
RX 11

Table of Contents

1. Introduction 1
- New in RX 11 1

2. RX Overview 2
- Navigating the manual 2
- Feature comparisons 2
  - RX 11 Audio Editor Feature Comparison 2
  - RX 11 Module Comparison 3
  - RX 11 Plug-in Comparison 4

3. Working with Files 6
- Opening Files 6
- Supported File Formats 6
  - Supported Audio File Formats 6
  - Supported Video File Formats 6
- Supported Channel Configurations 7
  - Multichannel File Support 7
- Creating New Files 7
- Managing File Tabs 7
- Saving Files 8
- Saving RX Documents 8
- Export Options 9
  - Export File 9
  - Export Selection 11
  - Export Regions to Files 12
  - Export Screenshot 12
  - Export History as XML 13
- File Info 13
- Closing Files 14
  - Close One File Tab 14
  - Close Other File Tabs 14
  - Close All File Tabs 14
  - Closing File Tabs With Unsaved Changes 14
  - Closing The Application With File Tabs Open 15

4. Recording in the RX 11 Audio Editor 16
- Overview 16
- Troubleshooting 16

5. Transport Controls & Displays 17
- Overview 17
- Transport 17
- Time and Frequency Readouts 18
- Transport Clock 18
- Time Format Display 18

6. Spectrogram/Waveform Display 19
- Overview 19
  - Anatomy of the Spectrogram Display 19
- Spectrogram Settings 20
  - Spectrogram Type 20
  - FFT Size 21
  - Enable Reassignment 21

Window ... 22
Frequency Scale ... 22
Frequency Overlap ... 22
Time Overlap ... 22
Color Map ... 22
High-Quality Rendering ... 23
Reduce Quality Above ... 23
Cache Size (MB) ... 23
Rulers ... 23
Amplitude Rulers ... 23
Color Map Ruler ... 23
Frequency Rulers ... 23
Waveform Displays ... 24
Waveform Overview ... 24

7. Text Navigation ... 26
Overview ... 26
Workflow ... 26
Search ... 27
Multiple Speaker Detection ... 28
Export Transcript ... 29

8. Interactive Tools ... 30
Overview ... 30
Zoom Tools ... 30
Zoom Selection Tools ... 30
Zoom Sliders ... 30
Navigation Tools ... 31
Grab & Drag Tool ... 31
Dragging Rulers and Using the Mousewheel ... 31
Channel Selectors ... 31
Channel Order (Multichannel) ... 32
Instant Process ... 32
Instant Process: Attenuate ... 32
Instant Process: De-click ... 32
Instant Process: Fade ... 33
Instant Process: Gain ... 33
Instant Process: Replace ... 33
Selection tools ... 33
View Clip Gain ... 34
Selection Modifiers ... 35
[Shift] - Add to selection ... 35
[Alt/Option] - Subtract from Selection ... 35
[Ctrl/Cmd] - Move Playhead without affecting Selection state ... 36
[Mouse Over] - Grab and Drag Selection ... 36

9. Undo History ... 37
Overview ... 37
Export History ... 37
Restore Selection ... 37
Expandable History List ... 38

10. Application Menus ... 39
File ... 39
Edit ... 39
View ... 40

Modules 42
Open Module Window 42
Learn 42
Render 43
Transport 43
Window 44
Help 45

11. Preferences 46
Audio 46
Display 47
Keyboard 48
Misc 49
License and Updates 50
Plug-ins 51
Assistant 52

12. Composite View 53
Overview 53
Workflow 53
Important Notes 53

13. Batch Processor 55
Overview 55
Batch Processor Presets 55
Input 55
Adding Files 55
Removing files 56
Module Chain 56
Output 56
Naming 57
File Format 57

14. Repair Assistant 58
Overview 58
Repair Assistant Module 58
Repair Assistant Plug-in 59
Voice Mode 59
Common Controls 59
Clean Up 60
Tone 60
De-Ess 60
De-Clip 60
Music Mode 61
Common Controls 61
Clean Up 61
Tone 61
De-Harsh 62
De-Clip 62
Repair Assistant Preferences 62

15. Common Module Controls 63
Overview 63
Presets 63
Preset directory locations 63
Preset Dirty Flag 64

Module Footer Controls 64
Preview 64
Pre- and Post-Roll 65
Compare 65

16. Module List 67
Overview 67
Saving module list filters 68

17. Module Chain 69
Overview 69
Controls 69

18. Ambience Match 71
Overview 71
Controls 71
Ambience Type 72
Complex 72
Static 72
Workflow 72
Using Ambience Match module in the RX Audio Editor 72
The Spectrogram in Ambience Match 72
Using Ambience Match as an AudioSuite Plug-In 73

19. Breath Control 74
Overview 74
Controls 74

20. Center Extract 76
Overview 76
Controls 76
Use Cases 77

21. De-bleed 78
Overview 78
Controls 79
Workflow 80

22. De-click 81
Overview 81
Controls 81
Instant Process Tool 82

23. De-clip 83
Overview 83
Controls 83
Using the De-clip Threshold Effect overlay in the Spectrogram/waveform view 85
More Information 85
Suggestions for severe distortion 85
Visual Examples 85

24. De-crackle 86
Overview 86
Controls 86

25. De-ess 87
Overview 87

Controls 87

26. De-hum 89
Overview 89
Mode Selection 89
Controls in Dynamic Mode 90
Controls in Static Mode 91
More Information 93
Ringing 93
Alternative Modules to use for Complex Hum Issues 94
Visual Example 94

27. De-plosive 95
Overview 95
Controls 95
Sensitivity 95
Strength 96
Frequency Limit [Hz] 96
Identifying plosives in the spectrogram display 96

28. De-reverb 97
Overview 97
Controls 97
More Information 99
Early Reflections 99
Using De-reverb as a real-time plug-in 99
Tips for Learning a Reverb Profile 100
Visual Example of De-reverb processing 100
Tips for dealing with complicated reverbs 101

29. De-rustle 102
Overview 102
Controls 102
Reduction Strength 102
Ambience Preservation 102
Separation Algorithm 103

30. De-wind 104
Overview 104
Controls 104
Reduction 104
Crossover Frequency 104
Fundamental Recovery 104
Artifact Smoothing 104
Tips 105
When to use the De-wind Module 105
Alternative Modules 105

31. Deconstruct 106
Overview 106
Controls 106
Tonal Gain 106
Noisy Gain 106
Separate Transients 107
Transient Gain 107
Tonal/Noisy Balance 107
Artifact Smoothing 107

More Information 107

32. Dialogue Contour 108
Overview 108
Displays 108
Waveform and Pitch Curve Display 108
Spectrogram Display 109
Playhead Indicators 109
Contour Curve 109
Contour Curve Axes 109
Contour Curve Readout 110
Contour Curve Editing 110
Add Nodes 110
Semitone Adjustments 110
Time Adjustments 111
Remove Nodes 111
Reset Individual Nodes 111
Controls 111
Pitch 111
Formant 111
Variation 111
Reset Curve 112
Curve Smoothing 112
Formant Scaling 112
Alternative Modules 112

33. Dialogue Isolate 113
Overview 113
Controls (STD and ADV) 113
Voice 113
Reverb 113
Noise 113
Sensitivity 114
Quality 114
Controls (ADV only) 114
Multi-band Interface 115
Alternative Modules 115

34. Guitar De-noise 116
Overview 116
Controls 116
Amp 116
Squeak 117
Pick 118
Processing Tips 119
Alternative Modules 119

35. Interpolate 120
Overview 120
Controls 120
Quality 120

36. Loudness Optimize 122
Overview 122
Metering and Readouts 122
Controls 123

Boost 123
Threshold 123
Ratio 123
Speed 123
Learn 123

37. Mouth De-click 125
Overview 125
Controls 125
More Information 125

38. Music Rebalance 126
Overview 126
Controls 126
Stem Split 127
Getting started with Music Rebalance in Logic 127
Using Music Rebalance (ARA) in Logic 128
Using Multiple Instances 128
Music Rebalance (ARA) Requirements & Limitations 129

39. Spectral De-noise 130
Overview 130
Controls 130
Learn 130
Adaptive Mode 131
Threshold (Noisy/Tonal) 131
Reduction (Noisy/Tonal) 132
Quality 132
Artifact Control 132
Noise Spectrum Display 132
Smoothing 133
Reduction Curve 133
Reset 133
Advanced Settings 133
Algorithm Behavior (Advanced Settings) 134
Smoothing - Advanced 134
Algorithm 134
FFT Size (ms) 134
Multi-Res 135
Noise Floor (Advanced Settings) 135
Dynamics (Advanced Settings) 136
More Information 137
Tips for getting the best results when learning noise profiles manually 137
Learning a Noise Profile From Multiple Selections 137

40. Spectral Repair 138
Overview 138
Attenuate 138
Controls 138
Replace 139
Replace Controls 139
Pattern 140
Pattern Controls 140
Partials and Noise 140
Partials and Noise Controls 141
Surrounding Region Display 141

Workflow 141
Applying Spectral Repair 141
More information 142
Visual example 142
Processing limitations 142
Spectral Repair as an alternative to De-click processing 142
Example use cases for Spectral Repair 143
Increase efficiency with the Find Similar Event Tool 143
Use the Compare Settings window to experiment with Spectral Repair processing 143
Adjust Surrounding Region Length and Before/after weighting to perfect your processing 143

41. Spectral Recovery 144
Overview 144
Controls 144
Learn 144
Spectral Patching 145
Low Gain 145
Low Cutoff 145
High Cutoff 145
High Gain 145
Render 145

42. Streaming Preview 146
Overview 146
Controls 146
Workflow and Presets 147

43. Voice De-noise 148
Overview 148
Controls 148
Voice De-noise Plug-in 149

44. Wow & Flutter 150
Overview 150
Controls 150
Wow 150
Flutter 151
Center Global Pitch 151
Display Wow/Display Flutter 152

45. Azimuth 153
Overview 153
Controls 153
Suggest and Adaptive modes are only available on stereo files 153
More Information 154

46. Dither 155
Overview 155
Controls 155
New Bit Depth 155
Noise Shaping 155
Dither Amount 155
Auto-blanking 156
Limit Noise Peaks 156
Suppress harmonics 156

47. EQ 157
Overview 157
Controls 157
EQ TYPE 157
FREQUENCY PRECISION 157
FREQUENCY and GAIN 157
Q/BANDWIDTH 158
More Information 158
Interacting with the EQ nodes 158
Choosing between Analog and Digital EQ Types 158
EQ curve display & EQ Type 158
Composite EQ Curve Display 158

48. EQ Match 159
Overview 159
Controls 159
AMOUNT 159
Workflow 159
How to Apply EQ Match to different selections 159
How to save a captured spectrum as an EQ Match Preset 160

49. Fade 161
Overview 161
Controls 161
Fade Type 161
Instant Process Tool 161

50. Gain 162
Overview 162
Controls 162
Gain 162
Instant Process Tool 162

51. Leveler 163
Overview 163
Controls 164
Numerical Readouts 164
Optimize For 164
Target Level 164
Responsiveness 165
Preserve Dynamics 165
Ess Reduction 165
Breath Control 165
Limiter 165

52. Loudness Control 166
Overview 166
Controls 166
Integrated loudness 166
Tolerance 166
Short-term/Momentary Loudness 166
True peak 166
Program loudness gate 167
Workflow 167
Loudness Standards 167

53. Mixing 168

Overview 168
Controls 168
Left Output Mix (%) 168
Right Output Mix (%) 168

54. Normalize 169
Overview 169
Target Peak Level 169

55. Phase 170
Overview 170
Controls 170
Adaptive Phase Rotation 170
Rotation (deg) 170
Suggest 170
More Information 170
Visual Example of Phase Rotation 171

56. Plug-in Hosting 172
Overview 172
Selection Based Processing 172
Plug-in Presets 172
Assign Presets to Keyboard Shortcuts 173

57. Resample 174
Overview 174
Controls 174
New sampling rate 175
Change tag only 175
Filter steepness 175
Cutoff shift 175
Pre-ringing 175
Post-limiter 175

58. Signal Generator 177
Overview 177
Silence 177
DC Offset 177
Tones 177
Noise 178

59. Time & Pitch 179
Overview 179
iZotope Radius 179
Controls 179
Algorithm 179
Formant Correction 180
Stretch Ratio 180
BPM Calculator 180
Pitch Shift 180
Transient Sensitivity 180
Noise Generation (Radius Mode Only) 181
Pitch Coherence (Radius Mode Only) 181
Phase Coherence (Mix Mode Only) 181
Adaptive Window Size (ms) (Solo Mode Only) 181
Shift Formants 182
Strength 182

Shift 182
Width 182
More Information 182

60. Variable Pitch 183
Overview 183
Displays 183
Waveform Display 183
Spectrogram Display 184
Playhead Indicators 184
Contour Curve Display 185
Contour Curve Axes 185
Contour Curve Readout 185
Contour Curve Editing 185
Add Nodes 186
Semitone Adjustments 186
Time Adjustments 186
Remove Nodes 186
Reset Individual Nodes 186
Reset Curve 186
Smoothing 186
Controls 187
Preserve Time 187
Pitch Coherence 187
Transient Sensitivity 188
Compare Settings 188
Alternative Modules 188

61. Variable Time 189
Overview 189
Displays 189
Waveform Display 189
Spectrogram Display 190
Contour Curve Display 190
Contour Curve Axes 190
Contour Curve Readout 191
Contour Curve Editing 191
Add Nodes 191
Time Stretch Ratio Adjustments 191
Time Adjustments 191
Remove Nodes 192
Reset Individual Nodes 192
Reset Curve 192
Smoothing 192
Controls 192
Pitch Coherence 192
Transient Sensitivity 193
Compare Settings 193
Alternative Modules 193

62. Find Similar 194
Overview 194
Controls 194

63. Markers & Regions 195
Overview 195

Controls 195
More Information 195

64. Spectrum Analyzer 196
Overview 196
Controls 196
Channel View Selection Buttons 196
Peak Finding 196

65. Waveform Statistics 197
Overview 197
Controls 197
More Information 198
True peak 198
RMS 198

66. RX Plug-ins 199
Plug-in Formats 199
Learning in Audiosuite 199
Method 1 - AS Learn (Audiosuite Learn) 199
Method 2 - UI Learn 200
Presets 200
History 201
I/O Meters 201
Options 201
General Options 201
I/O Options 202
Latency 202

67. RX Monitor 203
Overview 203
Workflow 203

68. RX Connect 204
Overview 204
Controls 204
DAW/NLE specific instructions 204
Adobe Audition CC RX Connect Workflow 205
Avid Pro Tools RX Connect Workflow 205
Steinberg Cubase and Nuendo RX Connect Workflow 206
Using RX as an External Audio Editor 207
Adobe Premiere Pro CC with RX as an external audio editor 207
Apple Logic Pro X with RX as an external audio editor 207
DaVinci Resolve with RX as an External Editor 208

69. RX Spectral Editor 209
Overview 209
Getting Started with Spectral Editor (ARA) 209
Troubleshooting 212
Controls 212
Processing Mode 212
Selection Tools 212
General Controls 212
Processing Modes 213
Attenuate 213
Replace 214
Gain 215

De-click ... 215
De-hum ... 216
Working with Comp Takes in Logic Pro ... 217
Keyboard Shortcuts ... 217
ARA Requirements and Limitations ... 217

70. Keyboard Shortcuts ... 219

71. Identifying Audio Problems ... 228
RX Processing Step Flowchart ... 228
Using the Spectrogram Display ... 230
Hum ... 230
Buzz ... 230
Hiss and other Broadband Noise ... 231
Clicks, Pops, & Short Impulse Noises ... 231
Clipping ... 231
Intermittent Noises ... 232
Gaps and Drop Outs ... 233

72. Authorization ... 234
Trial Mode ... 235
Demo Mode ... 235
Authorization Methods ... 236
More Authorization Help ... 236

73. License Information ... 237
Abseil ... 237
Anti-Grain Geometry ... 237
ARA ... 237
ARM_NEON_2_x86_SSE ... 238
base64 ... 238
Better Enums ... 239
Boost ... 239
Bravura ... 240
C++ Rest SDK ... 241
Caffe ... 243
cerberus ... 244
composable_kernel ... 244
coremltools ... 245
CpuInfo ... 246
Crashpad ... 246
curl ... 247
dcleblanc/SafeInt ... 247
Distributed Machine Learning Common Codebase ... 247
dlfcn-win32 ... 248
DLPack ... 248
DLPack: Open In Memory Tensor Structure ... 248
Eigen ... 248
emsdk ... 249
FarmHash ... 249
FLAC ... 249
FlashAttention ... 250
FlatBuffers ... 250
FreeBSD: getopt.c file ... 251
FreeType ... 251
FXDev ... 251

G3log: Asynchronous logger with Dynamic Sinks 252
gemmlowp 252
GLEW 252
google/glog 254
google/nsync 255
google/re2 255
google/sentencepiece 255
google/XNNPACK 256
gsl 256
HalideIR 257
HowardHinnant/date 257
huggingface/transformers 258
IcoMoon 258
Intel Math Kernel Library (Intel MKL) 258
Intel neural-compressor 259
intel/dnnl 260
Intel® Integrated Performance Primitives (Intel® IPP) 260
JDAI-CV/DNNLibrary 261
JsonCpp 261
LAME 261
libb64 262
libprotobuf-mutator 262
LibXML2 262
madler/zlib 263
magic_enum 263
Microsoft Cognitive Toolkit (CNTK) 264
microsoft/mimalloc 264
microsoft/wil 265
mpi4py 265
MurmurHash3 266
Netlib numeralgo na10 Aberth's method 266
nlohmann/json 267
nvidia/cutlass 267
OGG / Vorbis 267
onnx 268
onnx/onnx-tensorrt 268
onnxruntime 269
Open MPI 269
Opus 271
PIL 272
portmidi 272
posix pthread library 273
protocolbuffers/protobuf 273
pthreadpool 274
Pytorch / Caffe2 274
pytorch/cpuinfo 275
react-native 276
readerwriterqueue 276
Resemblyzer 277
Roboto font family 277
ruy 277
Skia 278
SQLite 278
TagLib 278
Tencent/rapidjson 279
Tensorflow 280

TinyXML 280
Tipue Search 281
vectorize 281
WebView2 282
xsimd 282
Yoga 283
ZeroMQ 283
zlib 284

1. Introduction



Welcome to RX 11!

RX 11 features new machine learning algorithms for dialogue isolation and dialogue de-reverb that work in real-time. We've added new tools to preview audio for streaming and to optimize loudness. Additionally, there are major improvements to Repair Assistant, as well as higher-quality source separation. We've also expanded the list of compatible hosts for the Spectral Editor ARA plugin.

New in RX 11

- Dialogue Isolate: Combines dialogue de-noising and de-reverbing in a single, realtime plugin and in an updated module. The Advanced plugin and module also include multi-band processing and a new offline high-quality mode.
- Streaming Preview: Hear music the way that the rest of the world will by applying the same loudness normalization and lossy data encoding that streaming platforms will apply.
- Loudness Optimize: Optimize perceived loudness and minimize measured LUFS of your track prior to posting it on streaming platforms with level normalization.
- Music Rebalance: A new neural net provides a largely improved quality of musical source separation as well as per-stem Sensitivity adjustment.
- Repair Assistant: The addition of machine learning technology provides better processing and faster learning of your audio.
- Dialogue Contour: Change the character and expressivity of a voice with new Formant and Variation controls, and see those changes reflected get in the new pitch curve.
- Mid/Side Channel Mode: Visualize and edit audio in Mid/Side configuration in the main RX window.
- Signal Generator: New **Pulses** tone shape, new **Modulate** pasting mode.
- Preferences: True peak calculation mode allows selection of the oversampling algorithm for maximal compliance of true peak level measurements.

2. RX Overview

iZotope's award-winning RX Audio Editor is the industry standard for audio repair, restoration, and enhancement. It offers a comprehensive suite of tools focused on alleviating common to complex audio issues. Post production professionals, audio engineers, and video editors alike use RX to transform problematic recordings into production-ready audio.

Navigating the manual

This help guide is shared by RX 11 Standard and RX 11 Advanced. The following tag is used throughout the manual to differentiate Advanced only features: **ADV**

Feature comparisons

The following tables outline the differences between Standard and Advanced. The rightmost column in each of the tables indicates features that are new or improved in RX 11.

