

1. When reducing bit depth, always dither. If the bit depth of the audio upon delivery is going to be lower (a 24-bit digital audio file being reduced to 16-bit for a CD master, for example), dither should always be applied as the final step before exporting.

RX 3 TIP • RX 3 includes iZotope's award winning MBIT+ Dither (also available in Ozone 5 and Ozone 5 Advanced), which manages your dithering process automatically.



TIP • If you're interested in learning more about the concept of dithering, iZotope has produced this informational video: youtu.be/vVNzylf9sGo. If you prefer to read, you can download iZotope's free dithering Guide at www.izotope.com/ozone/guides.

2. **Deliverables.** Whatever your audio repair and restoration solution, make sure you document the steps you take. Keep copies of the original audio, save presets, and note as much information as possible about the process. This is most useful for people working in audio forensics who might be required to verify and authenticate work. It also applies to freelance engineers who might need to clarify their time commitments.

RX 3 TIP • RX 3 saves your work as an RX 3 document, in addition to allowing you to export your edit history via XML. The inclusion of iZotope Insight (mentioned below) also allows you to measure and export graphs and spreadsheets of your audio loudness data—required by many large clients.

3. **Ensure loudness compliance.** If your audio repair and restoration work is part of a larger mix that's likely to be broadcast, you should ensure it complies with the latest loudness regulations in your territory. This avoids costly remixes and adjustments once you've delivered your audio.

RX 3 TIP • RX 3 Advanced includes iZotope Insight, a comprehensive metering suite that monitors and ensures loudness compliance. For more information, watch this useful video:
<http://youtu.be/XpmOLjDS0c>.

And for more information on Insight, visit www.izotope.com/insight.



4. Sample rate. Ensure everything is the correct sample rate, as required for delivery. A project destined for CD would require a sample rate of 44.1kHz, whereas the audio for a video project would typically require 48kHz. This may involve either downsampling or upsampling your audio. Please note: upsampling audio does not enhance sonic quality.



RX 3 TIP • RX 3 includes iZotope's Sample Rate Conversion (SRC) technology as part of the Resample module. It supports sample rates between 11.025kHz and 192kHz.

5. Control. Always retain a copy of your work that is of the highest quality. If your audio is going to be delivered across multiple formats (such as MP3, AAC, WAV or FLAC) it's a good idea to handle these conversions yourself. This will allow you to maintain quality control, and avoid a client taking your high fidelity work and using a bad encoder that will convert to these formats.

12: SUMMARY

We hope this guide has increased your knowledge of audio repair and restoration and, as a bonus, given you some ideas of how to use RX 3 effectively. When conducting audio repair and restoration, it's important to understand that each noise problem is different. Learning the fundamentals, such as those outlined in this guide, will allow you to quickly and effectively hone in on which tool you need for the job, and how best to use it.

Ultimately, there are no right answers, no wrong answers, and no rules—audio repair and restoration is about experimenting and having fun. To test what you've learned from this guide, be sure to download our audio examples and move on to Appendix C, where you'll find descriptions of how to repair real world audio problems. Enjoy!

Thanks,

iZotope, Inc.

P.S. Aside from RX 3, we also invite you to try out our other products. Each one is available for download as a free 10- day trial at www.izotope.com/trials.

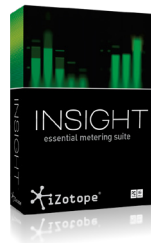
Here's a quick overview of several of our other products:



iZotope Ozone 5 Complete Mastering System

Ozone is a single, integrated plug-in that includes eight essential mastering tools: Maximizer, Equalizer, multiband Dynamics, multiband Stereo Imaging, Post Equalizer, multiband Harmonic Exciter, Reverb, and Dithering. Become the master of your audio!

www.izotope.com/ozone



iZotope Insight Essential Metering Suite (Included with RX 3 Advanced)

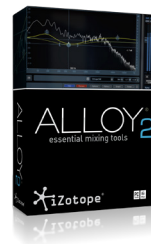
Insight is an extensive set of audio analysis and metering tools, perfect for visualizing changes made during mixing and mastering, troubleshooting mixes, and ensuring compliance with broadcast loudness standards.

www.izotope.com/insight



iZotope Nectar 2 Complete Vocal Suite

Make your vocals and dialogue sound professional in a broad range of genres with Nectar's complete set of vocal production effects. www.izotope.com/nectar



iZotope Alloy 2 Essential Mixing Tools

Alloy gives you futuristic tools, fast results and—most importantly—fantastic sound. Bring character and life to every element of your mix with Alloy.

www.izotope.com/alloy

ABOUT THE AUTHORS

In addition to the hard work of iZotope's development, QA and content development teams, we also wish to extend our thanks to Nat Johnson whose writing, restoration expertise and audio examples contributed greatly to the original version of this guide published in 2008.

APPENDIX A: GETTING SET UP TO REPAIR AND RESTORE AUDIO

EQUIPMENT

Software

iZotope RX 3 installs both the standalone audio editing application, and the separate plug-ins for use with a DAW.

If you prefer working inside a DAW, you can load many of the RX 3 modules as plug-ins, provided it supports one or more of the following formats in 32- or 64-bit: VST, VST3, AU, and AAX.

Spectral Repair only functions in DAWs that support offline processing, such as Avid's Pro Tools 11. In a DAW, you can open a spectrogram by using Spectral Repair as a plug-in.

iZotope audio repair and restoration technology is also licensed to a number of third party manufacturers. Certain algorithms and parameters are accessible within software such as Sony's Sound Forge—a good way to get started with the basics, if you don't yet own RX 3.

Many software solutions for audio repair and restoration, such as iZotope RX 3, have free trials. Taking advantage of these trials is an effective way of establishing your preferred solution.

Computer Requirements

Audio repair and restoration tasks can be some of the most CPU-intensive audio processes to perform. Therefore, powerful machines are preferred, as they will be faster at processing edits. Here are some requirements to keep in mind:

- A fast, multi-core processor. You'll need this kind of processor if you want to take advantage of RX 3's optimized speeds.

RX 3 TIP • RX 3 is already fast, but is additionally optimized to take advantage of multi-core processors. Thus the faster your processor, the faster your performance will be when crunching advanced audio algorithms.

- A large, high-resolution monitor. Spectrograms are incredibly visual by nature, so the more screen real estate you can dedicate to this way of visualizing audio, the better.

TIP • Dual monitor setups allow you to place a spectrogram on one screen, and your audio processing modules on another.

- **RAM.** If you work on large files, and have several audio or video editing programs running at once, a large amount of RAM is needed to cope with the load.
- **A backup hard drive.** With hard drive technology so affordable these days, the initial cost of a backup hard drive is far outweighed by the benefits of having everything backed up and retrievable—particularly useful if your audio repair or restoration project is for a big client!

Sound Card/Audio Interface

You can use the sound card that came with your computer to monitor audio when doing restoration work, but we recommend investing in a high-quality sound card or audio interface designed for professional-level audio recording and playback.

In most cases, this will let you monitor audio with less noise, and will give you higher quality A/D inputs for transferring analog media to your computer. It will also include professional standard outputs such as XLR or balanced quarter-inch phone jacks, which will let you connect to audio reference monitors with well-shielded cables.

Visit your favorite audio retailer to check out a wide range of interfaces at a wide range of prices. Pick one that suits your needs. If, for example, you'll be transferring a lot of recordings from vinyl 33rpm or 45rpm records, choose an interface that includes a ground screw and built in phono preamp.