RX 11 Audio Editor Feature Comparison

Features	Standard	Advanced	New/Improved
Batch Processor	X	X	
Clip Gain	X	X	
Composite View	X	X	
Expandable History list	X	X	
Find Similar	X	X	
First Time User Tour	X	X	Improved
Instant Process	X	X	
Markers & Regions	X	X	
Mid/Side Mode	X	X	New
Module Chain	X	X	
Module List View Filters	X	X	
MP3 Export	X	X	
Multichannel Processing		X	
Plug-in Hosting	X	X	
Recording & Monitoring	X	X	
Repair Assistant	X	X	Improved
Restore Selection	X	X	
Spectrum Analyzer	X	X	
Spectral Editing Tools	X	X	
Text Navigation	X	X	
Time/Frequency Feathering	X	X	
Waveform Statistics	X	X	

RX 11 Module Comparison

Module Name	Standard	Advanced	New/Improved
Ambience Match		X	
Azimuth		X	
Breath Control	X	X	
Center Extract		X	
De-bleed	X	X	
De-click	X	X	
De-clip	X	X	
De-crackle	X	X	
De-ess	X	X	
De-hum	X	X	
De-plosive	X	X	
De-reverb	X	X	
De-rustle		X	
De-wind		X	
Deconstruct		X	
Dialogue Isolate with De-reverb	X	X	New
Dialogue Contour		X	Improved
Dither	X	X	
EQ	X	X	
EQ Match		X	
Fade	X	X	
Gain	X	X	
Guitar De-noise	X	X	
Interpolate	X	X	
Leveler		X	
Loudness Control	X	X	
Loudness Optimize	X	X	New
Mixing	X	X	
Mouth De-click	X	X	
Music Rebalance	X	X	Improved
Normalize	X	X	
Phase	X	X	
Repair Assistant	X	X	Improved
Resample	X	X	
Signal Generator	X	X	
Spectral De-noise	X	X	
Spectral Recovery		X	
Spectral Repair	X	X	

Module Name	Standard	Advanced	New/Improved
Streaming Preview	X	X	New
Time & Pitch	X	X	
Variable Pitch	X	X	
Variable Time	X	X	
Voice De-noise	X	X	
Wow & Flutter		X	

RX 11 Plug-in Comparison

The RX 11 installer will give you the option to install plug-ins in the following formats:

- **AU**
- **AAX**
- **AAX Audiosuite**

> RX 11 Ambience Match and De-rustle plug-ins are AAX Audiosuite for Pro Tools only.

- **VST3**
- **AU ARA**
- **VST3 ARA**
- **AAX ARA**

> Spectral Editor ARA is available in Logic Pro (AU), StudioOne (VST3), and Pro Tools (AAX, via an upcoming update from Avid). Music Rebalance ARA is only compatible with Logic Pro (AU).

All plug-in formats are 64 bit only. VST2 is no longer supported.

RX Plug-ins	Elements	Standard	Advanced	New/Improved
Ambience Match			X	
Breath Control		X	X	
Connect		X	X	
De-click	X	X	X	
De-clip	X	X	X	
De-crackle		X	X	
De-ess		X	X	
De-hum	X	X	X	
De-plosive		X	X	
De-reverb	X	X	X	
De-rustle			X	
Dialogue Isolate		X	X	Improved
Guitar De-noise		X	X	

RX Plug-ins	Elements	Standard	Advanced	New/Improved
Monitor		X	X	
Mouth De-click		X	X	
Music Rebalance		X	X	Improved
Repair Assistant	X	X	X	Improved
Spectral De-noise		X	X	
Spectral Editor		X	X	Improved
Voice De-noise	X	X	X	

3. Working with Files

Opening Files

RX supports opening up to 32 audio files at a time. Files can be opened in the RX Audio Editor using the following methods:

1. Navigate to the File menu, select **Open…** and choose the files from the system dialog that appears. Alternatively, the following keyboard shortcuts can be used to launch the **Open…** system dialog: `Command`+`O` (Mac) or `ctrl`+`O` (Windows).
2. Drag and drop files into the main editor window to open them in a new file tab.
3. Drag a file from Finder/Windows Explorer onto the RX Audio Editor icon in the Dock/Desktop.
4. Click on the **Open file** button that appears in the RX Audio Editor window when no files are loaded.
5. Double-click on the RX logo in the middle of the RX Audio Editor interface when no files are loaded.

Supported File Formats

A number of different file formats can be opened and edited in the RX Audio Editor. The next three sections outline the supported file formats and channel count configurations that can be opened in the RX Audio Editor. For information about file formats when saving or exporting files in RX, see the Saving Files and Export Options sections below.

Supported Audio File Formats

The following audio file formats can be opened in the RX Audio Editor: **AAC**, **AAX (Audible Audiobook Format)**, **AIFF/AIF**, **BWF**, **CAF**, **FLAC**, **M4A**, **MP3**, **OGG**, **SD2**, **WAV**, **WMA**

Supported Video File Formats

The RX Audio Editor supports loading video file formats, but **does not support video playback.** When a video file is opened in RX, *only the audio data from that file will be imported*.

The following video file formats can be opened in the RX Audio Editor: **AVI**, **M4V**, **MOV** (*RX requires that Quicktime is installed in order to open Quicktime file formats, e.g. .mov files.*), **MPEG**, **MPV**, **WMV**.

FILE FORMAT DEPENDENCIES

Some file formats may have dependencies based on your operating system that may prevent you from importing them into the RX Audio Editor. For example, Windows native formats (like WMA and WMV) may not open on Mac and QuickTime formats (like AAC, MOV, and M4V) may require installing QuickTime on Windows and running the RX Audio Editor in 32-bit mode.

Supported Channel Configurations

The RX 11 Standard Audio Editor and RX 11 Advanced Audio Editor support opening mono and stereo audio files.

> **TIP: OPTION FOR OPENING SPLIT STEREO FILES IN ONE TAB**
>
> Mono audio files with (.L and .R) or (.1 and .2) extensions can be opened as either mono files (2 mono tabs) or split stereo (1 stereo file tab). See Misc for more information. Note that this option is only applicable to split stereo files and does not apply to split surround files.

Multichannel File Support

The RX 11 **Advanced** Audio Editor supports opening audio files with up 10 channels per file tab. Multichannel audio device settings can be configured in the "Audio" tab of the Preferences menu.

The channel selector labels can be configured by selecting an option in the Channel Order menu. To access the Channel Order menu, right-click on the Time ruler and navigate to the "Channel Order" sub-menu. The options available in the Channel Order menu depend on the number of channels in the active file tab.

Creating New Files

To create a new file in RX:

1. Open the File menu
2. Select "New..."
3. You will be prompted for the name, sample rate and channel count of the new file you are creating.

> **TIP: CREATE A NEW FILE FROM THE CONTENTS OF THE CLIPBOARD**
>
> If you have existing audio data in your clipboard (for example, if you have copied a selection from an existing file in RX), you can create a new file based on that audio data.
>
> Open the "File" menu, choose "New from Clipboard" or use the keyboard shortcut: `Command+Shift+N` (Mac) or `Ctrl+Shift+N` (Windows). *The new file will match the sample rate and channel count of the audio data present on your clipboard.*

Managing File Tabs

RX supports having up to 32 files open at once. You can navigate between tabs by clicking on a tab or using the following keyboard shortcuts:

- **Select tab to the right of current tab:** Control+Tab (Mac) or Alt+Tab (Windows)
- **Select tab to the left of current tab:** Control+Shift+Tab (Mac) or Alt+Shift+Tab (Windows)

If you have multiple files open, an arrow button () will appear to the right of the last visible tab. You can access file tabs that are not currently visible by clicking on the arrow button and selecting a tab from the menu.

▶ Close any tab by **clicking it with the center mouse button.**

Saving Files

There are a number of ways to save a file in the RX Audio Editor. The Save Operations include:

Name	Description	Default Mac Shortcut	Default Windows Shortcut
Save	**For uncompressed file formats (.wav or .aiff):** Overwrites the original file on disk	`Command+S`	`Ctrl+S`
	For compressed file formats: Opens the Export File dialog	`Command+S`	`Ctrl+S`
Save As…	**For uncompressed file formats (.wav or .aiff):** Save a copy of your file using the same file format	`Command+Shift+S`	`Ctrl+Shift+S`
	For compressed file formats: Opens the Export File dialog	`Command+Shift+S`	`Ctrl+Shift+S`
Save RX Document	Saves file as .rxdoc file extension (more information below)		
Save RX Document As…	Saves copy of your .rxdoc file		

AUTOSAVE

The RX Audio Editor will automatically save backups of your editing session by default. When the RX application is launched, it will open your most recent editing session. The option to turn it off is located under the Preferences > Misc tab as "Resume last editing session when app starts."

Saving RX Documents

You can save a file using the RX Document file format (.rxdoc) to archive your edits. An RX Document includes your original file, all the edits you've made to it, and your most recent selection and view state. **RX Documents can only be opened in the RX Audio Editor.** If you need to save your file so it can be opened somewhere else (like a DAW or media player), you need to export it in another format (like WAV or AIFF).

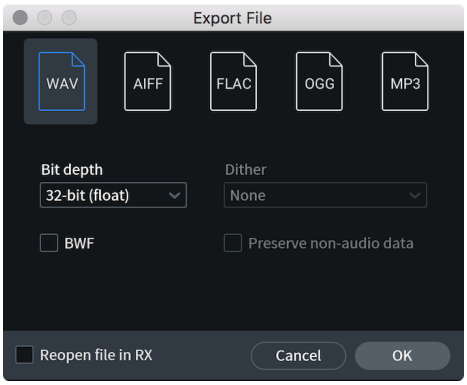
To save an RX Document, select File > Save RX Document… and select where you would like to store the file. Keep in mind that the size of the RX Document file can be very large, especially if your list of edits include multiple processes on the whole file.

Export Options

There are a number of different export options in the RX Audio Editor:

- Export File
- Export Selection
- Export Regions to Files
- Export Screenshot
- Export History as XML

Export File



1. Select File... > Export
2. Select the file format you want to Export to and adjust the associated settings as desired (available settings explained in the table below)
3. Click "OK"
4. In the system window, name your file and choose where you would like to save it to
5. Click "Save" to export your file. Checking the **Reopen file in RX** checkbox will open your exported file in the RX 11 Audio Editor after the export completes successfully.

Export Format Options

The following file formats are available when exporting files from the RX 11 Audio Editor:

- WAV - Uncompressed
- AIFF - Uncompressed
- FLAC - Compressed: Lossless Compression
- OGG - Compressed: Lossy Compression
- MP3 - Compressed: Lossy Compression

Uncompressed File Formats

The RX 11 Audio Editor allows you to export files to the following uncompressed file formats: WAV and AIFF.

WAV

The following options are available when exporting files to this file format:

- **Bit Depth:** Determines the bit depth of the exported file. Choices include: **16 bit, 24 bit, 32 bit (float), 32 bit (int).**
- **Dither**: Determines the Dither Type to be applied to the exported file. Choices include: **None, White Noise (TPDF), Noise shaping (MBIT+)**
- **BWF**: When selected, the file will be exported with extended information that is included in the file header of a Broadcast Wave File.
- **Preserve Non-Audio Data**: When selected, the exported file will retain the metadata of the original file.

AIFF

The following options are available when exporting files to this file format:

- **Bit Depth:** Determines the bit depth of the exported file. Choices include: **16 bit, 24 bit, 32 bit (float),** and **32 bit (int).**
- **Dither**: Determines the Dither Type to be applied to the exported file. Choices include: **None, White Noise (TPDF),** and **Noise shaping (MBIT+).**
- **Preserve Non-Audio Data**: When selected, the exported file will retain the metadata of the original file.

Compressed File Formats

The RX 11 Audio Editor allows you to export files to the following compressed file formats: FLAC (Lossless), OGG (Lossy), and MP3 (Lossy).

FLAC (Lossless)

The FLAC file format offers lossless compression. **The following options are available when exporting files to this file format:**

- **Bit Depth**: Determines the bit depth of the exported file. Choices include: **8 bit, 16 bit,** and **24 bit**
- **Dither**: Determines the Dither Type that will be applied to the exported file. Choices include: **None, White Noise (TPDF),** and **Noise shaping (MBIT+)**.
- **Compression Level**: Adjusts the compression strength of the FLAC encoder. Stronger compression requires more CPU time during file encoding but results in a slightly smaller file. FLAC compression setting does not result in any quality change to the signal since FLAC is a lossless format.

OGG (Lossy)

The Ogg Vorbis file format offers lossy compression. **The following options are available when exporting files to this file format:**

- **Quality:** Adjusts the bitrate of the Ogg Vorbis compression algorithm. Higher bitrate values result in higher audio quality, but also increase the file size
- **Prevent Clipping:** See the Prevent Clipping section below for more information.

MP3 (Lossy)

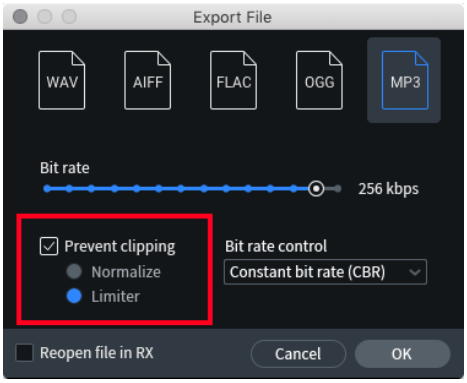
The mp3 file format offers lossy compression. **The following options are available when exporting files to this file format:**

- **Bit rate:** Adjusts the bit rate of the MP3 compression algorithm. Higher bit rates result in higher quality audio but will increase the file size.

- **Bit rate control:** Determines how (or if) bit rate varies over time. Choices include: **Constant bit rate (CBR), Average bit rate (ABR),** and **Variable bit rate (VBR).**
- **Prevent Clipping:** See the Prevent Clipping section below for more information.

Prevent Clipping

Predicts and prevents codec clipping when exporting audio in lossy formats (MP3 and OGG) by checking for decoded levels and adjusting levels of the original signal.



> **PREVENT CLIPPING PROCESSING TIME**
>
> **Prevent Clipping** may run significantly slower than regular encoding, since it computes the correct level adjustment depending on the amount of clipping occurring in the file. Files with **little to no codec clipping** usually will encode **quickly**, whereas **heavily clipping** files may take **longer**.

There are two types of file level adjustments that can be applied:

- **Normalize**: attenuates **overall level of the file** to ensure that the encoded/decoded file does not exceed 0 dBTP.
- **Limiter**: attenuates **parts of the file** that could become clipped to retain the level of non-clipping sections, while overall true peak levels are limited to 0 dBTP.

> **CHOOSING NORMALIZE OR LIMITER**
>
> The **Limiter** will leave larger sections of the file unchanged in level and will only attenuate sections that would experience clipping. However, like any dynamic processing, this may create pumping. The **Normalize** mode can completely avoid pumping at the expense of slightly reducing the overall level of the file.

Export Selection

This option will allow you to export only the audio that is contained within your current selection, as opposed to the entire audio file.

1. Select File > Export Selection, and the Export File dialogue box appears.
2. Follow the additional aforementioned steps.

Export Regions to Files

This option allows you to export each region in the currently selected file to an individual audio file.

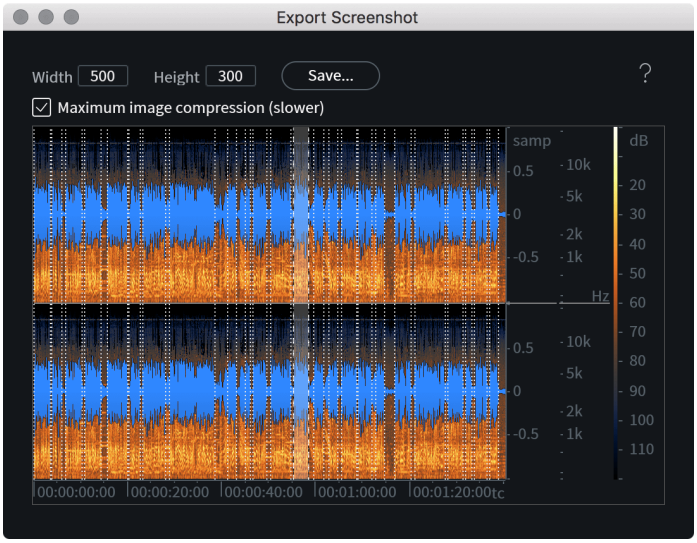
1. Open the **File** menu and choose **Export Regions to Files…**
2. Choose the export file format in the Export Files dialog and click “OK”
3. In the **Export Regions to Files** system dialog:
   a. Choose the destination for the exported files. By default, the destination is set to the location of the source file.
   b. Enter a suffix to append to all exported region filenames. When the default suffix of `(none)` is used, a suffix will not be added to the exported file names.

> **NOTE**
> The names of the exported files will match the names of the regions with the optional suffix appended to the filename. If any regions share the same name, a number will be appended sequentially to the exported filenames to avoid duplication.

4. Click **Save** to export.

Export Screenshot

This option allows you to export your current Spectrogram/Waveform display as a PNG image file. This can be very helpful for archiving any restoration process or for forensic documentation.



When clicking on Export Screenshot from the File menu, your current Spectrogram/Waveform view will be used for adjusting your screenshot size and position.

> The Spectrogram/Waveform transparency balance must be set before selecting File > Export Screenshot as this cannot be changed in this window.

To define the size of your screenshot, simply click and drag in order to enlarge or shrink the screenshot window. The dimensions of your resulting screenshot will update automatically, however these can also be entered manually by clicking once in either Width or Height.

The max resolution attainable for your screenshot will be limited by the individual computer's screen resolution.

When you are finished changing the dimensions of your screenshot, click on the Save button to name and save your .PNG screenshot to your chosen directory.

To save screenshots faster (at the expense of having a larger file on disk), disable Maximum image compression.

Export History as XML

Export the Undo history list of your current file tab to an .xml document.

File Info

The File Info window can be opened by navigating to Window menu > File Info. There are two sections in the File Info window: General Info and More Info. The More Info section lists information dependent on the file type. The following information is available in the General Info and More Info sections of the File Info window:

- **General Info**
  - Name: The current filename
  - Duration: Length of the file
  - Sampling rate: The original sampling rate of the file
  - Bit depth: The original bit depth of the file
  - Channels: Mono or stereo
  - Size on disk: Size of the file in bytes
- **More Info**
  - Timecode
  - Created by
  - Originator reference
  - Date created
  - Time created
  - BWF version
  - Coding history
  - Track Title
  - Artist
  - Album
  - Date
  - Track Number
  - Comment
  - Genre

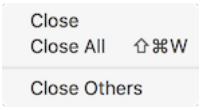
Closing Files

The following sections describe different methods availble for closing file tabs in the RX Audio Editor.

Close One File Tab

Single file tabs can be closed using the following methods:

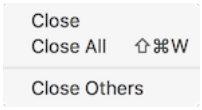
- Single-click on the 'x' button in the file tab display
- Select the "Close file" option in the File menu
- Right-click on any file tab and select "Close" from the context menu



- **Keyboard shortcuts:** `Command+W` (Mac) or `ctrl+W` (Windows)

Close Other File Tabs

To keep one file tab open and close all other file tabs: Right-click on the file tab that should remain open and select "Close others" from the context menu.



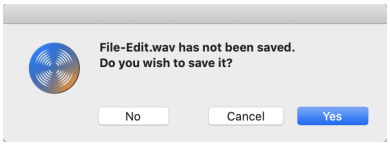
Close All File Tabs

All file tabs can be closed using the following methods:

- Select the "Close all files" option in the File menu
- Right-click on any file tab and select "Close all" from the context menu
- Keyboard shortcuts: `Command+Shift+W` (Mac) or `ctrl+Shift+W` (Windows)

Closing File Tabs With Unsaved Changes

If a file has been edited or processed in the RX Audio Editor and the changes have not been saved, a small dot will appear in the corner of the file tab to indicate that there are unsaved changes. When closing file tabs that have unsaved changes, a prompt will be displayed before the file is closed. The prompt will include options to save, revert changes or cancel before closing the modified file.



The following options are available in the prompt:

- **Yes**: The modified file will be saved as an RX Document file (.rxdoc), a system window will appear to select the save location for the file before closing the tab.
- **No**: Unsaved changes will be discarded and the file tab will be closed.

- **Cancel**: The prompt will be dismissed and the file tab will remain open.

Closing The Application With File Tabs Open

The RX Audio Editor application will open all file tabs present when it was last closed if the “Reopen previous audio files when app starts” option in the Preferences > Misc tab is enabled.

If this option is enabled and the application is closed when files with unsaved changes are present, a prompt will not be displayed. Any unsaved changes will be stored in the RX session data folder and will load the next time the application is opened.

If this option is disabled, a prompt will appear to save or discard changes when closing the application if any file tab has unsaved changes. A separate prompt will appear for each file tab with unsaved changes. If any of the prompt dialogs are canceled, the application will remain open.

4. Recording in the RX 11 Audio Editor

Overview

RX supports recording up to two channels at a time.

To record in the RX 11 Audio Editor:

1. Create a new file.
2. Press the Record button once to arm recording. The Record button will flash red when RX is armed to record. The meters to the right of the transport controls will update based on your input signal when recording is armed.
3. Before recording, you should ensure that your input levels are not clipping and allow for adequate headroom. Alternatively, you can enable input monitoring to set input levels without engaging record arm.
4. After adjusting your input levels, you can start recording by clicking the Record button again. When RX is recording, the Record button will display as solid red.
5. You can stop recording by clicking the Record button again.
6. After you have stopped recording you can edit and apply processing to the file.

RX 11 SESSION DATA FOLDER

After recording, your recorded audio data is stored in the RX 11 Session Data folder. You can set the location of the RX Session Data folder in Misc. If you use the recording functionality in the RX Audio Editor often, it is recommended that the RX Session Data folder be located on a drive with a sufficient amount of free space.

Troubleshooting

If you are having trouble recording in RX, try the following steps:

1. Enable **Input Monitoring** and look for activity on RX's level meters.
2. Close other audio applications, DAWs and NLEs open on your computer to make sure no other program is usurping the sound card.
3. Open **Preferences** > **Audio** and make sure the correct device is listed in **Input Device**. Also check in the **Channel Routing** dialog to make sure the correct inputs are selected.
4. Check your input source. Make sure the hardware connections between whatever you're recording from and the inputs on your audio interface are correct.

5. Transport Controls & Displays

Overview

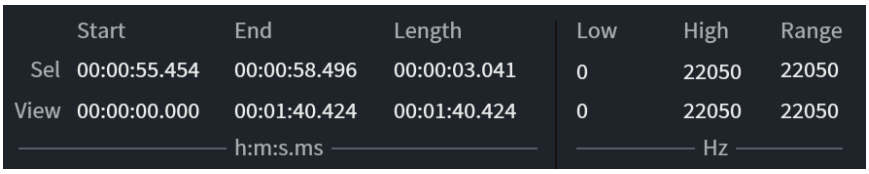
The following sections include information about the different transport controls, readouts, and displays you will encounter in the RX Audio Editor application.

Transport



Name	Description
INPUT MONITOR	When enabled, allows you to monitor input signal to set levels prior to recording. Input source is configured in the RX Audio Editor Preferences
RECORD	Begins recording into a new file. One click puts recording into an armed state for safely setting input levels. The next click begins recording. A third click stops recording. If you already have a file open, hitting Record will prompt you to create a new file for recording.
REWIND [Enter/ Return]	Brings you back to the start of the file.
PLAY [Spacebar]	Starts and stops playback. Starts or stops recording if Record is armed.
PLAY SELECTION ONLY	When you've made a selection of a time range, frequency range, or both, this button auditions just the selection (useful for isolating intermittent noises, etc.)
LOOP Ctrl+L (Windows) Cmd+L (Mac)	Enable this switch to loop the selected audio.
PLAYHEAD FOLLOWS PLAYBACK Ctrl+R (Windows) Cmd+R (Mac)	Toggles the behavior of the playhead on stop. If this is enabled, the playhead will return to the anchor sample (the position before playback began).
PLAYHEAD	To place the playhead, single click anywhere in the Spectrogram/Waveform display. To playback audio while positioning the playhead, click and drag the playhead icon or hold Ctrl (Windows) or Command (Mac) while clicking and dragging in the Spectrogram/Waveform display.

Time and Frequency Readouts



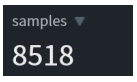
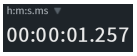
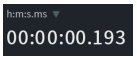
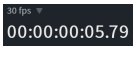
- The time display (on the left in the image above) shows the start, end and length time values of the current selection and the current view. The time format used in this display is selected in the time format menu, explained in the section below.
- The frequency display (on the right in the image above) shows the low frequency boundary, high frequency boundary and frequency range values for the current selection and the current view range.
- Clicking on any of these fields allows you to manually enter values.

Transport Clock

Indicates the current position of the playhead. Depending on time format display setting, this value is shown in hours/minutes/seconds, time code, or samples.

Time Format Display

You can change the time format display used in the time ruler and time readouts. It can be accessed by right-clicking the time ruler or by clicking the arrow button to the left of the currently selected time format label.

Format	Name & Description
	**Samples:** The sample counter, starting from 0
	**Time (h:m:s):** Time in hours, minutes, seconds, and milliseconds, starting from 0
	**Timecode (n fps):** The time code in hours, minutes, seconds, and frames, starting from 0
	**Source Time (h:m:s):** Time in hours, minutes, seconds, and milliseconds, starting from the clip's timecode origin
	**Source Timecode (n fps):** The time code in hours, minutes, seconds, and frames, starting from the clip's timecode origin. The value of n (the frame rate of the time code) is determined by the Time Scale Frame Rate setting in the Misc tab of the Preferences window

6. Spectrogram/Waveform Display

Overview

The RX Audio Editor features a rich visual environment for editing and repairing audio. The central focus of the interface is the Spectrogram/Waveform display. It combines an advanced Spectrogram with a waveform transparency overlay to provide frequency and amplitude information in one highly configurable window.

Using the spectrogram to identify audio problems

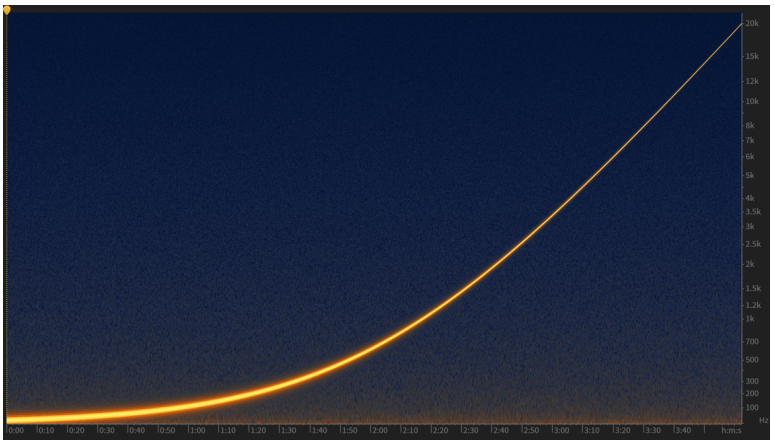
See the Identifying Audio Problems chapter for tips on using the spectrogram to spot common audio issues.

Anatomy of the Spectrogram Display

The spectrogram allows you to visualize both frequency and amplitude information of an audio recording in one display.

Frequency

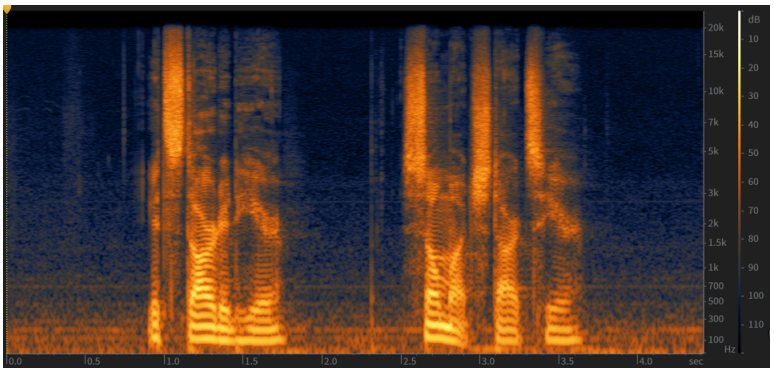
The Spectrogram shows frequency information across the vertical axis. Lowest frequency content is displayed at the bottom, highest frequency content is displayed at the top.



This image shows the spectrogram of a sine sweep over pink noise. The sine sweep starts at 20 Hz (bottom of the display) and sweeps to 20 kHz (top of the display) over 4 minutes.

Amplitude and Color

The amplitude of frequency content is indicated by variations in color in the Spectrogram. The color map ruler (to the right of the frequency ruler) shows the color being used to represent a given amplitude value.



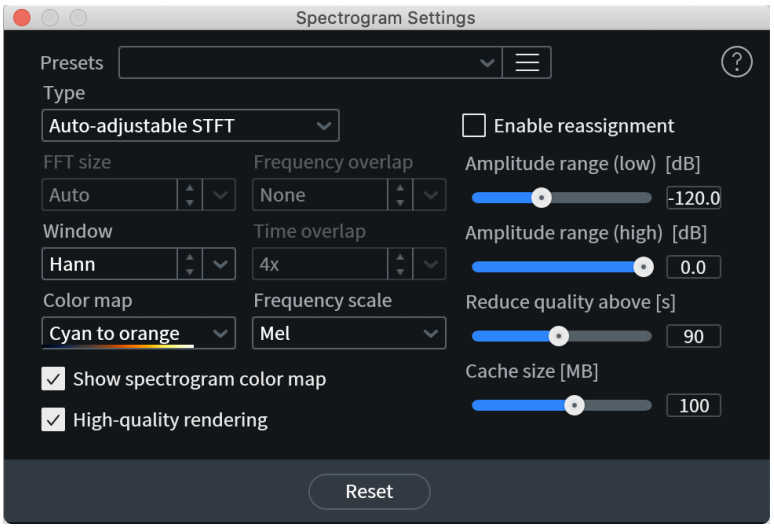
In this example, Louder events (speech) are indicated by brighter colors (yellow/bright orange) and quieter events (breaks in speech and noise floor) are indicated by darker colors (dark orange, blue, black)

Spectrogram Settings

The RX Spectrogram is highly configurable, you can adjust the default configuration, load a preset or save your own preset in the Spectrogram Settings window.

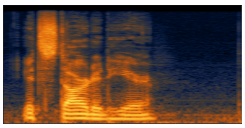
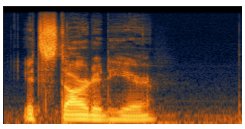
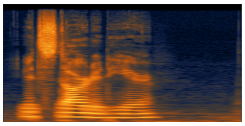
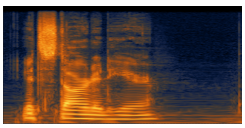
The Spectrogram Settings window can be opened using the following methods:

- From the "View" menu of the RX Audio Editor.
- By right-clicking on the spectrogram display and selecting **Spectrogram Settings** from the context menu.
- Using a keyboard shortcut: `command+shift+,` (on Mac) or `ctrl+shift+,` (on Windows).



Spectrogram Type

RX offers different methods for displaying time and frequency information in the Spectrogram. RX's advanced Spectrogram modes allow you to see sharper time (horizontal) and frequency (vertical) resolution simultaneously. There is always a trade-off of display quality versus processing time, so keep in mind that some modes will take longer to draw on the screen than others.

TYPE	DESCRIPTION
REGULAR STFT 	Most common spectrogram type (can be found in other editors) It has a fixed uniform time-frequency resolution. This is the simplest and fastest drawing mode in RX.
AUTO-ADJUSTABLE STFT 	Automatically adjusts FFT size (i.e. time and frequency resolution of a Spectrogram) according to the zoom level. For example, if you zoom in horizontally (time) you'll see that percussive sounds and transients will be more clearly defined. When you zoom in vertically (frequency), you'll see individual musical notes and frequency events will appear more clearly defined.
MULTI-RESOLUTION 	Calculates the Spectrogram with better frequency resolution at low frequencies and better time resolution at high frequencies. This mimics psychoacoustic properties of our perception, allowing the Spectrogram display to show you the most important information clearly.
ADAPTIVELY SPARSE 	Automatically varies the time and frequency resolution of the Spectrogram to achieve the best Spectrogram sharpness in every area of the time-frequency plane. This often lets you see the most details for a thorough analysis, but it's the slowest mode to calculate.

FFT Size

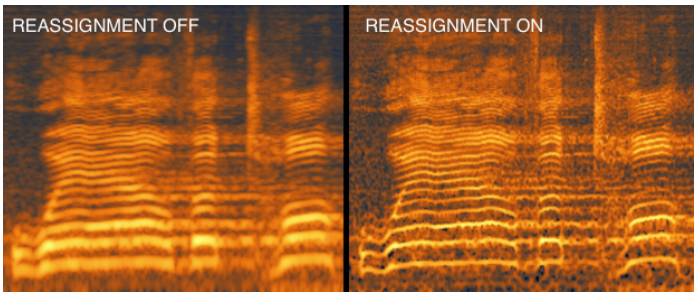
The greater the FFT size, the greater the frequency resolution, i.e. notes and tonal events will be clearer at larger sizes. However, choosing a larger number here will make time events less sharply defined because of the way this type of processing is done. Choosing Auto-adjustable or Multi-resolution modes allows you to get a good combination of frequency and time resolution without having to change this setting as you work.

Fast Fourier Transform (FFT)

A procedure for the calculation of a signal frequency spectrum. The greater the FFT size, the greater the frequency resolution, i.e., notes and tonal events will be clearer at larger sizes.

Enable Reassignment

Enables a special technique for Spectrogram calculation that allows very precise pitch tracking for any harmonic components of the signal. When used together with Frequency Overlap/ Time Overlap controls, this option can provide virtually unlimited time and frequency resolution simultaneously for signals consisting of tones.



Window

Selects between the different weighting functions (or windows) that are used for the FFT analysis. Window functions control the amount of signal leakage between frequency bins of the FFT. “Weak” windows, such as Rectangular, allow a lot of leakage, which may blur your Spectrogram vertically. “Strong” windows, such as Kaiser or cos3, eliminate leakage at the expense of a slight loss of frequency resolution.

Frequency Scale

Frequency scale adjustments can help you see useful information more easily. Different scales have different characteristics for displaying the vertical (frequency) information in the Spectrogram display.

- **LINEAR**: Displays frequencies spread out in a uniform way. This is most useful when you want to analyze higher frequencies.
- **LOGARITHMIC**: This scale puts more attention on lower frequencies.
- **MEL**: the Mel scale (derived from the word Melody) is a frequency scale based on how humans perceive sound. This selection is one of the more intuitive choices because it corresponds to how we hear differences in pitch.
- **BARK**: The Bark scale is also based on how we perceive sound, and corresponds to a series of critical bands.

Frequency Overlap

Controls the amount of oversampling on the frequency scale of Spectrogram. When used together with the Reassignment option, it will increase the resolution of the Spectrogram vertically (by frequency).

Time Overlap

This controls the time oversampling of the Spectrogram. In most cases, overlap of 4x or 8x is a good setting to start with. However, using higher overlap together with the Reassignment option will increase the time resolution of a Spectrogram, letting you see transient events clearly.

Color Map

The Spectrogram display allows you to choose between several different color schemes. There is no right or wrong color setting to use and we recommend you try them all to determine your preference. Sometimes certain color modes will make different types of noise stand out more clearly. Experiment!

High-Quality Rendering

Accurate max-bilinear interpolation of the Spectrogram (recommended). Turning this control off makes Spectrogram rendering slightly faster, but you'll lose some detail and clarity in the Spectrogram image.

Reduce Quality Above

RX's Spectrogram uses very accurate rendering, letting you see audio problems, such as clicks, even at low zoom levels. However, performing such rendering for long files can be somewhat slow. When the length of the visible Spectrogram is above the specified number of seconds, the Spectrogram calculation is changed to a fast and less accurate preview mode. When you zoom in, the Spectrogram calculation becomes accurate again.

Cache Size (MB)

Limits the amount of memory used by the Spectrogram.

Rulers

On the right side of the Spectrogram/Waveform display are the Amplitude ruler for the Waveform, Frequency ruler for the Spectrogram, and Color Map ruler for the Spectrogram.

Amplitude Rulers

You can right-click on the spectral Amplitude ruler to reveal a selection of amplitude scales:

- **dB:** Shows Waveform levels in decibels, relative to digital full scale (it is the most common type of scale used for spectrum analyzers).
- **NORMALIZED:** Shows Waveform levels relative to the full scale level of 1.
- **16 BIT:** Shows Waveform levels as quantization steps of a 16-bit audio format (−32768 to +32767).
- **PERCENT:** Shows Waveform levels as percentage from full scale.

Color Map Ruler

This ruler shows what color represents what amplitude in the Spectrogram. The range of this display is the dynamic range of the RX Spectrogram. You can click and drag the map to change the range and use the scroll wheel to make the range larger or smaller. This is useful for seeing very quiet noises without using gain to change the level of your audio.

Frequency Rulers

Right-clicking on the frequency ruler will display the frequency scale options:

- **LINEAR:** Linear scale means that Hertz are linearly spaced on a screen.

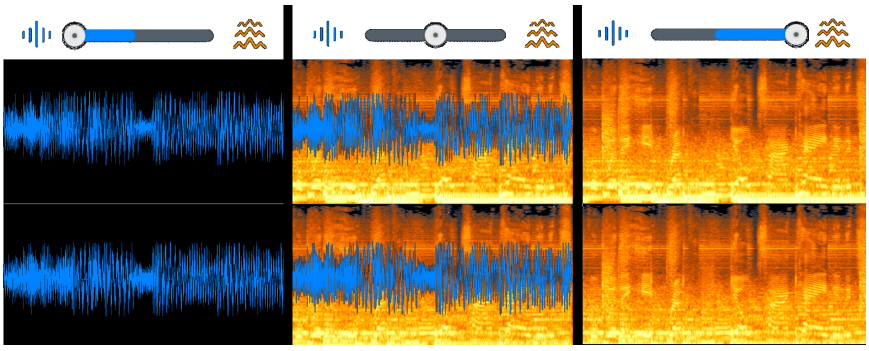
- **MEL *(default)* & BARK:** Mel and Bark are frequency scales commonly found in psychoacoustics, and reflect how our ears detect pitch. They are approximately linear below 500 Hz and approximately logarithmic above 500 Hz.
  - **MEL** scale reflects our perception of pitch: equal subjective pitch increments produce equal increments in screen coordinates.
  - **BARK** scale reflects our subjective loudness perception and energy integration. It is similar to Mel scale, but puts more emphasis on low frequencies.
- **LOG:** in this mode, different octaves occupy equal screen space. The screen coordinates are proportional to the logarithm of Hertz down to 100 Hz.
- **EXTENDED LOG:** this extends the logarithmic scale down to 10 Hz, so that it puts even more attention on lower frequencies.
- **PIANO ROLL OVERLAY:** A representation of how specific frequency ranges correlate to the western musical scale can be displayed by right-clicking on the Frequency ruler and selecting Show Piano Roll. If you would like to hide the frequency indicators so they don't obscure this piano roll, you can disable Show Frequencies and Ticks (which is enabled by default).

Waveform Displays

The Spectrogram Display features the Waveform Transparency Balance slider that lets you superimpose a Waveform display over the Spectrogram:



The slider lets you see both frequency and overall amplitude at the same time. This can be invaluable for quickly identifying clipping, clicks and pops, and other events.



Waveform Overview

An overview of the entire audio file's Waveform is displayed above the main Spectrogram/ Waveform display in order to provide a handy reference point when zooming and making audio selections in RX.



The Waveform overview will always display the entire audio file, and will also display any selections made in the main display. When zooming in on your audio, the currently visible audio region will also be highlighted in the Waveform overview. Click and drag on the highlighted region in order to scroll your main audio display left or right, and click and drag on the edges of the highlighted region in order to make the zoom tighter or wider. To zoom out fully, simply double click on the highlighted visible region.

NOTE

Hover over the waveform overview and use your mousewheel to scale the amplitude of the waveform display to provide a clearer overview. This will not affect the amplitude scaling in the main Spectrogram/Waveform display.

7. Text Navigation

Overview

Text Navigation converts speech into a transcription that is displayed above the spectrogram and shown in sync with the corresponding audio. Transcribed text is searchable and provides reference points for what is contained in the file. This eliminates the need to audition files in order to manually place markers.

TEXT NAVIGATION NOTE

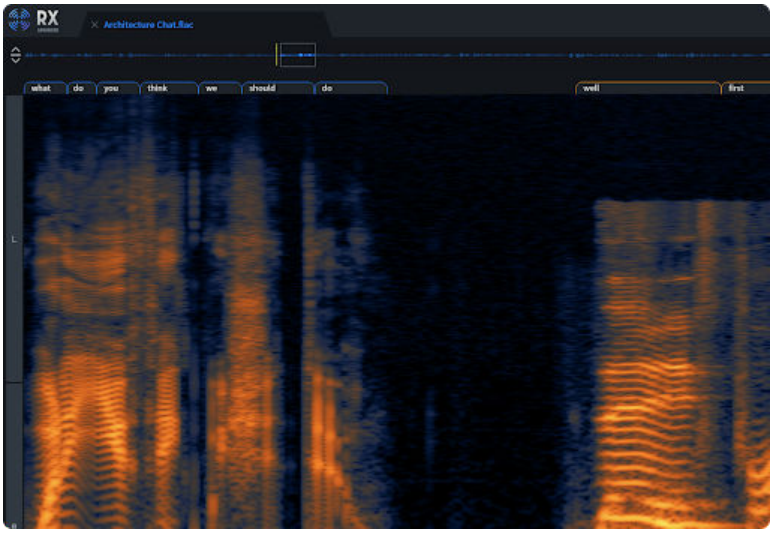
Text Navigation is designed to be an editing navigation tool and not a transcription service. It is optimized for American English. Accuracy may vary if there is noticeable background noise or a speaker has a non-American accent.