MONITORING AUDIO

Reliable audio monitoring, as you might imagine, is essential to successful restoration. We suggest a full-range speaker system that suits your needs, as well as your working space. It's a good idea to frequently check your work on other systems as well. As with mastering, the most important advice we can give you when it comes to monitors is to learn the unique characteristics of your pair and how they sound in your room. The best way to "learn" your monitors is to listen to lots of recordings that you know well. A good pair of headphones can sometimes help, as you may want to listen to low-level details that may not be obvious on loudspeakers.

Room Acoustics

You can have the best reference monitors money can buy, but if you're working in a tiny square room without acoustic treatment, you won't accurately hear the flat representation of frequencies. Even small details, such as where you place the speakers, will change their frequency response drastically. We recommend researching room acoustics online. Even if you're on a budget, there will always be a solution

that will help improve the sound of your room.

Companies that make acoustic treatment typically have a large number of resources available, and these resources are designed to help you determine what your space requires.

TRANSFERRING MEDIA TO THE COMPUTER

When you sit down and start restoring an audio project, you might be lucky enough to have a digital audio file that has either already been digitized, or that began life as a digital format. However, many audio restoration tasks require you to get an old recording from tape, vinyl or other source onto your computer. Here are some tips for getting the most usable digital transfers of your audio.

Start with the Master Copy

This can't be stressed enough: always utilize original, unprocessed source audio where possible. Files that have already been copied from record to tape, or digital files that have been converted to MP3 or another compressed format like lossy ATRAC, will be more difficult to repair. You'll face more audible artifacts that are introduced during the copying and compression process. Similarly, it may be much harder to get good results from audio that has already had audio repair work done to it.

Input Levels

Whatever your source—tape, disc, analog or digital—it's important to keep input signals at a moderate level if you're recording audio into a computer. A high level may result in inadvertent clipping at the soundcard or audio interface level. Too low a level and you may not have adequate headroom to accommodate changes in level as you apply certain processing functions during restoration. Since the noise floor of an analog to digital converter is fixed, you'll have a lower signal-to-noise ratio if you record at an unnecessarily low level.

TIP: Some CD and DAT players feature digital outputs, and many sound cards/audio interfaces have their own digital inputs. Taking advantage of this can be the perfect way to get digital audio directly into your computer without having to go through conversion to analog and back to digital.

RECORDING FROM ANALOG FORMATS

Vinyl LPs and 45s

After transferring a record to a computer, you can do a lot to remove noise and improve the quality of the recording. However, getting a good transfer from a vinyl record can be tricky. Keep these things in mind:

- The vast majority of turntables don't output a line level signal. They usually require a phono pre-amp. Phono preamps not only boost the level of audio coming in from the turntable, but they also

apply special equalization—a part of the RIAA standard for creating vinyl records. Without this gain and EQ stage, your transferred vinyl will likely sound bad.

- Most audio amplifiers designed to work with turntables have phono preamps built in. So too do professional DJ mixers (many of these also provide excellent balanced outputs suitable for connecting to your audio interface). As mentioned earlier, some audio interfaces even have built-in phono preamps.
- Some turntables now include direct USB outputs, getting around the preamp problem. However, keep in mind that many of these are very inexpensively made, and that the audio quality usually suffers during the A/D conversion process. Make sure you check the specs before you buy. A good turntable/audio interface combination may give you better quality and value for money.
- Make sure your turntable is grounded! Mixers and amplifiers with phono inputs will usually have a thumbscrew, which allows you to connect the ground wire from your turntable. Use it! If you don't, you'll probably be adding electrical hum to your transferred file.
- Clean your records! Removing dust with a vinyl brush and/or using a vinyl-cleaning product will make a huge difference in the quality of your vinyl transfer. Why spend hours removing dust pops on the computer? You can get rid of most of them beforehand by simply cleaning your record.
- Use a good stylus. If your record player has had the same needle for 30 years, chances are it's time for a replacement. Using worn out needles can affect the audio quality and cause damage to your records. Some turntables are now hard to find replacement needles for. Nevertheless, there are plenty of companies from whom you can purchase a replacement stylus.
- When you have a choice, choose a stylus that is made for archiving. Many available solutions are tailored to scratch DJs and sacrifice sound quality for tracking and other characteristics.

Shellacs, Cylinders and Older Phonograph Formats

There are a wide range of formats out there that preceded the modern vinyl LPs and 45 singles that are so commonplace today. In general, use caution when trying to transfer these types of recordings yourself, without the proper equipment. Many shellacs, for example, require different, non-standard stylus tips due to the varying widths of their groove wall. Some formats, like the Edison Diamond discs, have grooves that capture audio through vertical motion (called the hill and dale method) rather than lateral movement used by modern vinyl records. If you don't have the right equipment, consider finding someone that specializes in formats like this to do the digital transfer for you.

While we could write an entire book on these formats, it's simply out of the scope of this guide. If this is an area you're interested in, you'll find numerous web sites offering useful information and good tips on 78 RPM history, availability and playback equipment, as well as storage and general care advice.

Magnetic Tape

There are endless recordings out there made on cassette recorders and reel-to-reel machines. Whether you're trying to restore the Watergate tapes, working on reel-to-reel masters, or simply cleaning up old four-track cassette recordings of your high school garage band, there are a few things to keep in mind when transferring:

- The quality of your playback equipment will make a difference. If you connect your old ghetto blaster's RCA outputs to your computer, you'll not only introduce more noise into the transfer, you might have other unwanted problems like speed fluctuations.
- Slight changes in a tape's speed (called wow and flutter) can cause problems when you're trying to remove noise. This is because the noise can modulate along with the tape playback speed. It can make tonal and broadband noise difficult to remove as it changes over time. Work with a well-built and well-calibrated tape machine if possible.

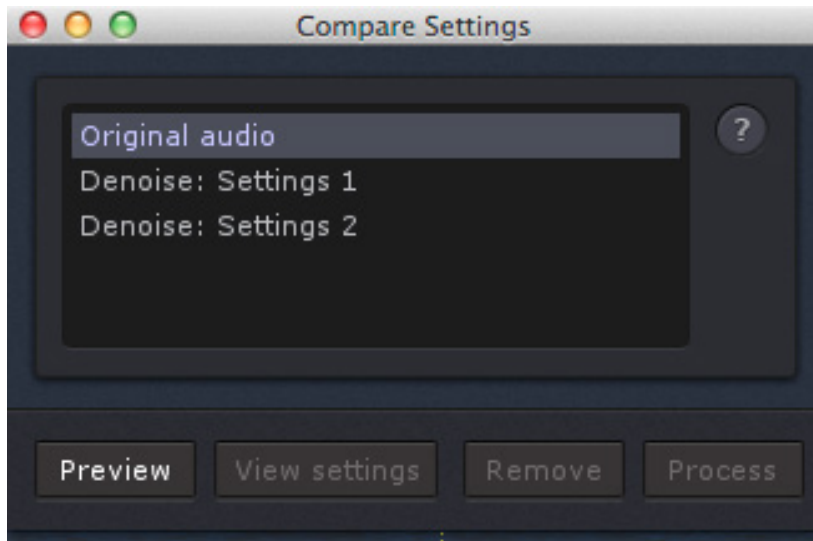
RX 3 TIP • RX 3 Advanced has a Deconstruct module. You can use this to reduce noise that changes sound as a result of wow and flutter, and also the time-variable wear found on older 33s, 45s and 78s.

- Tape heads can pick up debris and become magnetized over time, leading to poor playback quality. Cleaning your deck's tape head, capstan and rollers can help ensure the upmost quality for your recordings. You can find deck cleaning and demagnetizing kits online, and in specialty stores.
- Old mylar/polyester tape can stick to itself over time, which causes tapes to self-destruct upon playback. If you're pulling an old reel-to-reel tape out of a dusty archive, you may want to consult with an expert before trying to play it. Some tapes can be baked in an electric oven to remove moisture and revitalize the tape's adhesive, but we don't recommend you try this at home.

APPENDIX B: GENERAL RX 3 TOOLS

The following is a list of general RX 3 tools that are extremely useful for the workflow of the audio repair and restoration engineer.

COMPARING SETTINGS



The ability to easily, repeatedly and accurately compare settings—visually and aurally—with the RX 3 modules is a valuable tool and timesaver for every project you undertake. You'll find the Compare Settings feature to be an extremely useful reference when contrasting a wide range of choices with the original.

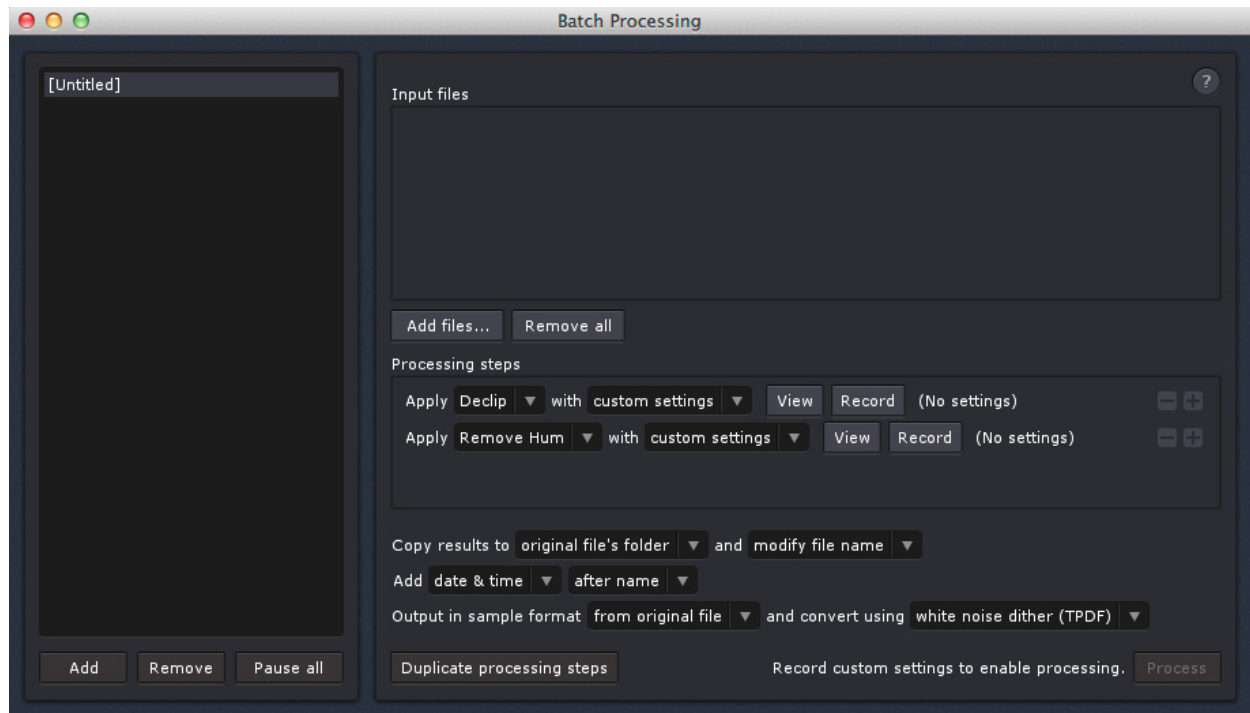
Using the Compare Settings Window

Find a setting in the module you feel works for a specific situation and then click on the Compare button. RX 3 will automatically open the Compare Settings window while processing the result of your settings in the background.

View and change your settings and click again on Compare. RX 3 will add a new item to the list for each time you click on it. You'll have an instant comparison with every setting listed in the window.

Notice that the Spectrogram display, as well as the sound, changes each time you select a setting in the Compare Settings window.

BATCH PROCESSING



RX 3 offers numerous time savers, and Batch Processing is one of its best. Simply defined, Batch Processing enables the automation of processing on file groups.

If you have several files that need to be processed in the same way, you can use Batch Processing to define the set of steps (e.g. denoise, then declick, then normalize), as well as the output format and naming convention for each. RX 3 will then automatically process all of the files in the background.

Apply Batch Processing to Files

Click File from the upper menu bar, and then click Batch Processing. Click the Add Files button to add one or more files for batch processing, then choose the processing module(s) you wish to use.

Select a preset for the module, or define your own settings. To view the settings for a particular batch-processing job, click the View Settings button.

RX 3 TIP • It's a good idea to use one file to manually establish the settings which can be applied to the other files that share the same problem.

Running a Batch Process

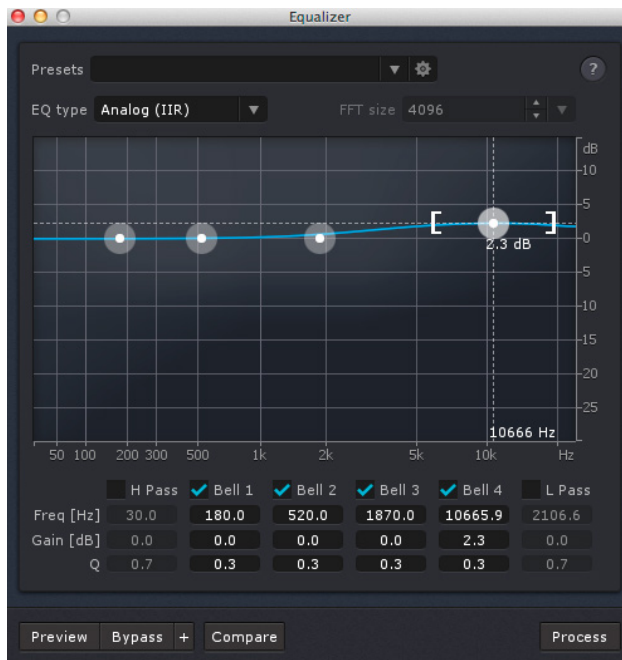
Once you're satisfied with the batch processing jobs, click Process to run them all. You'll see a progress dialogue box while RX 3 runs each job. To cancel the current job and all subsequent jobs, click Cancel.



WAVEFORM STATISTICS

Waveform Statistics supplies the audio engineer with useful information about the peak, RMS and loudness levels of the audio, as well as potential warnings about DC offsets and clipping.

It's accessible via View in the upper menu bar.



EQ

RX 3 includes a 4-band parametric equalizer module with adjustable notch filters and low-pass/high-pass filters. This is a handy tool for cutting or boosting certain frequencies, or for quickly removing unwanted low- or high-frequency sounds like rumble and hiss.

The EQ module comes in two EQ Types:

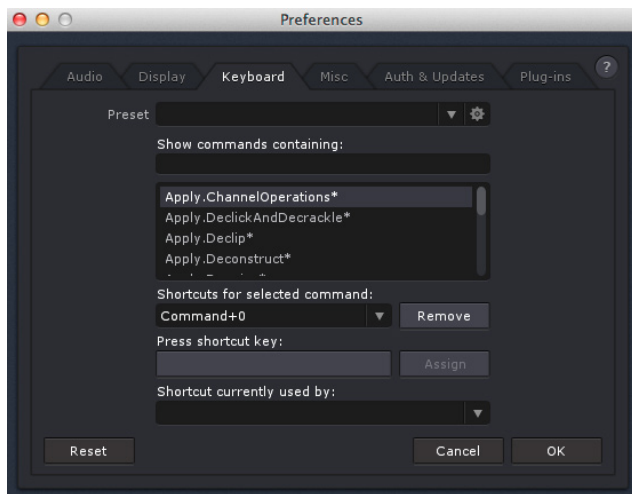
Analog EQ is a non-linear phase filter. The Linear-phase EQ uses a FIR (finite response) filter.