Workflow

To get started, drag or import audio into RX and click the Speech Recognition Word Lane button at the bottom left of the spectrogram.



Audio transcription begins immediately and runs in the background at approximately 7 - 9X real time. The transcription populates the tabs in the Word Lane above the spectrogram, indicating the transcription is in process.

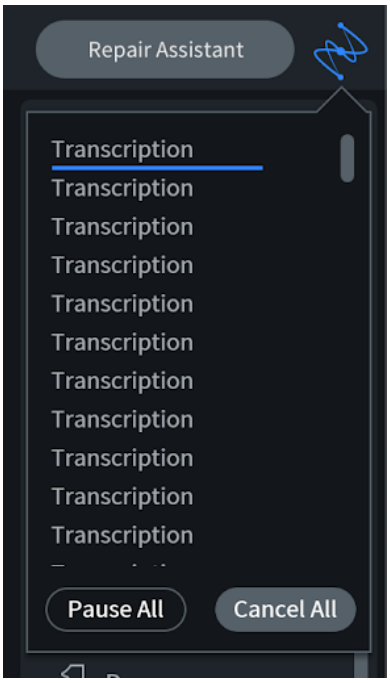


The audio material has to fulfill the following requirements:

- Audio must be dialogue or speech – the transcription of musical lyrics is not currently supported.
- Audio files must be at least 10 seconds long. The speech recognition button is disabled for files shorter than 10 seconds.

- Only American English is supported at this time

Clicking the iZotope logo opens a small window that shows the progress of the transcription. At the bottom of this window are buttons for pausing or canceling the transcription.



Once the transcription is finished, you can zoom in or out of the audio and the transcription expands or contracts accordingly.

- **Mac:** Command+= to zoom in, Command+- to zoom out, or swipe up or down with two fingers on a trackpad
- **Windows:** Ctrl+= to zoom in, Ctrl+- to zoom out

Click on a word tab to select the corresponding audio in the spectrogram. Drag the handles on either side of a tab to select surrounding words, a phrase, or a sentence.

When zoomed in close enough so there is a single word per tab, double-clicking on the word selects it and makes it editable. Type in a correction for a misspelled word or alter individual words to fit your editing needs.

- When editing, typing multiple words in the tab does not split the tab up
- Typing nothing in the tab will not delete the tab
- Right-click the Word Lane to Rescan. (Rescanning will overwrite any corrections)

Search

Text Navigation includes a fuzzy search to find words and variants of words, such as misspellings in the transcript. You can also search for individual letters.

Search is ideal for finding a replacement word, navigating to a specific section, or locating an alternate take.

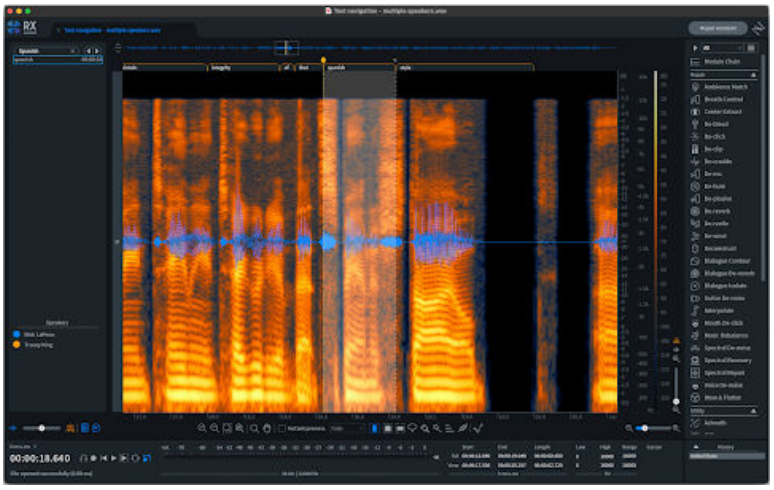
Searching three characters or less works like auto-complete, where the search looks for words that begin with those letters. Searching anything else works like a fuzzy search and will try to return results that are similar to the search query.

Click on the Text Navigation Pane button to bring up the search panel.



NOTE

The Word Lane button will also need to be enabled in order to see the transcription above the spectrogram.



Type a word in the search box and hit return/enter. If the word is found, every instance of it will be listed in the order in which it appears in the audio. Variants of the word are also listed. If the word isn't found, variants are listed only if they are identified in the audio.

Clicking on a word in the list moves the playhead to that word in the transcription and highlights the corresponding audio.

If you have a word that needs replacing, search for instances of that word, pick the best one, then copy it and paste it over the original.

Search can be used to target processing to a selected word. Dragging the handles on either side of the word targets processing to the selected audio.

NOTE

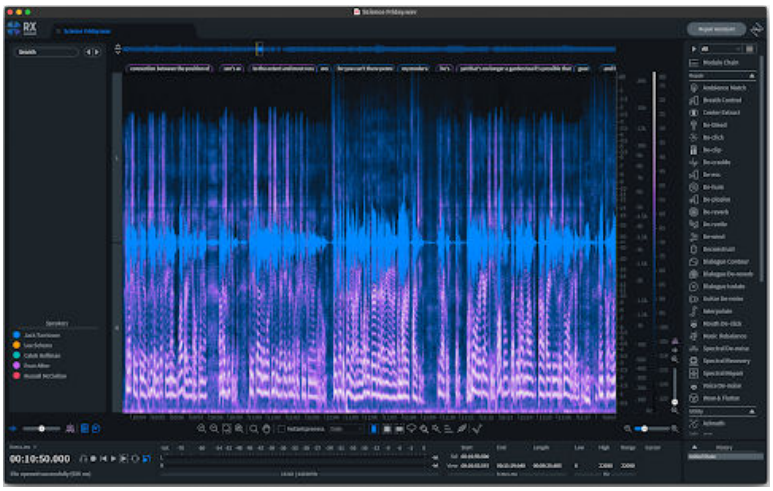
Search works with edited words, but typing more than a single word or phrase in a tab may negatively impact search results.

Multiple Speaker Detection

Built into Text Navigation is functionality that automatically detects when there is more than one speaker on a track and color codes the sections of speech associated with each speaker.

Multiple speaker detection runs after the text transcription pass has finished. Up to 8 speakers can be detected.

Each speaker receives a unique identifying color, which is indicated in the speaker pane and in the corresponding tabs of the transcription.

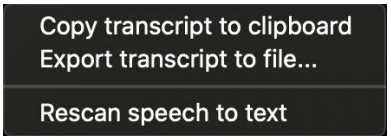


To select all instances of a speaker, click on the speaker name in the speaker pane. This makes it easy to target specific processing by speaker.

Double-click a speaker's name to edit or change it to fit the particular needs of a project.

Export Transcript

To access the Transcript Export menu, click the menu button at the top of the Text Navigation Pane or right-click in the Word Lane.



- **Copy transcript to clipboard:** Copy the transcribed text and paste it into a word processing application.
- **Export transcript to file:** Export the transcribed text as a .txt file.
- **Rescan speech to text:** Transcribe your file again.

8. Interactive Tools

Overview

Below the Spectrogram display is a toolbar that includes several options for working with the spectrogram/waveform display. The toolbar is split into three main categories, Navigation, Instant Process, and Selection.



Zoom Tools

You can zoom in horizontally and vertically on both the waveform and spectrogram views.

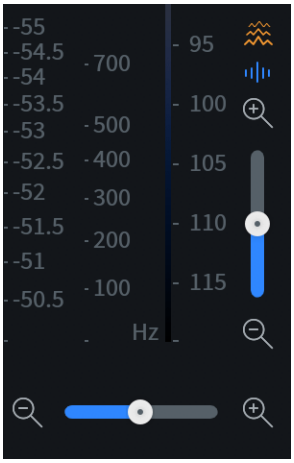
Zoom Selection Tools

The following zoom selection tools are available in the RX Audio Editor:



- **Zoom in:** zooms in on the time ruler.
- **Zoom out:** zooms out on the time ruler.
- **Zoom to selection:** zooms to fill the spectrogram/waveform display with the current selection.
- **Zoom to whole file:** resets zoom time and frequency zoom levels to default.
- **Zoom tool:** When enabled, the zoom level will update to fit the selection to the window when using the Time, Time-Frequency, or Frequency Selection Tools.

Zoom Sliders



Magnitude Ruler Zoom Slider

The vertical slider in the lower right-hand corner of the spectrogram/waveform display controls the zoom level for the spectrogram or waveform amplitude ruler(s). The buttons above the vertical slider select which ruler it will affect when adjusted.

Time Ruler Zoom Slider

The horizontal slider in the lower right-hand corner of the spectrogram/waveform display controls the zoom level of the time ruler.

Navigation Tools

Grab & Drag Tool



When zoomed in on an area, the Grab & Drag Tool [G] can be used to move through the time range by clicking and dragging on the Spectrogram.

Dragging Rulers and Using the Mousewheel

The rulers to the right and below the Spectrogram display can clicked and dragged to reposition the spectrogram and waveform to show a different time range, amplitude range or frequency range.

In addition the range shown can be adjusted with your mouse:

- ▶ To adjust the zoom, place your cursor on top of a ruler and move the mouse wheel.
- ▶ To reset any of the rulers to the default display range, double click on the ruler.

> **NOTE**
> All active selections will be accurately preserved and scaled when zooming the main display.

Channel Selectors

Channel Selectors appear on the left hand side of the spectrogram/waveform display when a file is loaded. Channels can be enabled or disabled individually by clicking on the channel selector buttons. Deselecting a channel will exclude it from playback, selections, and processing.

- If all channels are enabled, single-clicking on a channel selector will disable all other channels, only the channel that was clicked on will remain enabled.
- When some of the channels are disabled, single-clicking on a disabled channel selector will enable that channel without affecting the enabled disabled state of other channels.
- When any channel number of channels are disabled, double-clicking on any channel selector will quickly enable all channels.

KEYBOARD SHORTCUTS: STEREO FILE CHANNEL SELECTORS

Channel Order (Multichannel)

When working with multichannel files, the channel selector label order can be configured by right-clicking on the time ruler or by clicking on the arrow to the right of the time format display. The channel order options will update based on the number of channels in the active file tab. The "Discrete" option will label all channels as "M" (mono).

Instant Process

Instant Process [I] is a selection tool modifier. Instant Process is only available in RX Standard and Advanced.



When Instant Process is enabled, any new selection made will be immediately processed with the selected module in the Instant PRocess menu. The module settings applied by instant process reflects the current settings in the selected module.

When Instant Process is disabled, processing, editing and selection tools will function as they normally do.

NOTE

If you hold Shift while using Instant Process, this will allow you to build up additional selections. Once you release Shift, processing will occur. This is especially useful for tools such as the Magic Wand, which will pick up additional harmonics upon second click. If you don't like that selection, also hold Alt, and start redrawing a new selection. Release Shift once you're ready to process.

Instant Process offers several different modes, accessible via a drop down menu, which will instantly process the settings present in the named module/tab. The default settings are used, but if you define custom settings in that module, Instant Process will recognize and apply those custom settings. The modes are:

Instant Process: Attenuate

This mode will instantly apply the active settings from the Spectral Repair module's Attenuate tab. This is particularly useful if you see anything in the spectrogram you don't wish to remove entirely, but would rather quickly blend into the surrounding audio to make it less obvious or intrusive.

Instant Process: De-click

Applies the active settings from the De-click or Interpolate modules. De-click Instant Process will automatically remove all clicks present in your selection, which is particularly useful for editing a dialogue file, mismatching sample rate clicks and pops, and vinyl clicks.

If you make a selection under 4000 samples in length, this mode will automatically use the Interpolate module. Selection longer than 4000 samples will use the settings from the De-click module. The De-click module is effective on selections above 4000 samples in size, as it is able to identify clicks in relation to desirable audio, and then intelligently separate and remove the clicks. If a selection is less than 4000 samples in length, it is likely a small selection of an individual click, and Interpolate will fill the selection with audio information based on the surrounding audio.

Instant Process: Fade

This mode will instantly apply the active settings from the Fade module. This is particularly useful if you'd like to smooth over a transition or edit point within a complex audio file, especially if it's a limited bandwidth selection, such as choosing to fade in a certain harmonic or audio event in an audio file without changing the volume of the rest of the audio.

Instant Process: Gain

This mode will instantly apply the active settings from the Gain module. If you want to quickly adjust certain audio events up or down in volume, you can simply paint over it to see the immediate gain adjustment. For overall volume adjustment, use the Clip Gain line [Cmd+G / Ctrl+G].

Instant Process: Replace

This mode will instantly apply the active settings from the Replace tab in the Spectral Repair module. This is particularly useful if you see anything in the spectrogram you wish to remove entirely, as it will use the audio information that surrounds your selection to instantly and intelligently fill the gap.

Selection tools

The following selection tools are available in the RX Audio Editor:

Icon	Name	Description
	**Time selection tool [T]**	Select a range of time within the file (horizontally within the spectrogram)
	**Time-Frequency Selection tool [R]**	Makes rectangular selections in the spectrogram display to isolate sounds by time and frequency
	**Frequency Selection tool [F]**	Makes frequency only selection (vertical in spectrogram)
	**Lasso Selection tool [L]**	Makes a selection based on a free-form outline drawn with your cursor on the spectrogram
	**Brush Selection tool [B]**	Draw a free-form selection using a defined brush size in time and frequency in RX's spectrogram. The size of the Brush Selection tool can be adjusted by clicking and holding on the Brush Tool icon. **Note**: With the brush tool selected, you can also hold Control/Command and move the mouse wheel to make the brush size larger or smaller.

Icon	Name	Description
	**Magic Wand Selection tool [W]**	Automatically selects similar harmonic content surrounding the selected material. Click on the spectrogram to select the most prominent tone under the cursor when the magic wand is selected. Clicking on an existing selection with the Magic Wand tool will automatically select the overtone harmonics or related audio components of your current audio selection. **Note**: You can use the Brush or Lasso tools first to broadly define a sound and then use Magic Wand to refine your selection to include relevant harmonic material.
	**Harmonic Selection tool [Shift+Cmd+H / Shift+Ctrl+H]**	Duplicates your current selection to include harmonics above it. Start by selecting the fundamental frequency of audio with harmonics, then add or remove harmonic selections with this tool before processing.
	**Selection Feathering tool [Shift+F]**	Allows for crossfading of processed and unprocessed audio on both time and frequency axis. Adjust these controls depending on how precise you want the edits to be [Alt/Option+F]. Higher feathering values will introduce longer crossfades, while lower values will create sharper edits. Use the feather icon to enable/disable this feature.

When Feathering is disabled, the default time feathering of 20 ms remains on the selection to prevent the introduction of unwanted sonic artifacts.

When feathering is enabled, there is potential for the selection visualization to lag slightly.

View Clip Gain

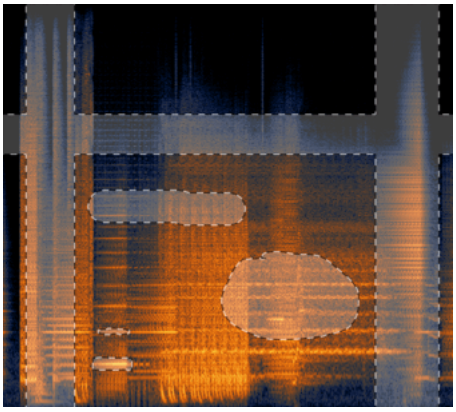
Quickly enable/disable the clip gain curve overlay.



Selection Modifiers

[Shift] - Add to selection

Hold down shift after making a selection in order to add another, separate selection. If any part of the new selection overlaps any other, the selections will be grouped into one. This can be especially powerful when combining several selection tools to create multiple selections of different size and shape.

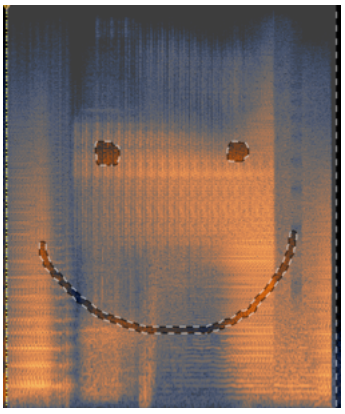


Multiple selections made using the Shift modifier key with different selection tools

[Alt/Option] - Subtract from Selection

Holding down Alt/Option will allow you use the currently chosen selection tool to remove or erase any portion of an existing audio selection. This can be especially useful with the Lasso or Brush tools, allowing you to edit or refine any piece of an existing selection.

This is also useful for refining complicated free-form selections. First make your lasso, brush, or magic wand selection, and then hold alt while using the time, frequency, or time and frequency tools to exclude entire time and frequency ranges from processing. Holding Alt/Option effectively turns the Brush tool into a selection eraser for broad refinements and Lasso into a selection "X-Acto knife" for detailed selection revision.



Selections refined to perfection using the Alt key modifier

LINKING SELECTION STATE TO UNDO HISTORY EVENTS

Using Ctrl/Cmd-Z to undo any particular process will also bring back the previous audio selection exactly as it was before applying any processing. In order to make use of this feature, be sure that Store Selections with Undo History is enabled inside of RX's Preferences > Misc menu.

[Ctrl/Cmd] - Move Playhead without affecting Selection state

Hold down Ctrl/Cmd to move the transport's playhead to any position without erasing your current audio selections. This can be especially useful with previewing or comparing complex audio selections without having to remake these specific audio selections.

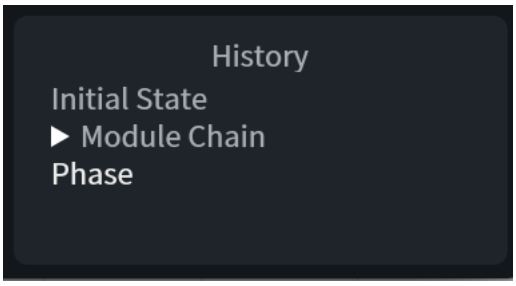
[Mouse Over] - Grab and Drag Selection

After using any of RX's tools to select a portion of your audio, when the mouse is subsequently placed on top of any selection, a Grab and Drag hand cursor will be displayed automatically, allowing you to change the position of that selection.

9. Undo History

Overview

In addition to the Undo and Redo commands, the Undo History list allows you to see a timeline of changes you've made and non-destructively revert back to earlier states. RX keeps a log of all your edits in this Undo History. When modifying Clip Gain or processing with the Module Chain related edits will be added under parent items, named "Clip Gain" and "Module Chain" respectively, in the Undo History list.



NOTE

You can rename items in the undo history by double-clicking them. You can save an .rxdoc of your current file to retain the undo history list, you can find more information on RX documents in the Working with Files chapter.