The analog filter applies a very different (some say warmer) character compared to the linear phase filter, which is very precise and

designed to minimize phase shifts in audio. You can easily switch between the two filters to hear which best suits your project.

PRESETS

Groups of settings in each of RX 3's modules and plug-ins can be saved and recalled as presets. Once you've made modifications to the settings in a module, you can name the preset and save it with a custom filename. Moreover, you can easily export your presets as .XMLs and store them on your computer—perfect for backup, or to share with other RX 3 users.



KEYBOARD SHORTCUTS

Similar to presets, RX 3 includes default keyboard shortcuts. For greater flexibility, you can customize these to better suit your personal preferences. The keyboard shortcuts menu is simple, and it will save you time and increase your efficiency.

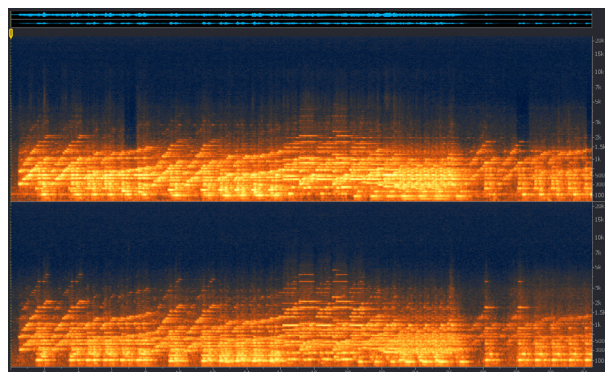
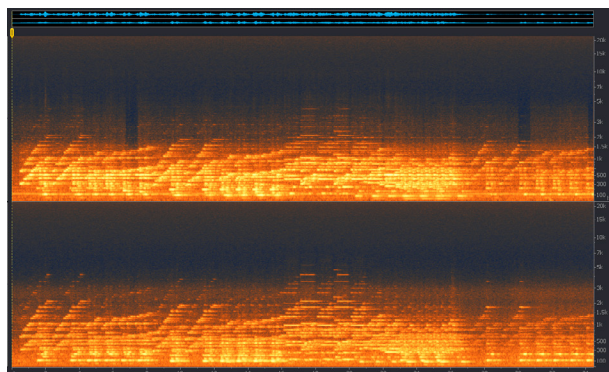
APPENDIX C: REPAIRING EXAMPLE AUDIO

Let's take a look at some real world examples of audio problems and some tips and tricks to help go about fixing these common problems. To try out these fixes for yourself, [download the example files](#).

EXAMPLE 1: REMOVING BROADBAND NOISE FROM A CONCERT RECORDING

Details:

Player-organ recording in Boston: Vienne Organ Symphony No. 1- Allegro (STEREO). Originally recorded on Ampex 440-B, 1/4" analog tape, zero noise reduction used during session.



Comments:

The principal problem was a steady, low-level background noise resembling tape hiss. The challenge was to reduce the noise without sacrificing upper register harmonics that appear during quieter passages. The settings used in Denoiser, particularly threshold, were conservative.

The Spectral Denoiser has some tools that may help:

- The 'Broadband' and 'Tonal' reduction sliders can be delinked. There are not really any tonal elements in the noisy signal, and reducing them would perhaps affect the low-level organ harmonics as they fall into



the noise floor. The 'Reduction Curve' has been enabled, and it reducing the high and low ends more harshly, since this is where the tape hiss is occurring. The curve is also being used to back off the mid range, where the organ is centered.

- In RX 3 Advanced, you can see the 'Enhancement' slider. 'Enhancement' is being used to enhance signal harmonics that fall below the noise floor.

TIP • Examining the waveform of a quiet passage, and checking your program aurally against the original, helps verify that your material is not being compromised during processing.

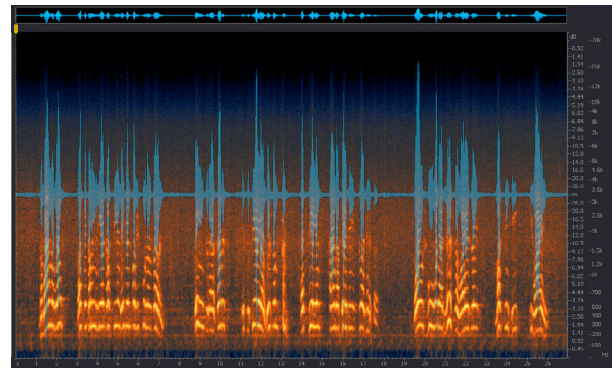
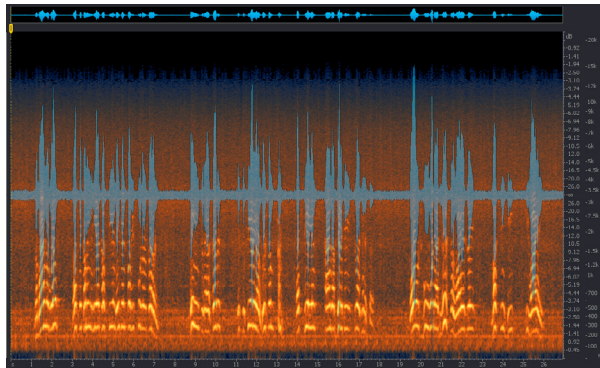
Goals:

- Train RX 3's Spectral Denoiser using a sample of pure noise. Use Denoiser to reduce background hiss and low-level hum.
- Try not to damage the musical quality of the performance—sometimes leaving a little noise behind is better than taking the life out of the recording! You may want to be conservative with the noise reduction controls, and use the smoothing options to prevent artifacts.

EXAMPLE 2: RESTORING AN HISTORICAL SPEECH: MAKING VOICE MORE IN- TELLIGIBLE

Details:

This file contains the voice of famed philanthropist Andrew Carnegie, recorded in 1914 (Internet source) (MONO).

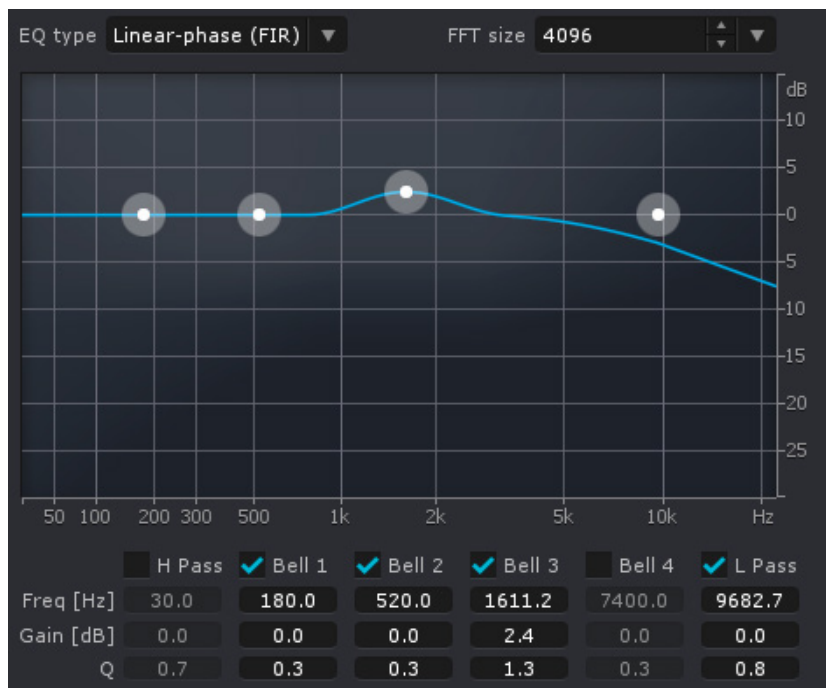


Comments:

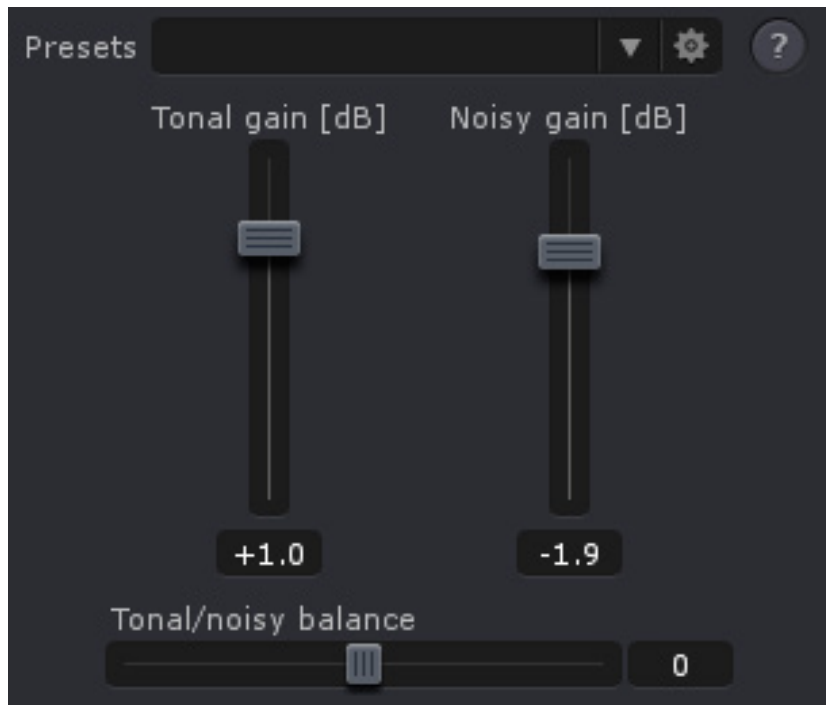
Some of the words in this speech, recorded by Carnegie in 1914, are somewhat difficult to decipher—owing to a pronounced Scottish accent, made even more muffled by someone’s rudimentary attempt at noise reduction. In addition, there’s a steady background noise present. raise the voice out of the mud using a slight boost—roughly between 1K and 9K—and then using Denoiser. The EQ compensation was left intact following the use of Denoiser, as this left the speech a bit clearer.

Goals:

- Use gentle linear-phase equalization to lift the voice out of the background noise.



- Use Deconstruct to boost the tonal elements of the signal (voice) and reduce the variable noise/crackle.



- Use the Spectral Denoiser to reduce steady-state background noise.

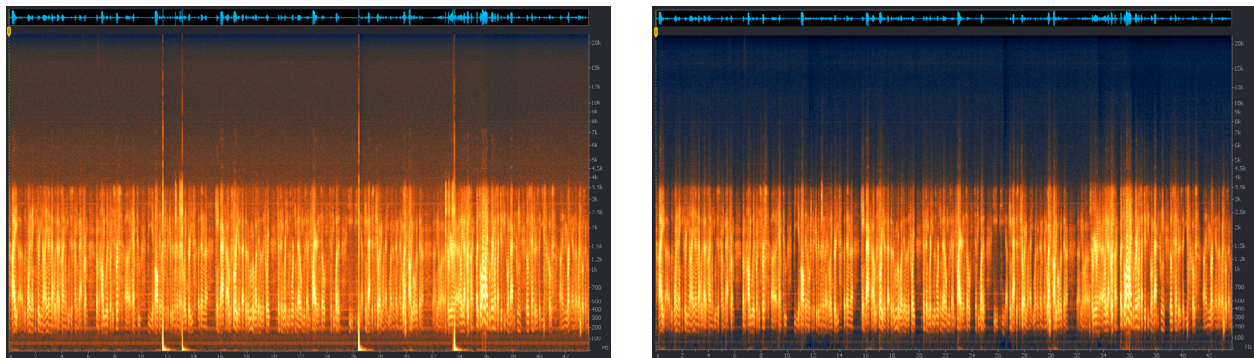


TIP • Using an approach that incorporates both EQ and Denoiser can make buried dialogue much more intelligible.

EXAMPLE 3: CLEANING UP A PHONE INTERVIEW WITH DECLICK AND SPECTRAL REPAIR

Details:

Interview with Francis Ford Coppola, recorded directly to DAT from phone patch (MONO).

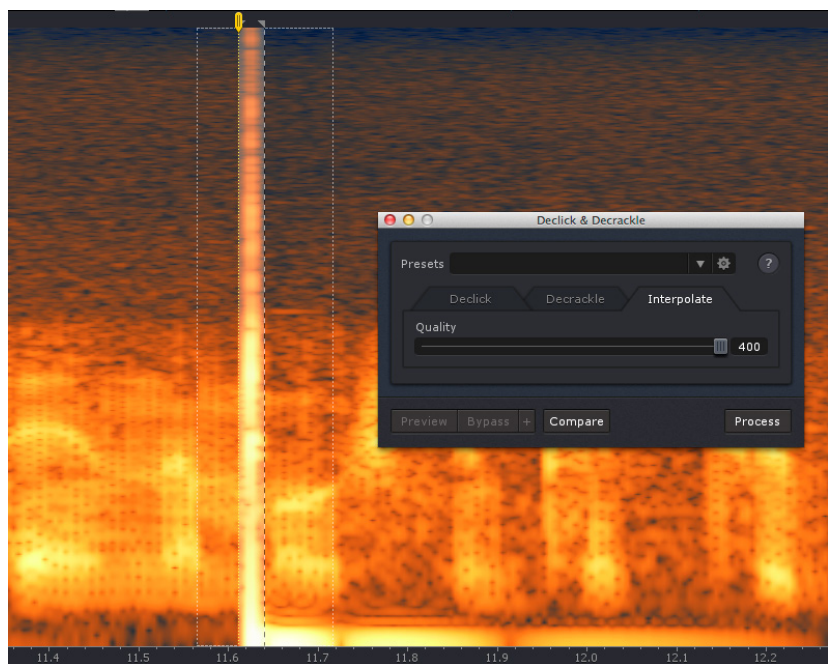


Comments:

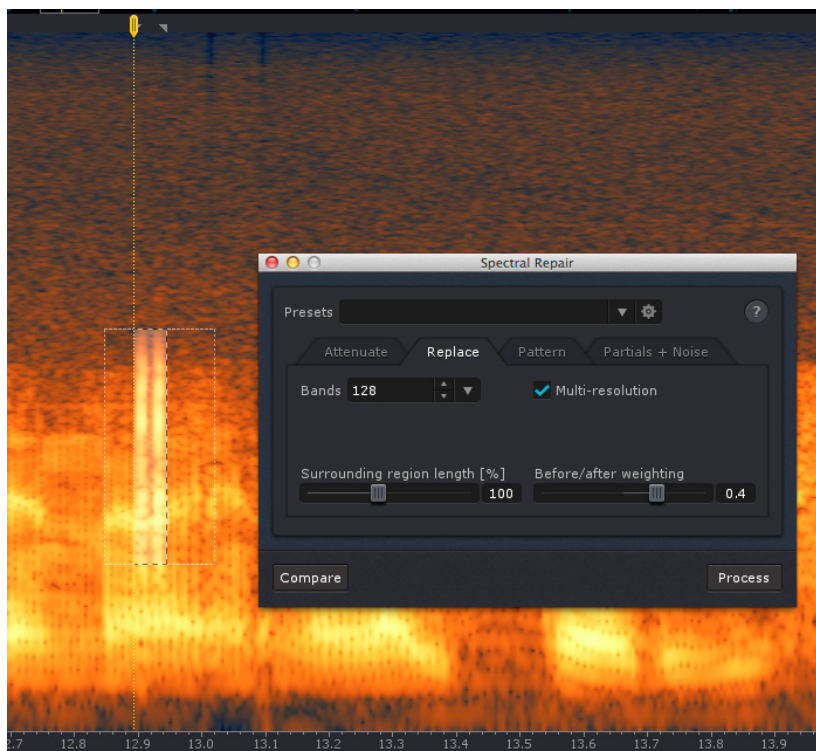
The problems were caused mainly by mechanical contact with the telephone handset. For instance, at 11.62 seconds, where Coppola says the words ‘motion picture business’, a loud click is audible.

Goals:

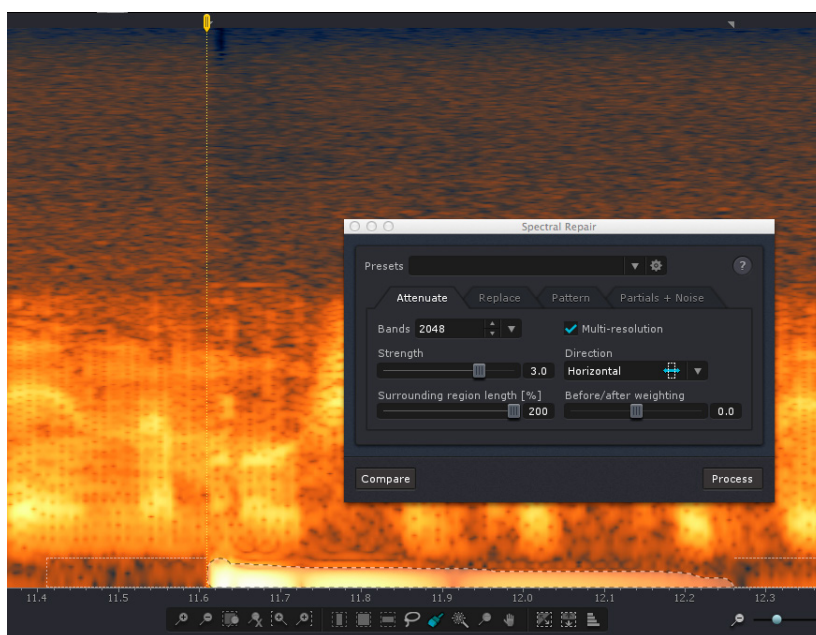
- Use Declick’s ‘Interpolate’ mode to remove the full bandwidth



- Use Spectral Repair's 'Replace' mode to remove the limited bandwidth handset noises.



- Use Spectral Repair's 'Attenuate' to remove the low frequency thumps.



- Use the Dialogue Denoiser to reduce the level of background noise in the dialogue.

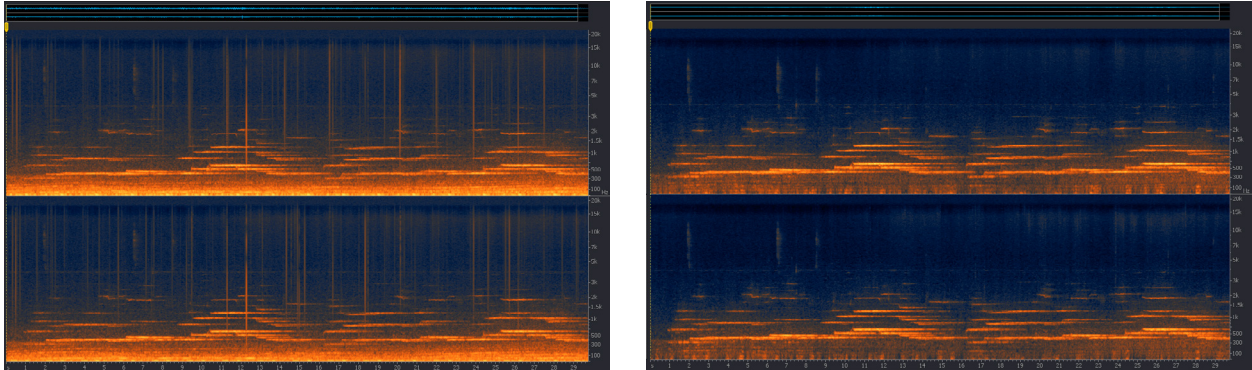


RX 3 TIP • Each of the phone handset noises is comprised of a vertical spike and a horizontal thump. Removing the vertical spikes with Declick's 'Interpolate' mode still results in an audible thump. A quick pass with Attenuate mode will remove the low frequency sound.

EXAMPLE 4: REMOVING CLICKS AND POPS FROM A CONCERT ON RECORD

Details:

Kings College Choir, Cambridge, England. Source: 33 1/3 RPM LP (1965, STEREO).

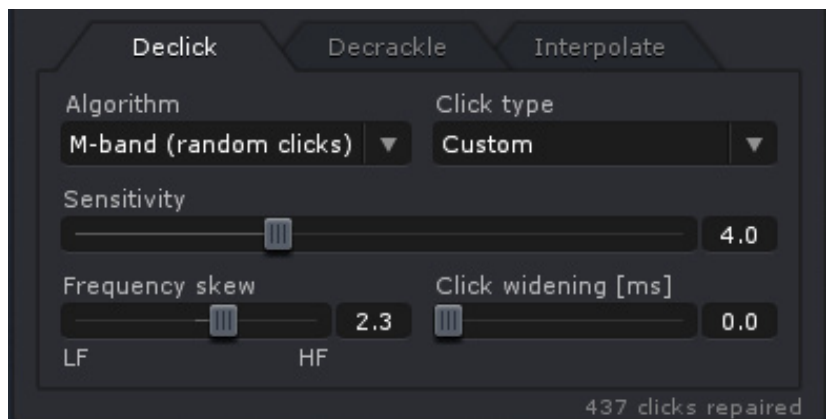


Comments:

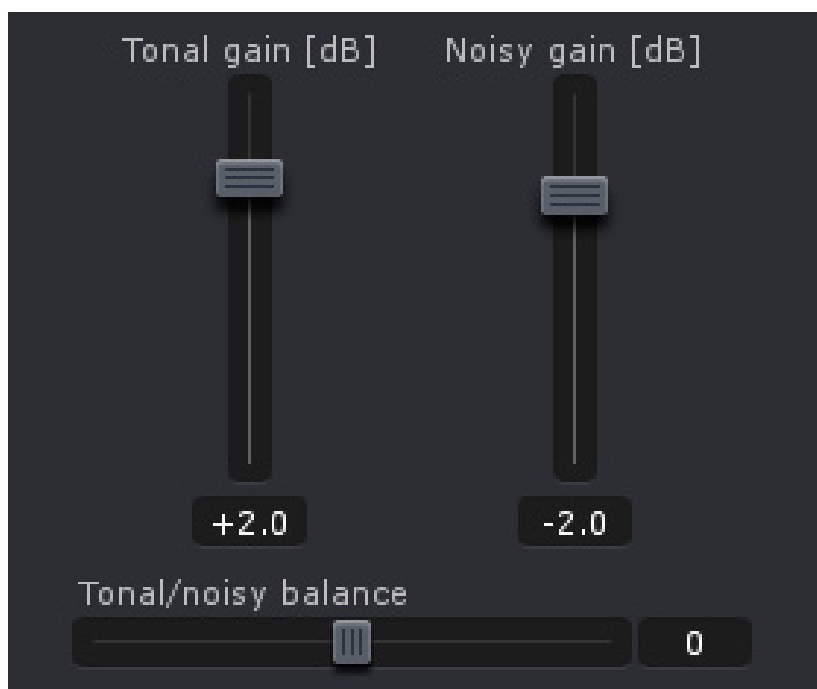
This vinyl disc was transferred for playback as a CD. Fortunately, the only problems encountered were mostly light vinyl clicks and heavy room ambience—a result of the massive chapel at Cambridge University being the recording venue. Declick was employed using the ‘M-band (random clicks)’ algorithm and one or two spots in manual mode.. The rumbling room tone was reduced by the Spectral Denoiser.