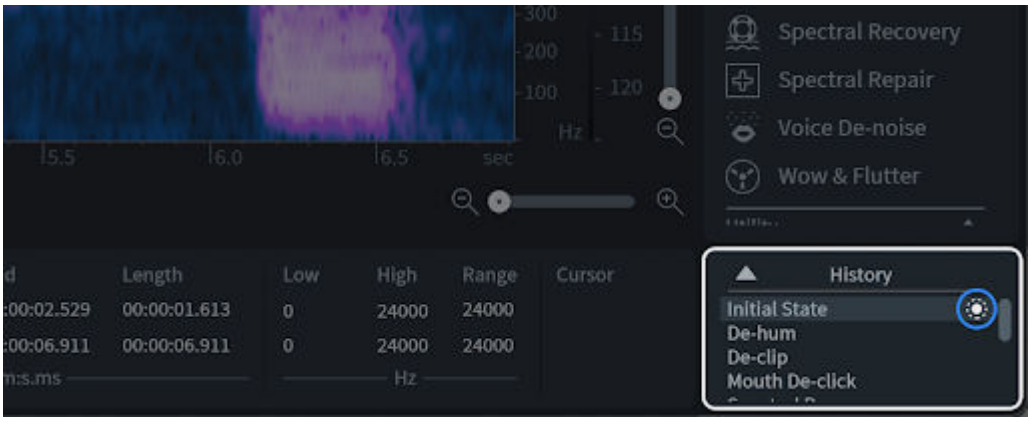
Export History

You can use the Export History feature to save an XML file, listing the entire undo history for your particular file. For forensic and archival purposes, it is often useful to have an official record of all edits that were made to a particular file. When a file's history is exported, the following information will be stored in an XML file:

- RX Version Number
- Time and Date
- Corrected File
- Number of Channels
- Sampling Rate/Bit Depth
- Edit History: Parameters and Selections

Restore Selection

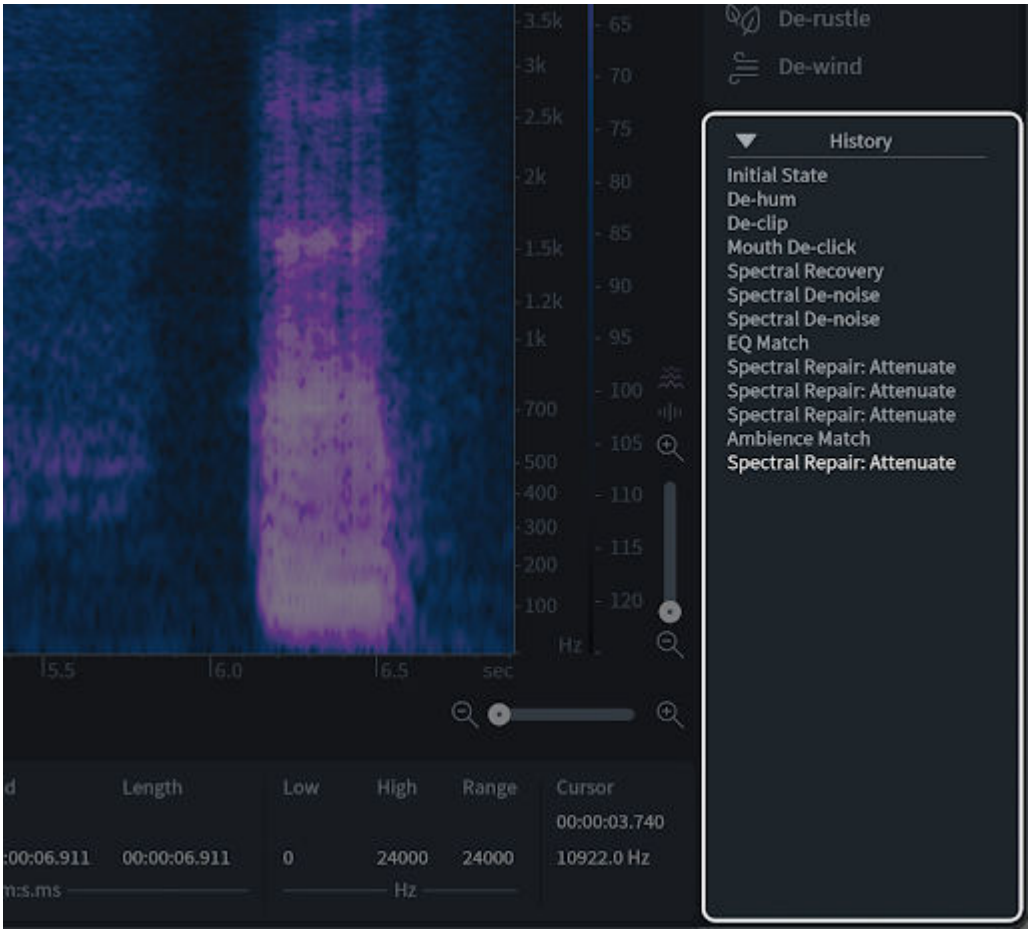
You can now restore selected regions to earlier states in the Undo History. When hovering your cursor over an item in the History Pane with a selection made in the spectrogram, an icon appears in line with the text. Clicking this icon will restore your selection to that point in history.



This is the same as going back to that point in undo history, except only for your current selection.

Expandable History List

The History list is now expandable, so you can see more of your edit history without scrolling.



10. Application Menus

File

The file menu provides options for creating, importing and exporting files. For information about file management in the RX Audio Editor, see the Working with Files chapter.

Edit

The Edit menu includes the following options:

- **Undo [Ctrl/Cmd-Z]** Reverses the last action taken.
- **Redo [Ctrl/Cmd-Shift-Z; Ctrl/Cmd-Y]** Cancels the undo.

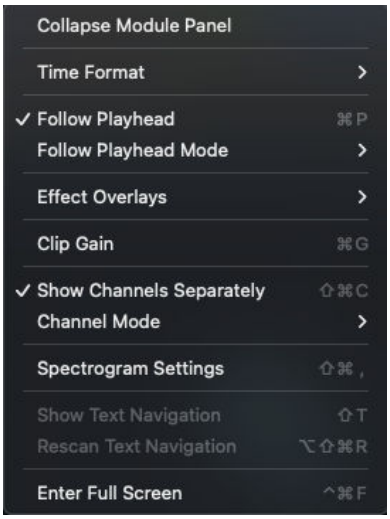
> **UNDO HISTORY**
>
> The RX Audio Editor includes an Undo History event list in each file tab. The Undo History list stores a list of all processing or edit operations applied to a given file.

- **Cut[Ctrl/Cmd-X]** Removes the currently selected audio and stores it temporarily on the Clipboard.
- **Copy [Ctrl/Cmd-C]** Makes a copy of the currently selected audio and places it on the Clipboard.
- **Paste[Ctrl/Cmd-V]** Places audio that has been copied or cut to the Clipboard at the current cursor point.
- **Paste Special** Provides additional options for placing the Clipboard data:
  - **Insert** [Ctrl/Cmd-Alt/Opt-V]: Inserts the audio from the Clipboard and moves audio in the project [does not overwrite]
  - **Replace** [Ctrl/Cmd-Alt/Opt-Shift-V]: Replaces audio in the project with audio from the Clipboard
  - **Mix [Shift-V]** Combines the audio from the Clipboard with audio in the project
  - **Invert and Mix [Alt/Opt-V]** Inverts the audio in the Clipboard and then mixes it with audio in the project. This is useful when you want to compute the difference between two signals.
  - **To Selection [Alt/Opt-Shift-V]** Pastes audio from the clipboard only within the selection bounds, regardless of the copied audio's length. If the copied audio is longer than the new selection, the audio will be cropped to fit. If the selection is longer than the audio being pasted, silence will be inserted to fill the remaining space.
  - **To Clip Gain Only [Ctrl/Cmd-Shift-V]** Pastes only clip gain information to the current selection
- **Deselect [Ctrl/Cmd-D]** If audio is selected, deselects it and places the anchor sample at the start of the selection.
- **Reselect [Esc]** Restores the last selection if you have no current selection.
- **Select All [Ctrl/Cmd-Shift-D]** Selects the entire open file.
- **Invert Selection [Ctrl/Cmd-Shift-I]** Selects everything that isn't currently selected.
- **Invert Selection Frequencies [Ctrl/Cmd-I]** Selects everything in the current time range that isn't selected. This is useful for refining processing by first selecting what you don't want to process, then inverting the selection frequency.

- **Select Harmonics [Ctrl/Cmd-Shift-H]** Refines the current selection to include more harmonics. For this feature to work well, try it with a simple selection that includes only the fundamental harmonic of what you are trying to select. You can also use the Magic Wand tool to automatically refine a selection to include the appropriate harmonics.
- **Begin Selection at Playhead [Left Bracket]** *during playback only* If audio is currently selected, this will automatically adjust the selection to begin at the current playback position.
- **End Selection at Playhead [Right Bracket]** *during playback only* Automatically create a selection between the current playhead position and the original anchor playhead position.
- **Delete Selection [Delete on a time selection]** Deletes the selected audio and closes the space with audio from either side of the timeline.
- **Silence [Delete on a frequency, time-frequency, or freeform selection]** Deletes selected audio and replaces it with silence.
- **Trim to Selection [Ctrl/Cmd-T]** Deletes all audio except for the selected audio.
- **Add Marker or Region [M]** This will create a new marker point at the current location of the cursor/playhead or create a new region if any audio is selected.
- **Edit Cursor Mode** Changes the behavior of the editor cursor to select by time and/or frequency, or to zoom. These modes can also be selected from the Cursor Mode buttons.
  - **Select Time [T]** Makes a time selection
  - **Select Time/Freq [R]** Makes a Rectangular time-frequency selection
  - **Select Freq [F]** Makes a frequency selection for the duration of a file
  - **Lasso [L]** Selects everything in a freely defined area
  - **Selection Brush [B]** Selects everything in a predefined radius
  - **Selection Wand [W]** Intelligently selects material similar to whatever is under your cursor [magic wand]
  - **Zoom Time [Z]** Zooms in time
  - **Zoom Time/Freq [Shift-Z]** Zooms in time and frequency
  - **Zoom Freq [Alt/Opt-Z]** Zooms in frequency
  - **Grab Time [G]** Grabs and drags the view in time
  - **Grab Time/Freq [Shift-G]** Grabs and drags the view in time and frequency
  - **Grab Freq [Alt/Opt-G]** Grabs and drags the view in frequency
- **Snap [Ctrl/Cmd-Shift-;]** Snap selections to the boundaries selected in the “Snap To” sub-menu.
- **Snap to:**
  - Markers
  - Ruler Coarse
  - Ruler Fine
  - Zero Crossings
  - All
  - None

View

The View menu includes the following options:



- **Collapse/Expand Module Panel** Collapses the module list panel of RX into a row of icons.
- **Time Format** RX's time scale and playhead location counter can be set to show different time units. Learn more about changing the time format in the Transport chapter
- **Follow Playhead [Ctrl/Cmd+P]:** Toggles whether or not the current view follows the playhead position during playback.
  - In **Page mode**, the view will follow the playhead one view length at a time.
  - In **Continuous mode**, the view is centered on the playhead as it moves across the file.
- **Effect Overlays:** This sub-menu allows you to turn special display features for the De-clip and Spectral Repair modules on and off. For an overlay to be visible, you need to have the option selected in the view menu, and the corresponding module window needs to be open.
  - **De-clip Threshold** When the waveform is visible, the threshold settings and controls appear as white lines within the display. This display can be used to adjust the de-clip threshold settings.
  - **Spectral Repair Source Regions** Displays effective region bounds when using the Spectral Repair module.
- **Clip Gain:** Enables/disables the Clip Gain envelope in the main editor window. The Clip Gain curve can also be toggled on or off by clicking the "View Clip Gain" button to the right of the selection tool buttons, or by using the following keyboard shortcuts: Command+G (Mac); Ctrl+G (Windows).
- **Show Channels Separately:** Toggles how channels are displayed in the spectrogram/waveform view. When enabled, each channel will be drawn in its own lane in the main editor view. When disabled, the channels are drawn in one summed view. This allows for greater vertical resolution, especially when working with multichannel files. This option can also be toggled on or off by: clicking the Channel View button to the left of the mini waveform overview display or by using the following keyboard shortcuts: Command+Shift+C (Mac); Ctrl+Shift+C (Windows).
- **Channel Mode:** Selects between Left/Right and Mid/Side channel mode. This affects the waveform and spectrogram display as well as any processing that may be applied. This option is only available if Show Channels Separately is enabled.

> **NOTE**
> Music Rebalance and Dialogue Isolate are not designed for use in Mid/Side mode.

- **Spectrogram Settings:** Opens the Spectrogram Settings window.

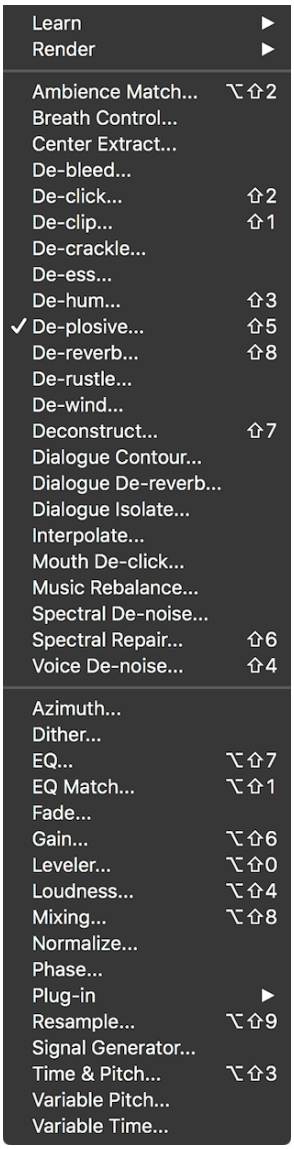
- **Enter full screen:** Enables full screen mode.

Modules

The Modules menu includes options for opening module windows, running Learn on the current selection without opening the module window, and rendering module settings on the current selection without opening the module window.

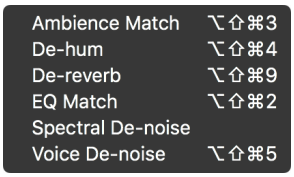
Open Module Window

Open a module window by selecting it from this menu.



Learn

Run a Learn pass on the current selection without opening the associated module window.



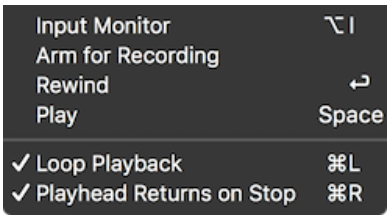
Render

Render settings for any module on the current selection without opening the associated module window.



Transport

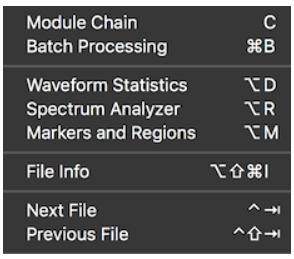
The Transport menu includes the following options:



- **Input Monitor [Alt/Option-I]** Enables input monitoring. When input monitoring is enabled, the input signal of RX will be routed to the output signal of RX.
- **Arm for Recording / Record / Stop Recording [Alt/Option-Space]** Runs the next possible step for recording. If you have not opened a new file, Arm for Recording will open the New File dialogue box for you.
- **Rewind [Return]** Sets the playhead to the beginning of the file.
- **Play/Stop [Space]** Starts or stops playback. If Input Monitoring is enabled, starting playback will temporarily suspend Input Monitoring.
- **Loop Playback [Control/Command-L]** Toggles playback looping. If nothing in the file is selected, the end of the file will loop back to the beginning.
- **Playhead Follows Playback [Control/Command-R]:** Toggles the behavior of the playhead on stop. If this is enabled, the playhead will return to the anchor sample (the position before playback began).
  - This is useful for comparing processing. If this is disabled, the anchor sample will be set to the current playhead position. This is useful for moving through a file while listening for irregularities.

Window

The Window menu includes the following options:



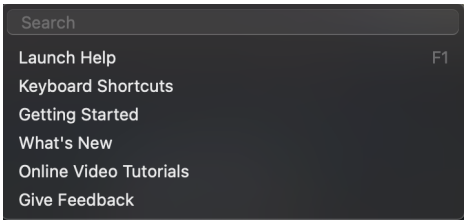
- **Batch Processing [Ctrl/Cmd+B]:** This gives you access to file based batch processing, as explained in the Batch Processor chapter.
- **Waveform Statistics [Alt/Opt+D]:** This gives you access to informational readouts on a variety of amplitude measurements, as explained in the Waveform Statistics chapter.
- **Spectrum Analyzer:** The Spectrum Analyzer displays an analytical view of your audio. More information is located in the Spectrum Analyzer chapter.
- **Markers and Regions [Alt+M]:** Markers and regions allows you to define and save particular points or selections in time for your audio file. More information is located in the Markers and Regions
- **Close All Floating Windows/Reopen Closed Windows [Ctrl/Cmd+Opt+W]:** Closes or reopens all floating windows.
- **File Info:** View metadata and other information about the audio file, as explained in the Working with Files chapter.

- **Next File (Control-Tab)** Changes RX's current file tab to the next file in the window order.
- **Previous File (Control-Shift-Tab)** Changes RX's current file tab to the previous file in the window order.

Help

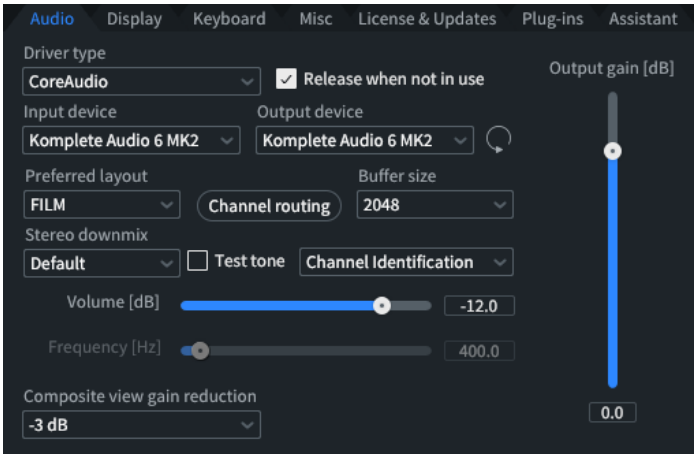
The Help menu includes the following options:

- **Launch Help**: Open the product help documentation 8 **Keyboard Shortcuts**: Open the Keyboard Shortcut Guide
- **Getting Started**: Launch the First Time User Tour
- **What's New** See the new features and tools RX 11
- **Online Video Tutorials**: View articles and tutorials on iZotope.com
- **Give Feedback**: Send feedback directly to the iZotope team



11. Preferences

Audio



- **Driver Type:** Selects the audio device driver type (for example: ASIO, CoreAudio, or RX Monitor)

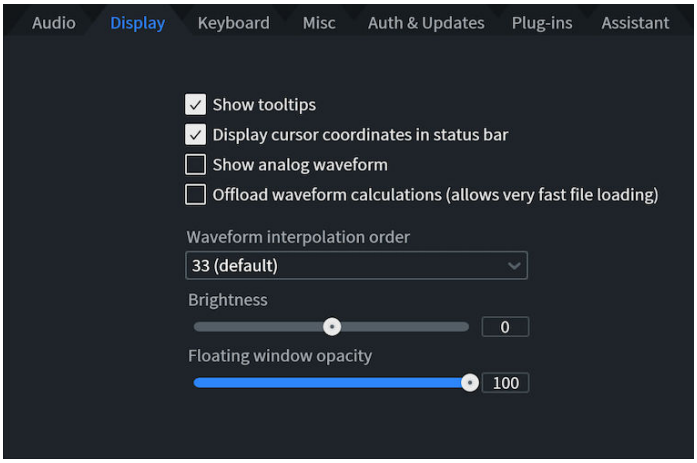
> Some hardware devices monopolize the audio drivers when sending audio clips to RX via RX Connect. If you are not able to hear the audio sent to RX from your DAW with RX Connect, instantiate RX Monitor on a track in your session and change the audio driver to RX Monitor in the Driver type menu.

- **Input/Output Device:** Choose the device/sound card you want RX to use for playback and recording.
- **Device Refresh:** Rescan for newly connected audio devices.
- **Buffer Size:** The total playback buffer size. In general, lowering these buffer sizes will improve meter responsiveness and lower latency, but increase CPU needs. Raising buffer sizes will lower CPU cost but increase latency. It's worth exploring these ranges to find values that work best on your system.
- **Num Buffers:** Number of playback sub-buffers. (Windows MME Only.)
- **Composite View gain reduction:** Nondestructively reduces the output gain of all clips included in the Composite View tab by the amount specified in the dropdown.
- **Channel Routing:** Opens the Channel Routing window. Input and output channel routing can be configured in this dialog when working with ASIO/CoreAudio driver types or to configure multichannel output routing.
- **Preferred Layout (Multichannel only):** Selects the default channel ordering option when multichannel files are loaded in the application. **ADV**
- **Stereo Downmix (Multichannel only):** Selects the stereo downmix configuration to use when **monitoring** multichannel files with a stereo audio output device. **ADV**

> RX does not support downmixing or upmixing files when saving or exporting. The downmix option only applies to playback of multichannel files on stereo systems.

- **Configure Driver:** Launches the manufacturer's driver configuration dialog.
- **Release when not in use:** Auto-closes the audio device when playback in RX stops, freeing it for use in other audio applications. Disable this if playback from RX isn't responsive enough.
- **Test Tone:** The test tone generator is useful for testing your speakers, audio hardware and listening environment. Tones at set frequencies or at a custom frequency can be used as test tones, as can white or pink noise. In addition, a Channel Identification mode will identify left and right speakers.
  - **Enable:** Starts playback of a test tone.
  - **Type:** Sets the type of test tone to play.
  - **Volume:** Sets the volume of the test tone.
  - **Frequency:** Sets the frequency of the test tone.
- **Output Gain:** Output gain allows you to nondestructively adjust the playback level of RX 11 Audio Editor.