Goals:

- Use the Declick module to remove the clicks and pops.



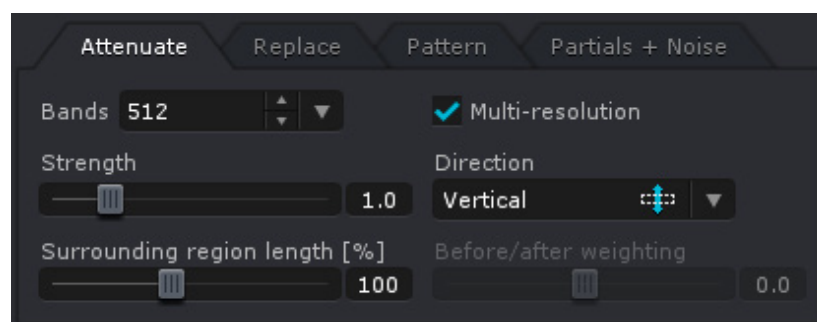
- Use the Deconstruct module to reduce the variable noise components while adding a slight lift to the tonal noise components.



- Use the Spectral Denoiser to reduce background noise. Try to retain the musical character of the performance while removing as much of the offending noise as possible.

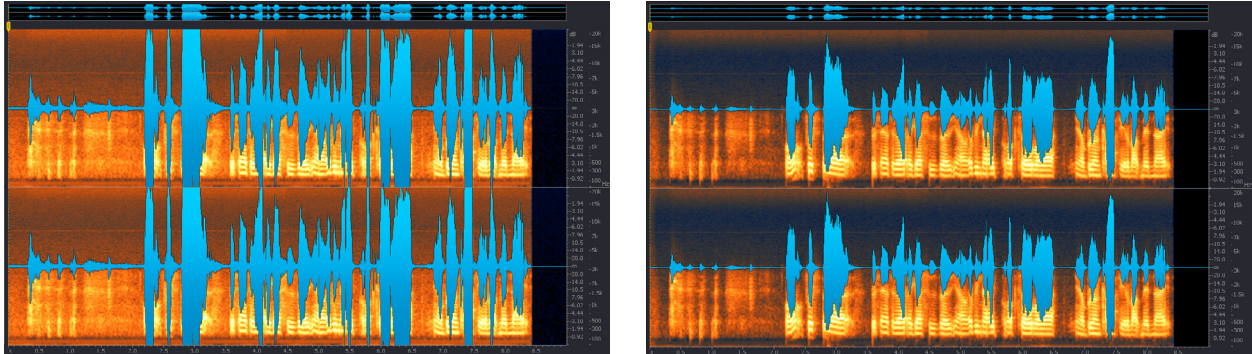


- Use Spectral Repair to reduce the low traffic hum that creeps in occasionally.



EXAMPLE 5: REMOVING CLIPPING FROM A PHONE INTERVIEW

DETAILS:



This phone interview, from an iZotope Podcast with producer Morgan Page, was recorded with a phone line breakout box and a USB audio interface.

Comments:

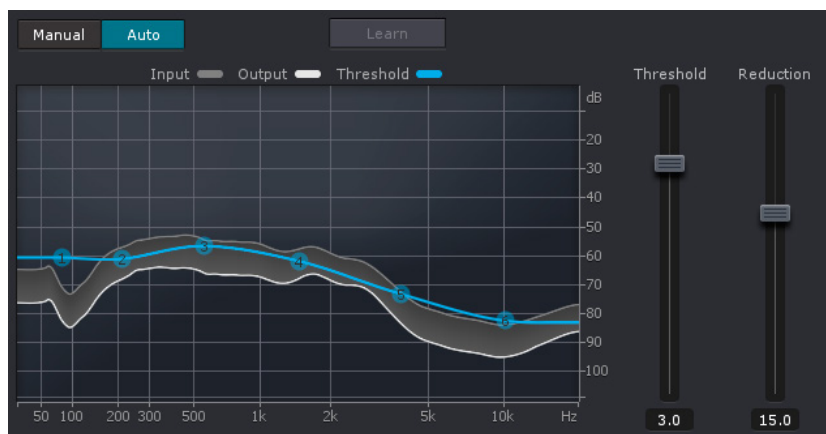
This is a classic example of input gain being set too high, resulting in some nasty clipping. In this example, the distorted sections can be clearly heard, but viewing them with the spectrogram display and waveform overlay highlights them in clearer detail.

Using the Declip module can remove the clipping entirely. The suggest feature will help suggest the optimum position for the thresholds.



Goals:

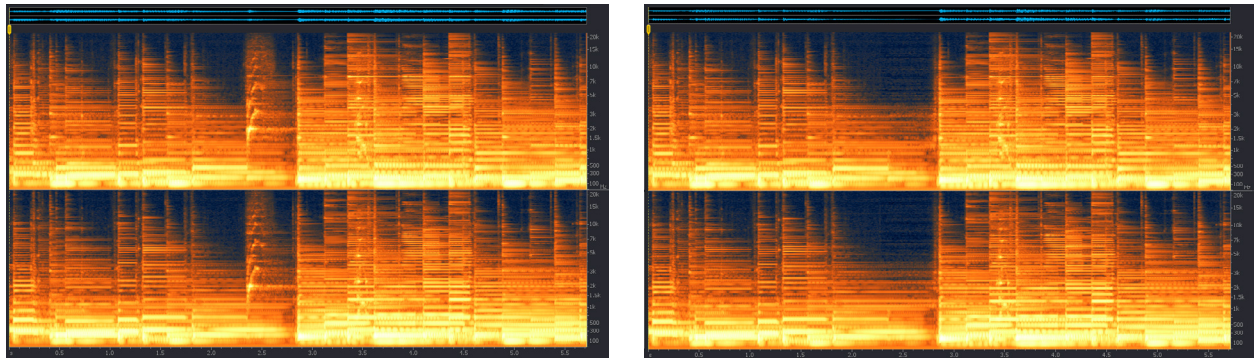
- Use the Declip module to remove the clipping.
- Use the Dialogue Denoiser to reduce the noise.



EXAMPLE 6: REMOVING GUITAR STRING SQUEAKS WITH SPECTRAL REPAIR

Details:

This recording by guitarist Jamie Robertson contains a few loud squeaks—the result of sliding one’s hand up and down the fretboard.