Display



- **Show tooltips:** When enabled, hovering over an RX feature with the mouse cursor will show a short description of the feature.
- **Display cursor coordinates in status bar:** When enabled, the time coordinate of the cursor is shown in the status bar at the bottom of the RX main window. The amplitude of the audio at the cursor position and the frequency at the cursor position is also shown.
- **Show analog waveform:** When digital audio is played back, it is converted to analog. The peak values in the analog waveform can be larger than the peaks in the digital waveform, leading to clipping in the output of a digital-to-analog converter. When Show analog waveform is enabled, RX will compute an analog waveform in the background. Any peaks will be highlighted in red on top of the existing digital waveform.

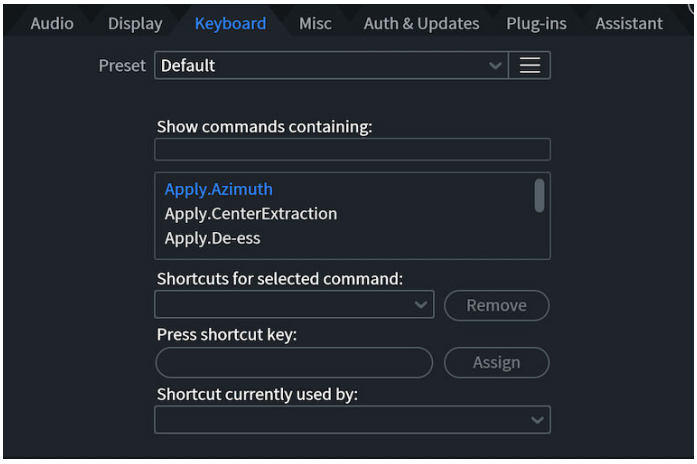
> **NOTE**
>
> RX will automatically display an analog waveform when zooming in at extreme zoom levels.

- **Offload waveform calculations:** When enabled, RX's waveform display will be computed in the background. This allows very large files to be loaded very quickly, but it slows down RX's waveform displays.

- **Waveform interpolation order:** If you zoom into the waveform so that individual samples become visible, RX will display an upsampled analog waveform as well as the individual digital samples. The interpolation order controls the quality of upsampling. Higher values yield more accurate analog waveforms at the expense of CPU usage.
- **Brightness:** Adjusts the general brightness of the RX interface, allowing you to make RX more readable on your specific display.
- **Floating window opacity:** Changes the opacity for RX's floating windows. This can be useful if you wish to leave floating windows on top of the spectrogram and waveform without completely obscuring the display.

Keyboard

> **CUSTOMIZING KEYBOARD SHORTCUTS**
>
> While RX includes default keyboard shortcuts, you can also customize them to your liking. Refer to the Keyboard Shortcut Guide to reference a list of default key commands and internal shortcut command names used by RX. Referring to the guide can help you quickly identify the name of the key commands you want to customize and search for them in the "Show commands containing" field (explained below.)

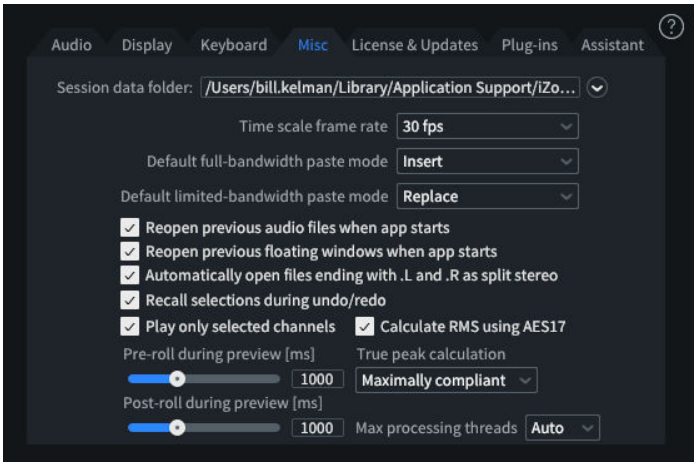


- **Presets:** Save groups of key assignments with this tool.
- **Show commands containing:** Lets you search by keyword for a command you want to assign to a keystroke.
- **Shortcuts for selected command:** Shows if there are any keystrokes assigned to the command selected in the above menu.
- **Remove:** Removes the currently assigned keystroke from a command.
- **Press Shortcut Key:** To assign a new keystroke to a command, select the command from the menu, then click in this field and press a key or combination of keys.
- **Assign:** Assigns the entered keystroke to the current command. The shortcut will only be assigned to the current command if you press this button.

- **Shortcut key currently used by:** Lists commands that the current keystroke is assigned to.

> **USING THE ALT MODIFIER ON WINDOWS**
>
> On Windows systems, by default, "Alt + a letter" will open the corresponding menu for your currently open application. `Alt` + `V` for example will open RX's View menu drop down. By default, none of RX's shortcuts should conflict with these keyboard shortcuts, however if you wish to assign `Alt` + `V` to another operation, it will take precedence over the View menu.

Misc



- **Session data folder:** Allows you to choose a different folder to save RX's temporary session data. These files are created to allow actions to be undone and sessions to be recalled in RX. Because these can be very large, it is best to set this to the drive on your computer with the most free space.
- **Time scale frame rate:** This sets the frame rate used to draw the time scale when RX is set to display the time code (see View menu or right-click the time ruler to change this setting). Choose from a list of standard frame rates or click in the combo box to define a custom frame rate.
- **Default full-bandwidth paste mode:** This controls RX's behavior when pasting a full-bandwidth audio selection. Insert will move aside existing audio, Replace will overwrite existing audio, and Mix will add to existing audio.
- **Default limited-bandwidth paste mode:** Similar to the full-bandwidth paste mode, this controls RX's behavior when pasting a limited-bandwidth audio selected.
- **Reopen previous audio files when app starts:** When enabled, RX will open all of the files (including edits, processing and undo history events) that were present when RX was last closed. Disabling this option will open the RX Audio Editor in its default state (no files loaded.)
- **Reopen previous floating windows when app starts:** When enabled, any floating windows that are open when the application is closed will be reopened the next time the application is opened.
- **Automatically open files ending with .L and .R as split stereo:** Mono audio files with (.L and .R) as well as (.1 and .2) extensions will be opened as stereo files when this option is enabled.

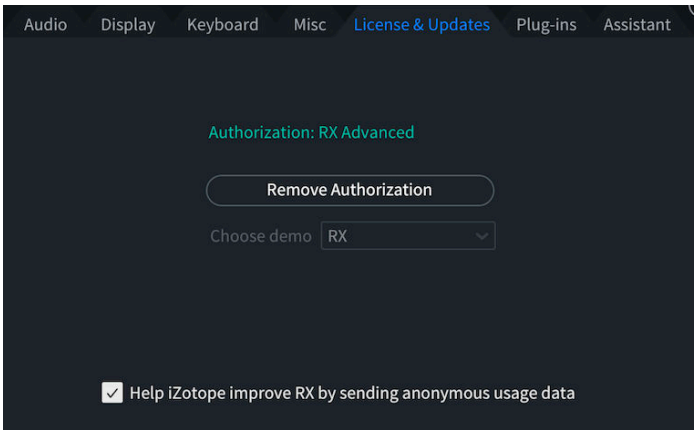
- **Recall selections during undo/redo:** When this is enabled, RX will recall the selection used for an item in the undo history. When stepping through the undo history events, selections that were used for each event will be restored along with the audio.

  DISABLING THE SELECTION UNDO/REDO OPTION

  Sometimes it is useful to turn this off if you need to compare undo history items and not break your current selection (like a useful loop).

- **Play only selected channels:** If only a single channel of audio is selected and this option is enabled, all other channels will be muted during playback.
- **Calculate RMS using AES-17:** Uses the AES-17 1998 standard for RMS calculations (0 dB is a full scale sine wave) in the level meter, Waveform Statistics and Leveler modules. The other option is when 0 dB is the RMS of a full-scale square wave. These options differ by 3 dB.
- **Pre- and Post-Roll during preview (ms):** When Previewing audio processing in any module, the specified time amount will be added to the beginning and end of the previewed selection in order to provide contrast between unprocessed and preview-processed audio.
- **True Peak Calculation:** RX provides several modes for measuring true peak levels in Waveform Statistics:
  - **BS.1770 baseline** mode uses a "minimal" 4× oversampling with the specific filter from the BS.1770 standard.
  - **HQ oversampling** mode uses a higher (at least 9×) oversampling factor with a flat passband response. It does not exhibit the passband ripples of the BS.1770 filter, but may sometimes produce lower readings than the baseline filter.
  - **Maximally compliant** mode runs both of the above measurements and takes the higher of two detected true peak levels. In this way, it achieves compliance with most existing true peak measurement software.
- **Max Processing Threads:** When RX is applying resource-intensive processing it can slow down your computer. At the expense of speed, adjusting the Max Processing Threads to a lower value limits the memory resources RX consumes in order to have a smoother experience while doing other tasks as RX runs in the background. This is especially useful when using the Batch Processor. The default value is 'Auto', and the minimum value is 1 thread.

License and Updates



- **Authorization**: Provides options to authorize or de-authorize RX (explained in the Authorization chapter).

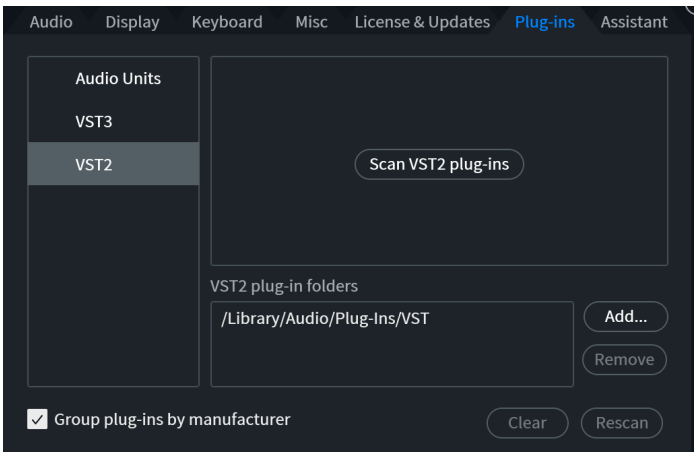
- **Help iZotope improve RX by sending anonymous usage data**: Enables sending anonymous usage data to help us improve our products.

Plug-ins

RX 11 Audio Editor supports the use of the following plug-in formats in the “Plug-in” module:

- **VST2:** Windows and Mac (Intel and Rosetta only)
- **VST3:** Windows and Mac
- **AU (AudioUnit):** Mac Only

> **UPDATE YOUR PLUG-INS**
>
> Please make sure your third-party plug-ins are updated to their latest version in order to ensure that RX will be able to scan them correctly. Contact your plug-in manufacturers for updated installers, if necessary.



> **ABOUT THE ABOVE SCREENSHOT**
>
> Only Windows users and Mac users working on Intel-based systems (or Silicon-based systems running Rosetta) will see the option for VST2. For all other systems, this format is unsupported and will not appear as an option.

- **Plug-in Lists:** Displays plug-ins that have been scanned for use in the “Plug-in” hosting module of the RX Editor.
- **Enable:** Enables that plug-in format for use in the RX Audio Editor. This will trigger plug-in scanning to begin in the background.
- **Disable:** Disables the associated plug-in format. This will clear the scanned plug-in list for that format. Re-enabling that plug-in format will prompt RX to re-scan that plug-in format.

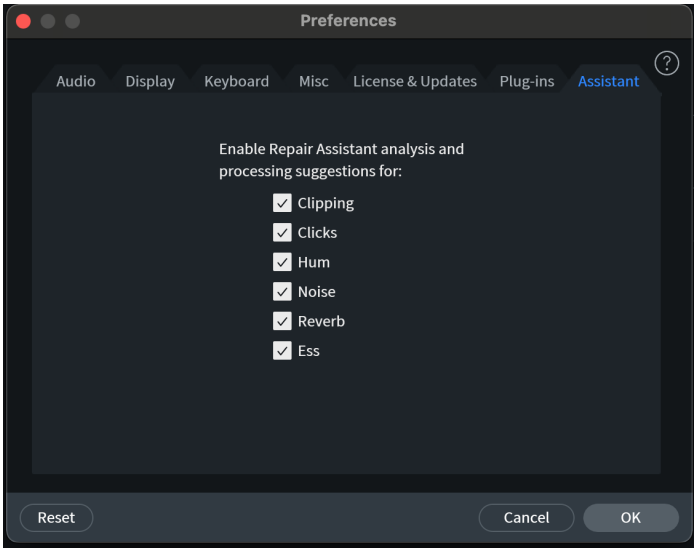
> **NOTE**
>
> If a plug-in failed scanning for any reason, the plug-in’s name will be prefixed with an error tag (ex: [Crashed] or [Failed]) to help troubleshoot the failure.

- **VST plug-in folders:** Allows you to add or remove custom VST2 plug-in folder paths. RX uses the system VST2 plug-in folder by default. If you are using a custom directory for VST2 plug-ins, use this option to ensure that those VST2 plug-ins will be scanned.

> **NOTE ABOUT SUB-FOLDERS WHEN SCANNING FOR PLUG-INS**
>
> RX will scan the first level of sub-folders in the custom VST2 folder. If some of your plug-ins do not show up when you scan them, and you know they're in a subfolder of your plug-in folder, try moving them up one directory level.

- **Group plug-ins by name in plug-in menus:** When enabled, the RX plug-in menu will group plug-ins by common first words, usually the manufacturer's name. When disabled, the RX plug-in menu will appear as a single, alphabetically sorted list.
- **Rescan:** If RX detects that a plug-in is unstable, it will blacklist it and prevent it from being opened. The rescan option allows you to clear the blacklist of unsupported plug-ins and rescan all installed plug-ins in case an RX update or an update from the plug-in manufacturer resolves the issue.

Assistant



- **Enable Repair Assistant analysis and processing suggestions for:** Informs what common audio problems will be considered when Repair Assistant analyzes a selection. When an option is checked in this list, the associated audio problem will be included in analysis and processing suggestions. Options include:
  - Clipping
  - Clicks
  - Hum
  - Noise
  - Reverb
  - Ess

12. Composite View

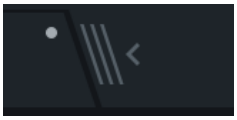
STD & ADV

Overview

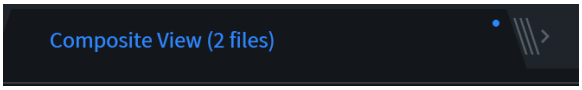
The Composite View feature in the RX Audio Editor combines all active tabs into one “Composite” tab that allows you to apply the same processing to multiple files simultaneously. Composite View can be a valuable tool for increasing efficiency when performing repetitive spectral editing functions.

Workflow

To enter Composite View mode, click on the Enter Composite View button.



When you are done making changes in Composite View, click on the Exit Composite View button to continue working with your individual tracks.



Important Notes

- **Composite View is designed to function as a bulk editor, it is not intended for mixing.** A maximum of 32 files can be collapsed into a Composite view tab.
- **Composite View assumes that all files start at the same point.**

  USE THE SIGNAL GENERATOR MODULE TO INSERT SILENCE

  If needed, you can adjust file start times by inserting Silence with the Signal Generator or using edit commands to modify timing before collapsing your files into a Composite View tab.

- **Composite View learns on the first track only.**
- **Sample Rates:** Composite View requires all files to have matching sample rates.

  RESAMPLE FILES USING THE RESAMPLE MODULE

  Use the Resample module to conform each file tab to the same sample rate value.

- **Spectrogram Display:** Composite View displays a summed composite spectrogram/waveform display for all files collapsed into the Composite tab. To view individual file displays, simply exit Composite View.
- **Editing and Processing:** All processing and edits applied in Composite View will be applied to all tracks.

- **Undo:** To revert edits made in Composite View on a subset of tracks you can exit composite view and use the Undo History list to undo changes made to individual files.
- **Playback:** Composite View plays back the sum of all files included in the Composite tab.
  - Channel selection states in individual tabs are not respected in Composite View playback or processing.
- Markers and Regions are not supported in Composite View.
- Compare Settings functionality is not supported in Composite View.
- De-bleed is not recommended for use in Composite View.
- Wow & Flutter will not work in Composite View.
- Loudness Control will not work in Composite View.

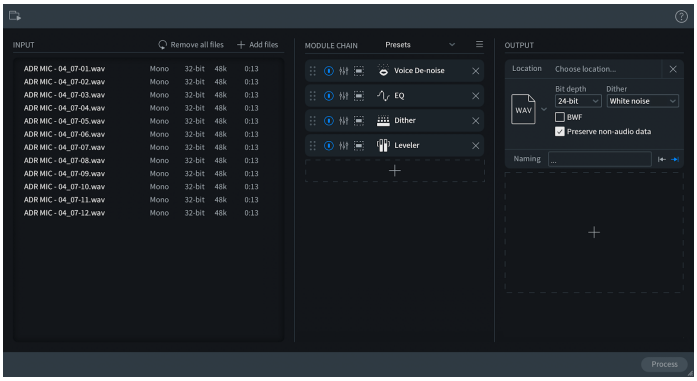
13. Batch Processor

Overview

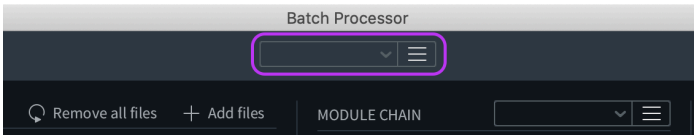
The Batch Processor allows you to process multiple files at once with a custom module chain, streamlining time-consuming tasks. The Batch Processor can run simultaneously with the RX Audio Editor, allowing you to process in bulk while working in RX.

The Batch Processor window can be opened using the following methods:

- Navigate to the RX Audio Editor “Window” menu and select “Batch Processor”.
- Open the Batch Processor window using the default keyboard shortcuts: Cmd+B (Mac); Ctrl+B (Windows).



Batch Processor Presets



Click here to save, load, or update a Batch Processor preset. Output options and module chain contents can be saved together as presets.

Learn more about saving presets

More information about saving presets is covered in the Common RX Module Controls chapter.

Input

The Input section allows for files to be added for processing. Added files will display the channel count, bit depth, sample rate and file length.

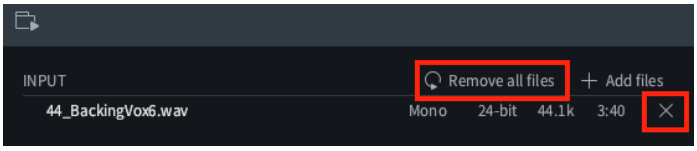
Adding Files

Files can be added to the Batch Processor window by using one of the following methods: Click the Add files button at the top of the Input section. Drag and drop files or folders from Finder (Mac) or Windows Explorer (Windows) into the Input files section of the Batch Processor window. Folders that are added will retain their folder structure in the output directory.

TIP

Add multiple files from the same folder by holding the Command key (Mac) or ctrl key (Windows) when selecting files from the Finder or Windows Explorer dialog.

Removing files



Individual files can be removed from the Batch Processor by clicking the 'X' that appears when hovering over a file. Clicking the Remove all files button at the top will clear all files from the input section.

Module Chain

Processing steps can be configured in the **Module Chain** area of the Batch Processor window. You can create a new Module Chain or recall a Module Chain preset by selecting it from the dropdown menu.

Learn more about the Module Chain

See the Module Chain chapter for more information about the controls available in the Module Chain.

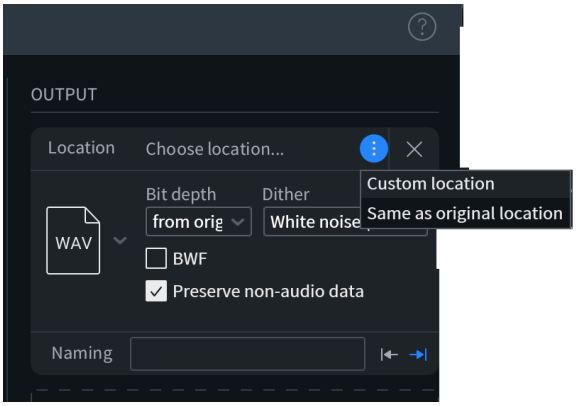
NOTE

Wow & Flutter is not available in the Batch Processor.

Output

The Output Section includes options for selecting the location of the processed files, modifying file names, and setting the output file format.

Use the "Choose folder..." button to select the destination for the processed files from the system dialogue that appears. To send files to the same folder they are sourced from, click the ellipsis button in the output panel and choose 'Same as Original Location'.



Input folder structure will be retained in the output location of your choice.

Naming

To modify a file name, add the desired text to the naming field and select the Before "|<-" or After "->|" icons.

File Format

Choose the output file format and configure the options associated with the selected file format (e.g. bit rate, quality, etc…). Available File Formats:

- WAVE
- AIFF
- FLAC
- Ogg Vorbis
- MP3

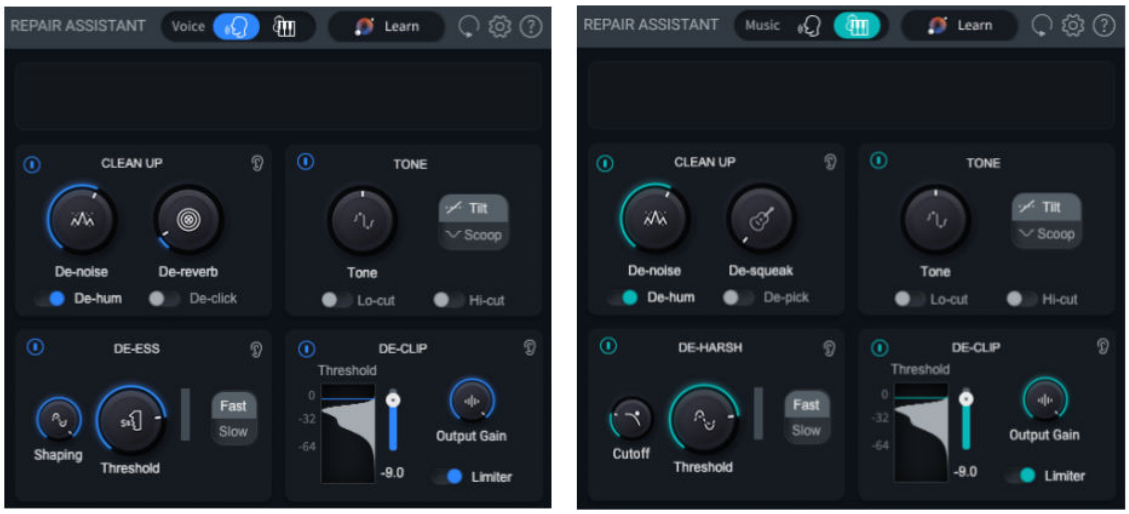
14. Repair Assistant

Overview

Repair Assistant is designed to detect common audio problems in a recording and quickly provide suggested settings for improving them. The Repair Assistant interface provides access to parameters that allow the user to fine-tune the suggestions to their own liking. Repair Assistant is available in the RX Audio Editor as a module and as a plug-in.

Repair Assistant Module

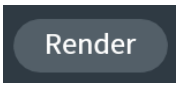
To open Repair Assistant in the RX Audio Editor application, simply click the **Repair Assistant** button in the upper right-hand corner of the RX Audio Editor window.



After opening the module, choose between **Voice** and **Music** modes. This changes the type of processing chains generated and is used in analysis to measure noise amounts and determine processing components and settings. **Voice mode** is intended for spoken word or dialogue. **Music mode** is intended for a full-mix music recording, individual instruments, or isolated sung vocals.

Next, click **Learn** to analyze your audio. Repair Assistant is designed to be trained on audio that contains dialogue, vocals, music, and sound effects. It isn't necessary to learn on noise-only segments of audio like with other modules such as Spectral De-noise. Repair Assistant analyzes any selection made in the Spectrogram (if no selection is made it analyzes as much data as necessary in order to make an effective suggestion) and automatically adjusts the settings for optimal repair of your audio. After analysis is complete, you will be able to adjust the various parameters of the tools that are working inside the assistant.

To apply Repair Assistant's recommended processing to the current selection, select the processing option in the Repair Assistant window and click the **Render** button in the bottom right-hand corner of the window.



The Open Module Chain allows you to save Repair Assistant's recommendations as a Module Chain preset, or make changes to the generated settings before rendering. To open Repair Assistant's recommendations in the Module Chain, click the **Open Module Chain** button in the Repair Assistant footer area (this feature is available in RX 11 Advanced only).

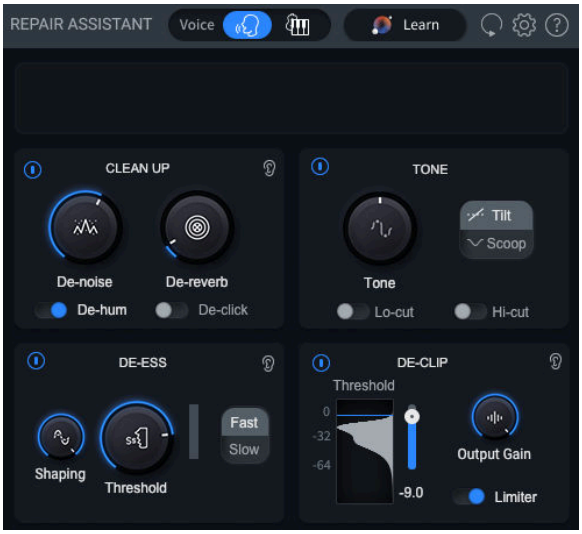


Repair Assistant Plug-in

The experience in the Repair Assistant Plug-in is very similar to that of the Audio Editor, but the plug-in analyzes your audio as you play audio, as opposed to analyzing the entire file/selection. To use the plug-in, simply choose your **Mode (Voice or Music)**, click **Learn**, and play your track. Keep playing until Repair Assistant has finished listening to your audio and created its suggestion.

> When using AudioSuite, clips must be longer than 7 seconds. Learning a clip less than 7 seconds long will cause the plug-in to 'learn' indefinitely and appear to hang in this state. When using AAX, you may loop shorter clips in order to run learning.

Voice Mode



Common Controls

- The On/Off button in the upper left corner of each section enables/bypasses that section.



- The Audition button in the upper right corner of each section allows you to hear only what is being removed by that section.



Clean Up

- **De-noise**: Controls the applied amount of Dialogue Isolation processing.
- **De-reverb**: Controls the applied amount of De-reverb processing.
- **De-hum (on/off)**: Activates/deactivates De-hum processing (De-hum will only operate if hum is detected during learning).
- **De-click (on/off)**: Activates/deactivates De-click processing.

Tone

- **%**: Controls the amount of boosting/cutting bi-directionally (based on EQ shape).
- **% (Tilt)**: Positive boosts the high end and lowers the low end; Negative boosts the low end and lowers the high end.
- **% (Scoop)**: Positive boosts the high end, boosts the low end, and lowers the mid range; Negative lowers the high end, lowers the low end, and keeps the mid range as-is.
- **Tilt/Scoop**: Choose which EQ shape to use.
- **Lo-cut (on/off)**: Activates/deactivates lo-cut filter.
- **Hi-cut (on/off)**: Activates/deactivates hi-cut filter.

De-Ess

- **Shaping**: Determines how much the spectral shape of the sibilant is changed.
  - A setting of 0% leaves the natural shape of the sibilance.
  - A setting of 100% flattens the shape of the sibilant toward a specified spectral profile.
- **Threshold**: Determines the level at which the De-ess module begins compressing sibilance.
- **Fast/Slow**: Sets the attack and release times for the processing.

De-Clip

- **Threshold**: Defines the level used for detection of clipped intervals. Generally, this should be set just below the actual level of clipping. To set the threshold, move the Threshold slider until it lines up with the place in the histogram just below where clipping is concentrated.
- **Output Gain**: Selects the gain to be applied to the selection after De-clip.
- **Limiter (on/off)**: Applies a true peak limiter after processing to prevent the processed signal from exceeding 0 dBFS.

Music Mode



Common Controls

- The On/Off button in the upper left corner of each section enables/bypasses that section.



- The Audition button in the upper right corner of each section allows you to hear only what is being removed by that section.



Clean Up

- **De-noise**: Controls the applied amount of Adaptive De-noising (best for eliminating room tone in vocals or mic'd acoustic instrument recordings).
- **De-squeak**: Controls the amount of processing for removing fret and finger noises/squeaks that occur when sliding up or down the guitar neck between notes or chords.
- **De-hum (on/off)**: Activates/deactivates De-hum processing (De-hum will only operate if hum is detected during Learning).
- **De-Pick (on/off)**: Activates/deactivates De-pick processing

Tone

- **% (Tilt)**: Positive boosts the high end and lowers the low end; Negative boosts the low end and lowers the high end.
- **% (Scoop)**: Positive boosts the high end, boosts the low end, and lowers the mid range; Negative lowers the high end, lowers the low end, and keeps the mid range as-is.
- **Tilt/Scoop**: Choose which EQ shape to use.

- **Lo-cut (on/off)**: Activates/deactivates lo-cut filter.
- **Hi-cut (on/off)**: Activates/deactivates hi-cut filter.

De-Harsh

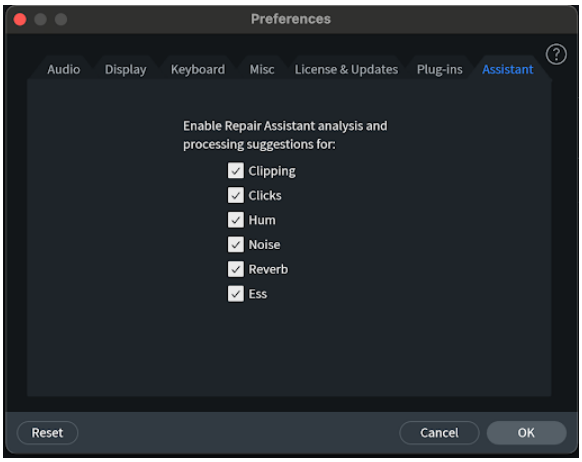
- **Cutoff**: Specifies the crossover point between musical signal (to be preserved) and "harshness" (to be reduced). The Cutoff frequency value functions as the lower boundary for harshness detection.
- **Threshold**: Determines the level at which the De-harsh module begins compressing harshness.
- **Fast/Slow**: Sets the attack and release times for the processing.

De-Clip

- **Threshold**: Defines the level used for detection of clipped intervals. Generally, this should be set just below the actual level of clipping. To set the threshold, move the Threshold slider until it lines up with the place in the histogram just below where clipping is concentrated.
- **Output Gain**: Selects the gain to be applied to the selection after De-clip.
- **Limiter (on/off)**: Applies a true peak limiter after processing to prevent the processed signal from exceeding 0 dBFS.

Repair Assistant Preferences

To bypass analysis by any of the Repair Assistant modules, click on the gear button in the upper right-hand corner of the Repair Assistant window to open the Assistant Preferences. All modules are selected by default, and deselecting any module will remove it from Repair Assistant's analysis.



These preferences apply only to Repair Assistant in the RX 11 Audio Editor application, not the plug-in. For more information about the RX Preferences menu, see the Preferences chapter.

15. Common Module Controls

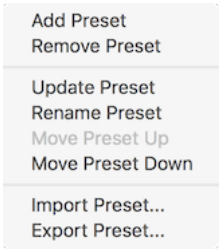
Overview

The RX Audio Editor is designed to give you a range of processing options. Most of the modules in RX feature multiple processing modes, ranging from fast algorithms that sound great on most material to very time-intensive algorithms for critical applications. Understanding the Presets, Preview and Compare controls will help you save time, especially when taking advantage of RX's more powerful processing modes.

Presets



Each module in the RX Audio Editor features a preset menu that allows you to choose between factory presets and custom presets that you have saved. Any preset saved in a module in the RX Audio Editor can be opened in the corresponding RX plug-in, when applicable.



- **Add Preset:** Creates a new preset.
- **Remove Preset:** Removes a preset from the drop-down list.
- **Rename Preset:** Changes the name of a preset.
- **Set Preset Shortcut:** Allows you to define a keyboard shortcut for any preset in order to recall and apply different module settings quickly.
- **Import Preset:** Allows you to import presets (from another machine or another user, for example).
- **Reload Preset:** Rescans your preset directory for this module in order to refresh the available preset list.
- **Explore Preset/Reveal Presets in Finder:** Opens Windows Explorer or Finder window to the location of your presets on disk.

Preset directory locations

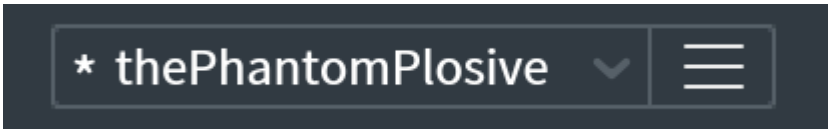
- Windows factory preset install location: *C:\Program Files\iZotope\RX Pro Audio Editor\Presets*
- Windows user preset save location: *C:\users\username\Documents\iZotope\RX Audio Editor\User Presets*
- Mac factory preset install location: */Library/Application Support/iZotope/RX Pro Audio Editor/Presets/*
- Mac user preset save location: *~/Documents/iZotope/RX Audio Editor/User Presets/*

If you previously installed RX Pro, then your presets folders will be in these locations:

- Windows user preset save location (RX Pro previously installed): *C:\users\username\Documents\iZotope\RX Pro Audio Editor\Presets*
- Mac user preset save location (RX Pro previously installed): *~/Documents/iZotope/RX Pro Audio Editor/Presets/*

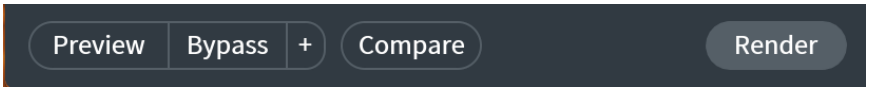
Preset Dirty Flag

Making changes to a preset will add an asterisk [*] to the preset name to show that it has been modified.



You can use Update Preset to commit your changes, or Add Preset to create another preset with your new settings.

Module Footer Controls



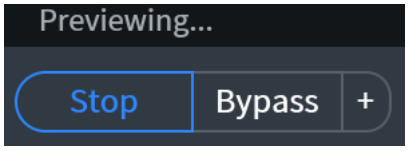
Preview

Many RX modules feature a Preview button in the bottom panel of the module window. Some modules don't support Preview because of the time-intensive nature of their processing, in many of these cases the Compare settings option is available (explained below).

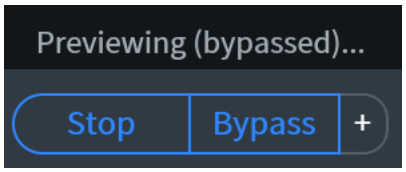
Preview allows you to make adjustments to controls and hear the results without the need to process and undo processing multiple times to achieve your desired results.

Preview will apply to your active selection or if no selection is made, Preview playback will start from the current playhead position. When Loop is enabled in the Transport, Preview playback will loop.

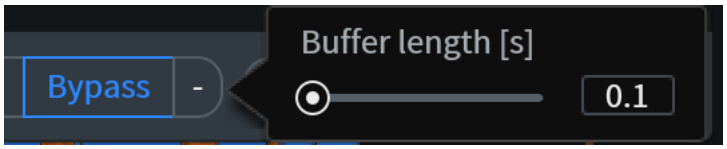
- **Preview [Shift-Space]:** Plays a pre-rendered preview of the module's current settings on the selected audio. During preview, module settings can be adjusted and adjustments will be heard within the length of the current preview buffer. For most modules, the Preview buffer size is about a half of a second long, but the Preview Buffer Length can be adjusted by accessing the Preview Options, explained below.



- **Bypass [Shift-B]** Bypasses module processing during preview.



- **Preview Options [+]:** Allows adjustment of the Preview Buffer size.



For more CPU-intensive settings, like the highest quality algorithms in Spectral De-noise and the highest quality De-click settings, RX can buffer playback to allow you to preview these slower than real-time processes.

NOTE ABOUT PREVIEW BUFFERING

When Previewing module processing, the active buffer length for preview rendering is tinted red in spectrogram/waveform display.

Pre- and Post-Roll

When previewing an effect, it is often very helpful to hear a small portion of the unprocessed audio before hearing your processed selection. This can provide a much clearer comparison and allow you to more easily discern whether or not the current processing settings are producing the desired effect.

By default, RX will play back one second of unprocessed audio before and after the current selection when previewing your processing. The Pre- and Post-roll times can be defined in the Preferences > Misc window. RX can play up to ten seconds of audio before or after the previewed selection. Pre- and Post-Roll will also occur when previewing a looped region of audio

HOW TO DISABLE PRE- OR POST-ROLL

Set the Pre- and Post-roll times in the Misc tab to 0.

TIP

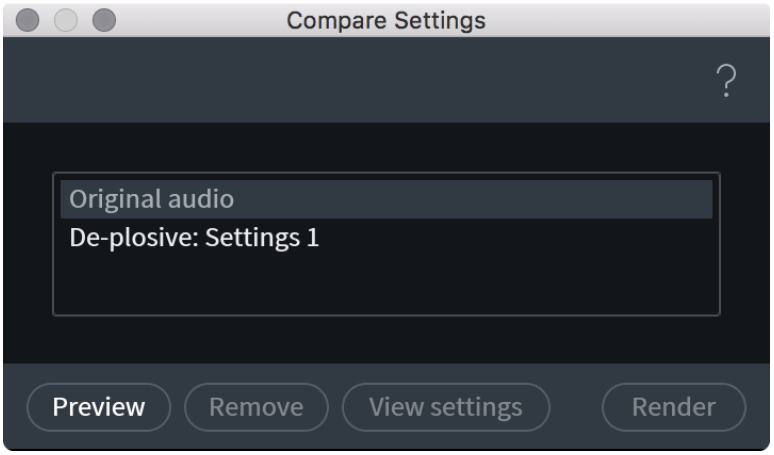
Pre- and Post-roll can also be simulated manually by holding Control (Windows) or Command (Mac) to set the playhead to any desired position while preserving your audio selection. Once the playhead is set, clicking on Preview in the desired module will then start the Preview playback from the playhead position.

Compare

When you want to quickly try a lot of different settings, use the Compare feature. In some cases, you might not know what settings of a module will give you the best, most transparent results. By hitting the Compare button instead of the Process button, you can audition multiple settings of the same module and then audition the results side by side in the Compare Settings.

While one group of settings is processing in the background, you can return to the module and try a different group of settings. Learning to use the Compare Settings tool can save you from having to apply and undo a process multiple times just to find the right settings, making it a valuable time-saving feature.

Another advantage to using the Compare Settings tool is seeing the effect your settings have in the spectrogram/waveform display and spectrum analyzer.

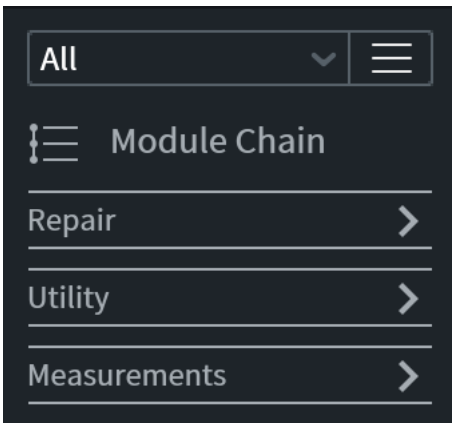


- **Process Comparison List:** Each time settings are sent to the Compare window, a new list item is created, by default titled “Settings 1,” “Settings 2,” etc.
- **Preview:** To hear (and see) the result of an item in the list, select that item and hit Preview.
- **View Settings:** Updates the controls to reflect the settings selected in the Process Comparison List
- **Remove:** Remove an item from the list.
- **Rename:** Allows you to rename items with more descriptive names.
- **Render:** Apply the selected list item to the audio file.

16. Module List

Overview

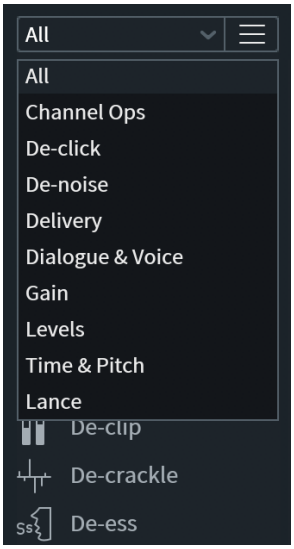
The Module List in the RX Audio Editor groups modules into the following categories: **Repair**, **Utility**, and **Measurement**. The modules displayed in the module list can be customized by saving a module list filter. The module categories can also be collapsed by clicking on the category header.