Comments:

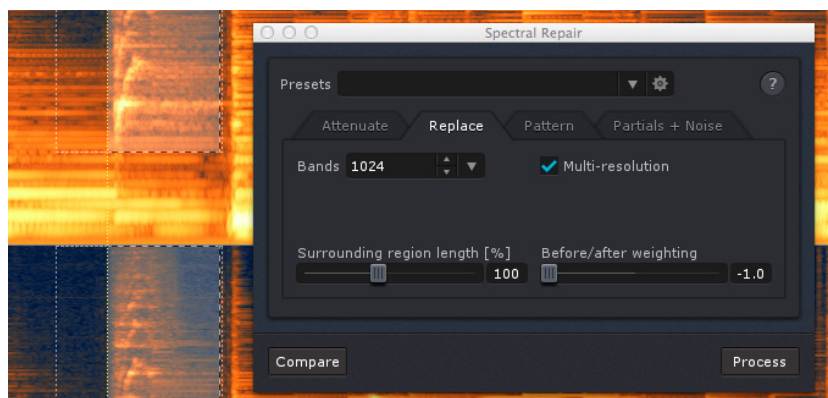
These situations are always tricky. Should you remove all minor blemishes if your tools allow you to do so? Or should you leave these squeaks in the final product to keep the recording’s character intact? Sometimes removing only the most offending events is the best tactic, but we’ll let you decide how far to go!

Identify the squeaks using the Spectrogram Display. Draw a box around them using the Time/Frequency selection tool. Use Replace in Spectral Repair to select the squeaks and apply the process.

RX 3 TIP • As these squeaks often occur right before a new chord or note is played, you may want to set the before/after weighting control all the way to the left. This will ensure that the repair is made solely using the sustained material from the earlier chord.

Goals:

- Use RX 3 Spectral Repair’s Replace to substitute squeaks for natural guitar sustain. Aim for a realistic result that does not sound unnatural.



APPENDIX D: TIPS FROM THE PROS

We asked a few of our trusted friends and RXperts to contribute their advice and tidbits with regards to audio repair, restoration and editing.

Learn more about them at www.izotope.com/artists

BILL JACKSON | Emmy winner for Outstanding Sound Mixing for HBO's *Entourage*

One of the tricks I use with RX is to carefully Denoise using the D (best) algorithm, then resample and Denoise using the C algorithm. This helps when there is broadband noise as well as a high frequency hiss that you want to reduce.

While mixing a film, I use RX Spectral Repair so frequently that I leave it open almost all of the time. It is a great tool for checking a clip to see what unwanted tones are present as well as low frequency bumps and rumbles, and then reduce or remove them by essentially painting them out.

Recently, I was asked to remove some unwanted electric dolly motor noise from an upcoming PBS show. They had previously tried to remove the noise with another software program, but the producer wasn't happy with the results. Using the standalone version of RX3, I was able to completely remove the motor noise with Spectral Repair, and since I had all the RX tools available at once, I easily removed some pops and clicks, distortion, and even used the RX EQ and RX gain on specific sections to balance the dialog tone and level throughout the scene. The producer was very happy with the results.

iZotope plug-ins are such a big part of my mixing that they are incorporated into my templates. I can't possibly achieve the quality of mix I'm delivering any other way.

STEVE LEVINE | Music producer, and producer of the BBC radio series, *The Record Producers*

Using vintage equipment on a session can be very creative, however many synthesizers are past their prime and are often plagued with hum, buzz and noise... that's where RX comes to the rescue, allowing you to incorporate these gems into your work whilst maintaining a modern technical level.

When denoising vintage synths, do it in two stages—first get rid of the hum and see how it sounds in the track then tackle the hiss. Be gentle, as you want to keep as much “air” as possible whilst getting rid of any annoying hiss.

Sometimes on my radio show the interviewee has limited time available so any shortcomings in the recording can be easily fixed with RX in post (great for fixing bad digital clips or annoying background

noise).

Another hidden gem in RX is the 24/16 bit conversion. This must be one of the best available—the BBC require 16 bit masters—so I use this.

Sometimes in our radio show we have to use 78 shellac discs. The declicking in RX can make unusable discs perfect for the program.

CHRIS SHAW | Audio engineer for major artists like Bob Dylan, Leonard Cohen, and Public Enemy

I find that Spectral Repair is really good at taking the edge off vocals that are recorded a bit too hot or when a singer's voice is a bit too shrill as he/she hits the top of their range. You can easily see and select the upper harmonics (usually in the 3-5k range) and reduce them without having to resort to complex EQ automation.

Spectral Repair excels at handling difficult de-essing problems. Many times a singer will have a “whistle” embedded in the “ess” frequencies which can be hard to alleviate with a standard de-esser. By using Spectral Repair you can reduce the whistle whilst leaving the “ess” frequencies intact.

I've used the DeClip module to restore the attack and dynamics to drum tracks that were too hot on many archival recordings made with early 16-bit recorders. This immediately removes most of the harsh artifacts made with “vintage” digital gear.

You can use Spectral Repair to remove hi-hat bleed on snare and tom tracks. “Find Similar Events” in the edit menu greatly speeds up the process.

Nothing is better at removing hum from guitar and bass tracks better than RX Denoise. It has saved many unusable tracks for me.

MIKE THORNTON | Audio guru, *Sound on Sound* columnist, and Deputy Editor of Pro Tools Expert

Using two lightweight passes with the denoiser will give better results with fewer artifacts.

Try using the denoiser and target the buzz frequencies with the envelope curve rather than using the Remove Hum module.

The Remove Hum module can be used to remove other problem sounds, not just hum and buzz. I used it to remove mobile phone text alerts.

If you are working on low-frequency problems in Spectral Repair change the Frequency scale to Log or

even Extended Log to display more detail at lower frequencies. If the problems are at higher frequencies try setting the Frequency scale to Linear to display more detail at higher frequencies.

The Declick tool is not just a one trick pony for cleaning up vinyl recording. It's great for removing digital clicks caused by clocking problems, mouth clicks on voiceovers, and some kinds of distortion can be cleaned up with the Declick tool.

I have just taken a recording made in a bedroom with very little soft furnishings, and actually had a ping to it. Ran it through Dereverb and was able to get a result close to as if it had been recorded in a radio studio.

JASON GRAVES | BAFTA award-winning composer for the *Dead Space* video game score

Definitely audition all four algorithms before committing. I always default to the 'D (best, slowest)' algorithm, even though it takes a little longer to preview in real time, but sometimes B or C actually sound better to my ears.

If you're involved in the original recording, always preroll at least three or four seconds of room tone. Then you'll have the perfect handle to train RX and zap the noise effortlessly.

If you're hearing too many artifacts reduce the Noise reduction slider by a few dB and increase the Smoothing by a few dB. Sometimes even 2-4dB sounds significantly more natural.

For extreme denoising cases, I've found two passes with half the dB in Noise reduction yields more transparent results.

This is hands-down THE best audio repair software out there. I use it on a daily basis—nothing sounds as natural and transparent as RX3. Literally every virtual instrument I've recorded and built has been run through RX. You can hear it on all the strings and brass in the *Dead Space* franchise and the latest *Tomb Raider* game. Or should I say you CAN'T hear it!

ADAM AYAN | Gateway Mastering's multi-award-winning mastering engineer

Knowing what to denoise and what to leave alone can be tricky in of itself. My general rule of thumb is if a noise or sound takes me out of the musical moment, or distracts me from enjoyment of the music then it must go!

BOB BRONOW | Emmy-winning sound mixer for *Deadliest Catch*, *Ax Men*, and more

There's always the old chestnut: The squeaky wheel does NOT get the grease. It gets RX, and is never heard again.

Sed minto officitur, quo volori dolorest amust faccuscia denis nulpariatem excerum ut quibusae. Nequodit



Copyright © 2013 iZotope, Inc. All rights reserved. iZotope®, RX®, Ozone®, Insight®, Nectar®, Alloy®, MBIT+™ and the iZotope logo are trademarks or registered trademarks of iZotope, Inc. All other trademarks contained herein are the property of their respective owners. Other products and company names mentioned herein may be trademarks of their respective companies. This material is provided for information purposes only; iZotope, Inc. assumes no liability related to its use. Use of the software is subject to a related license agreement.