Repair	Utility	Measurements
Ambience Match **ADV**	Azimuth **ADV**	Find Similar
Breath Control **STD** & **ADV**	Dither **STD & ADV**	Markers
Center Extract **ADV**	EQ **STD & ADV**	Spectrum
De-bleed **STD** & **ADV**	EQ Match **ADV**	Waveform Stats
De-click	Fade	
De-clip	Gain	
De-crackle **STD** & **ADV**	Leveler **ADV**	
De-ess **STD** & **ADV**	Loudness Control **STD & ADV**	
De-hum	Mixing	
De-plosive **STD & ADV**	Normalize	
De-reverb **STD & ADV**	Phase	
De-rustle **ADV**	Plug-in	
De-wind **ADV**	Resample **STD & ADV**	
Deconstruct **ADV**	Signal Generator	
Dialogue Contour **ADV**	Time & Pitch **STD & ADV**	
Dialogue De-reverb **ADV**	Variable Pitch **STD & ADV**	
Dialogue Isolate **ADV**	Variable Time **STD & ADV**	
Guitar De-noise **STD & ADV**		
Interpolate **STD & ADV**		
Mouth De-click **STD & ADV**		

Repair	Utility	Measurements
Music Rebalance **STD & ADV**		
Spectral De-noise **STD & ADV**		
Spectral Recovery **ADV**		
Spectral Repair **STD & ADV**		
Voice De-noise		
Wow & Flutter **ADV**		

Saving module list filters



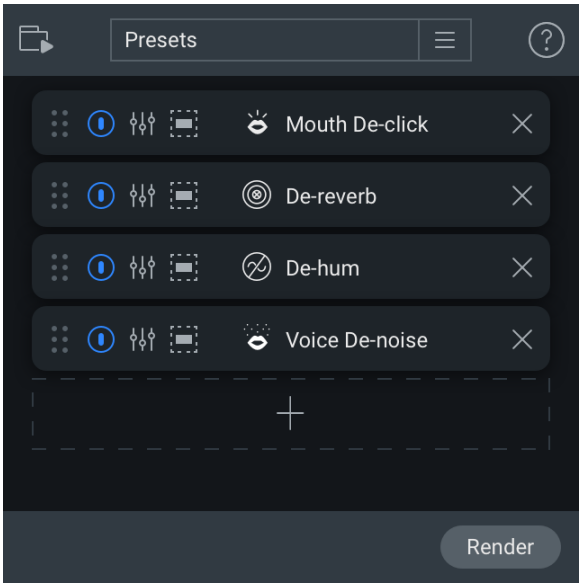
1. From the menu on the right of the List Filter Selector, select “Add List Filter”
2. Choose a name for your new List Filter, hit return/enter to save the name
3. Use the checkboxes to select the features you want to include
4. Click the “Done Editing" button to save your new filter

17. Module Chain

Overview

The Module Chain section allows for the processing of multiple modules in series. You can choose from a variety of factory presets to get started or configure your own custom chain and save a preset to use it again in the future. This section includes options for adding, removing, rearranging and customizing the processing of a module.

Controls



	Name	Description
	Preset Menu	Save or load a module chain preset. More information about saving presets is covered in the Common RX Module Controls chapter.
	Add Module	Click the [+] to add a module from the module list to the end of the module chain.
	Remove Module	Click the "X" icon to remove the associated module from the module chain.
	Power Button	Enables/disables module processing.
	Reorder Module	Click and drag a module panel up or down within the Module Chain to change its order in the signal flow.
	View Module	Opens the module window with settings loaded for the associated step in the module chain.
	Frequency Selection Settings	Options for processing a module only over a selected frequency range.

NOTE

Wow & Flutter is not available in the Module Chain.

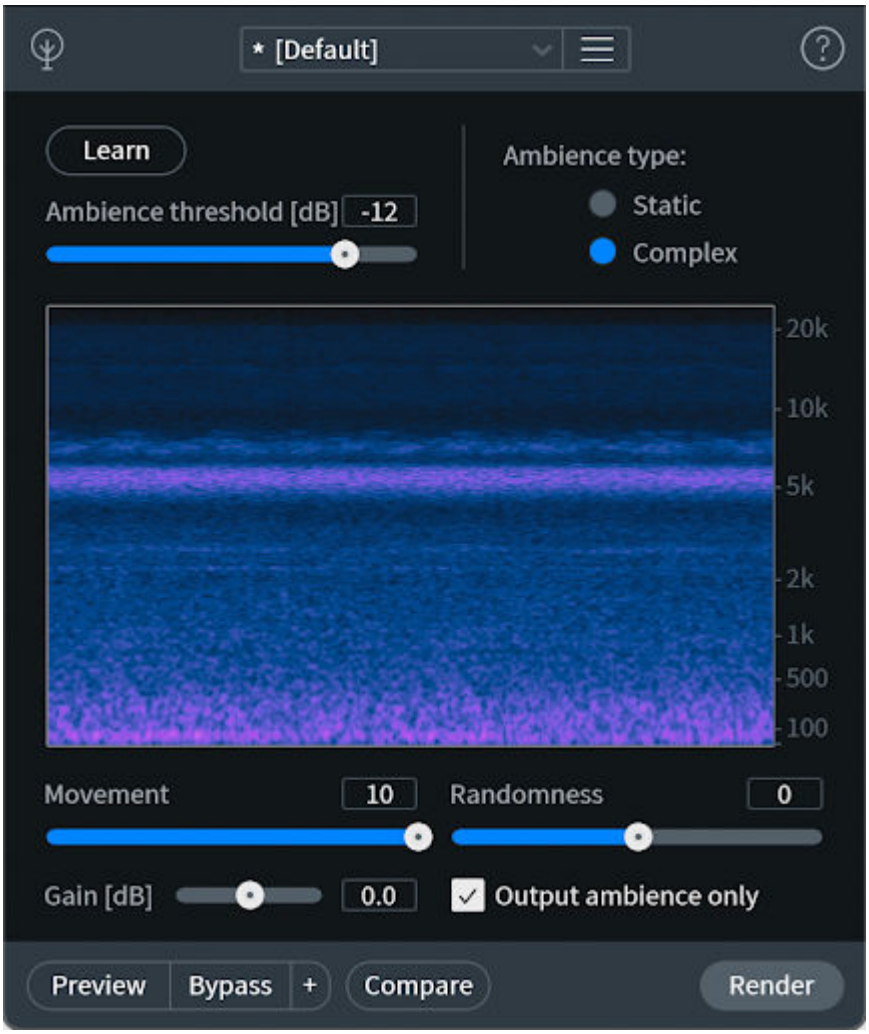
18. Ambience Match

ADV | Module and Plug-in (Audiosuite only)

Overview

The Ambience Match module lets you match the noise floor of one recording to another recording. For example, you can apply the ambience of a location recording to your ADR tracks.

Controls



- **AMBIENCE THRESHOLD:** Discards dialogue or noises above the ambience threshold when learning for complex ambience profiles. Complex mode only.
- **AMBIENCE TYPE:** Chooses between Complex or Static ambience.
- **MOVEMENT:** Higher values have more variety. Lower values remove distinctive elements for a more uniform sound. The elements that are removed as the movement value is reduced are greyed out in the Ambience Match spectrogram. Complex mode only.

- **RANDOMNESS:** Higher values allow the generated ambience to differ more from the learned input and encourage more variation. Lower randomness will try to match the learned input closely and may sound more like copy-pasting. The default Randomness setting will typically provide the most natural sounding results.Complex mode only.
- **GAIN:** Trims the level of the generated ambience.
- **OUTPUT AMBIENCE ONLY:** When checked, the selected region will be replaced by the generated ambience. When unchecked, the ambience will be mixed with the selected audio.

Ambience Type

Complex

Ideal for variable ambiences with movement and texture, Complex mode generates a more dynamic ambience based on your learned source. Each time you render, Complex mode will generate a new random ambience.

Static

Perfect for unchanging ambience like room tone, air conditioners, or fans, the static Ambience Match algorithm analyzes an audio selection and finds the lowest common denominator (the noise that is common across the audio file) and treats that as the ambient profile.

Workflow

Using Ambience Match module in the RX Audio Editor

To train Ambience Match, provide it with a selection of raw noise. If there is no single fragment of raw noise, or you want to save time, selections with speech can be used. In Complex mode, the Ambience threshold can be used to ignore speech or other loud sounds in the selection. In Static mode the algorithm will intelligently discard speech and only leave noisy parts in the noise print.

The Spectrogram in Ambience Match

Ambience Match has a helpful spectrogram that reflects your learned noise profile. In Static mode, the spectrogram will display what will be rendered after you learn. In Complex mode, the spectrogram will display your learned audio, and takes into account the ambience threshold setting.

To match the ambience between selections:

1. Open the Ambience Match module in the module list
2. Make a selection in a file.
3. Click Learn.
4. Make another selection.
5. Adjust the Gain level as desired. The Gain control adjusts the level of synthesized ambience.
6. Select Output Ambience Only if you want the selection replaced with only the ambience from the first selection.
7. Click Render.

NOTE

The Ambience Match module cannot reduce the amount of ambience that already exists in the selection, it can only increase it. To reduce the ambience, use the Spectral De-noise module.

To create an Ambience Match preset:

1. In Ambience Match, click the gear icon to the right of the preset drop-down menu.
2. Select Add Preset.
3. Enter the name for the new preset.
4. Press Enter.

Using Ambience Match as an AudioSuite Plug-In

In addition to applying Ambience Match inside of the RX Audio Editor, it can also be used as an AudioSuite plug-in inside of Avid's Pro Tools or Media Composer.

When using Ambience Match inside of Pro Tools or Media Composer, we recommend not learning from audio that contains fades within the selection or the handles. As Ambience Match establishes an ambient profile using the lowest common denominator, learning from audio that's being faded in may result in inconsistent detection of the noise floor. Handles can be preserved by using Ambience Match in clip-by-clip mode.

NOTE ABOUT LEARNING AMBIENCE ON FADES

When using Pro Tools, you may see an inconsistent result as a result of Pro Tools adding dithering to fades, which varies based on the session's bit depth. Since this dithering noise almost certainly doesn't match the material's noise profile, this will throw off Ambience Match. If you're running a session in 16 bits, the dither added by a fade will be sufficient enough to affect the detection algorithm. The problem is less pronounced in 24 or 32 bits. To adjust the bit depth of your session, go to Setup > Session in Pro Tools.

19. Breath Control

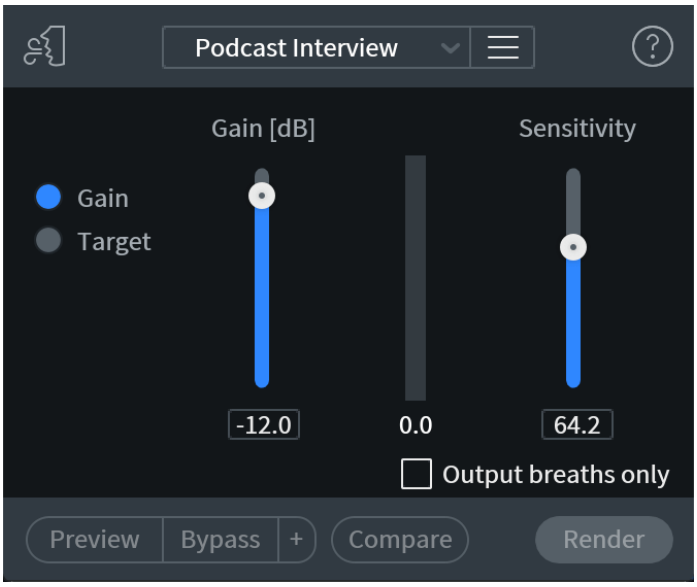
STD & ADV | **Module & Plug-in**

Overview

The Breath Control module intelligently detects breaths in dialogue or vocal recordings and suppresses them. Removing and reducing breaths in recordings can be a time-consuming process for dialogue editors and music producers alike. Breath Control can help reduce the time spent on repetitive editing without sacrificing the quality of your dialogue or vocal recordings.

The Breath Control module analyzes your file and distinguishes breaths from dialogue or sung vocals based on their harmonic structure. If any piece of the incoming audio matches a harmonic profile similar to a breath, the module will suppress that portion of the audio until sung vocals are detected. Different from a 'Threshold' based process in which the module is only engaged once the audio has risen to a certain volume, Breath Control will perform its analysis regardless of level. This allows for accurate breath recognition for a multitude of quiet or loud vocal styles with minimal adjustment of the module's controls.

Controls



- **GAIN MODE:** When Gain Mode is selected, the Breath Control module will reduce the gain of detected breaths by an absolute amount, regardless of the level of the breath. In some cases, this is desirable when trying to handle heavy breathing or as a way of removing all breaths from a particular spoken or sung vocal take. Depending on your settings, this can result in unnatural sounding results as the very quiet breaths may be inaudible, while the loud breaths will be reduced to a normal level.
  - **GAIN (dB):** Sets the desired amount of gain reduction applied to all detected breaths, regardless of level.

- **TARGET MODE:** When in Target mode, the reduction amount of the 'Target' slider will set the level of detected breaths. This can result in more natural sounding breath reduction as detected breaths are only reduced when necessary. Loud and abrasive breaths will be reduced heavily, while quiet, natural sounding breaths will be left at the same volume.
  - **TARGET LEVEL [dBFS]:** Sets the resulting desired level of all detected breaths above the set target.
- **SENSITIVITY:** This controls how sensitive the breath control plug-in is when detecting the harmonic structure of breaths in your audio.
- **OUTPUT BREATHS ONLY:** When enabled, only the audio of the detected breaths will be passed to the output of the module. This can help when setting the Sensitivity control in order to make sure that only the breathing in your audio is being processed.

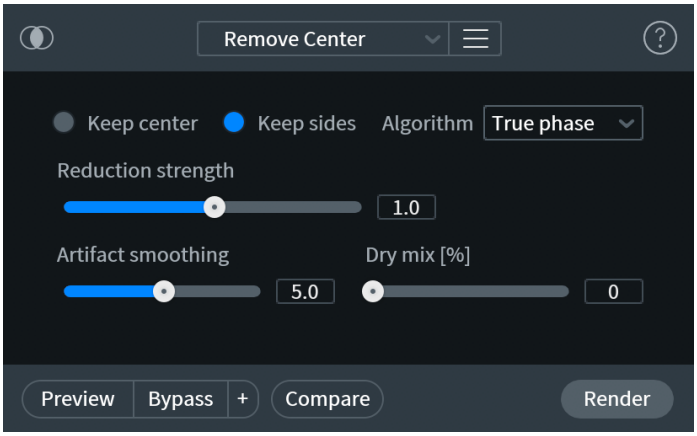
20. Center Extract

ADV

Overview

The Center Extract module preserves (using **Keep Center**) or removes (using **Keep Sides**) the center channel of a stereo file. Extracting the center will retain the center of a stereo field and attenuate everything on the sides, such as signals panned to the left or right. See the Use Cases section below for more contextual examples and additional information about Center Extract processing.

Controls



- **KEEP CENTER:** When the signal you want to preserve is even in both channels and noise is uneven between channels, extracting the center can remove a lot of noise.
- **KEEP SIDES:** If you want to preserve the wide stereo information and remove the center information, you can keep the sides of the signal instead.
  - **ALGORITHM:** Two different algorithms are available:
    - **TRUE PHASE:** Cancels the center with phase information and retains the original panning of the sides.
    - **PSEUDO PAN:** Extracts the side information and artificially stereo-izes it into two channels.
- **REDUCTION STRENGTH:** Controls the level of the preserved signal. Lower values will retain more information, higher values will discard more information.
- **ARTIFACT SMOOTHING:** Helps to reduce or eliminate the "musical noise" that is often characteristic of FFT-based processing. Musical noise can be described as how something may sound underwater. Increase this slider if your output sounds watery, but decrease it when too much smoothing makes your audio sound dull.
  What is an FFT?
  Fast Fourier Transform: a procedure for the calculation of a signal frequency spectrum. The greater the FFT size, the greater the frequency resolution, i.e., notes and tonal events will be clearer at larger sizes. However, when using FFT-based processing, the more audio you remove from your source, the more likely you are to create undesirable artifacts.

- **DRY MIX [%]:** Controls the amount of unprocessed signal mixed into the processed signal. Useful for reducing artifacts introduced by processing by preserving the original characterisitics of your audio.

TIP: USE AZIMUTH BEFORE CENTER EXTRACT FOR BEST RESULTS

It is often a good idea to make sure stereo channels are balanced by running Azimuth correction before using Center Extract.

Center Extract is not available when Composite View is active.

Center Extract processing is not available on mono files. The nature of Center Extract processing makes it unapplicable to mono files because they lack stereo field information.

Use Cases

- **Using Center Extract as an alternative to Mid-Side Encoding**: Center channel extraction will preserve a stereo image if the side channels are retained. This can make it more desirable in some cases than Mid-Side encoding (which would sum left and right hard pans into one channel).
- **Use Keep Center to Reduce noise in stereo files transferred from a mono source**: A mono record transferred to a stereo tape would have side channel noise that would be suppressed by extracting the center channel using Keep Center.
- ***Use Keep Sides to Remove vocals from a stereo recording**:
  - The lead vocal track, in many popular mixes, is typically panned to the center of the mix. Panning something to the center results in equal information being present in the side (Left and Right) channels.
  - Using the “Keep Sides” processing mode will retain the unique side channel information present in a file, and reduce the Center channel.
  - This is useful for karaoke-style removal of vocals from a song, especially because the process results in a coherent stereo image.

21. De-bleed

Overview

De-bleed reduces the leakage of one signal into another, such as when vocals bleed into a guitar microphone, or when a click track fed into headphones bleeds into an open mic. The De-bleed module learns a bleed relationship between two tracks. In the descriptions below, these two tracks are referred to as:

- **BLEED SOURCE TRACK:** This track contains only the bleed source audio.
- **ACTIVE TRACK:** This track contains the audio present in your source track (bleed), mixed in with the audio you want to preserve.

De-bleed relies on a relationship existing between your Active track and Bleed Source track in order to function properly and provide the best results. In order to properly establish the relationship between the Active and Bleed Source tracks, the following requirements must be met:

- **At least two files are open in the RX Audio Editor:** De-bleed requires 2 tracks (a Bleed Source Track & an Active Track) to process.
- **The Bleed Source and Active Track have matching sample rates**: The sample rates of the Bleed Source Track and Active Track must match in order to process using De-bleed. Using tracks of different sample rates may indicate that the tracks are not related.

  > **TIP: HOW TO CORRECT SAMPLE RATE DIFFERENCES**
  >
  > If you need to correct the sample rates of your bleed source and active track selections in the De-bleed module, you can use the Resample module in the RX Audio Editor to correct sample rate differences between your files, if necessary.

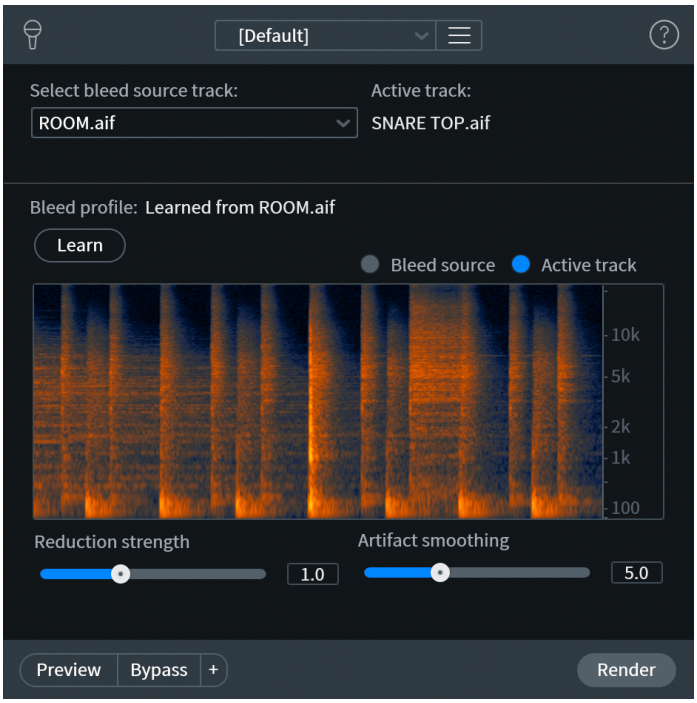
- **The Bleed Source and Active tracks are time aligned**:
  - The Bleed Source Track and Active Track must be time aligned within a few milliseconds of each other.
  - *Time aligned* means that if the two files were played back together, they would sound in sync. So if the same audio events are occurring at the same points in the timeline for both the Bleed Source and Active tracks, this means they're time aligned.

You can adjust the length or timing of files in the RX Audio Editor in the following ways:

▶ To remove the active selection, use the **Cut** operation from the **Edit** menu. The key commands are Command+X (Mac) or Ctrl+X (Windows).

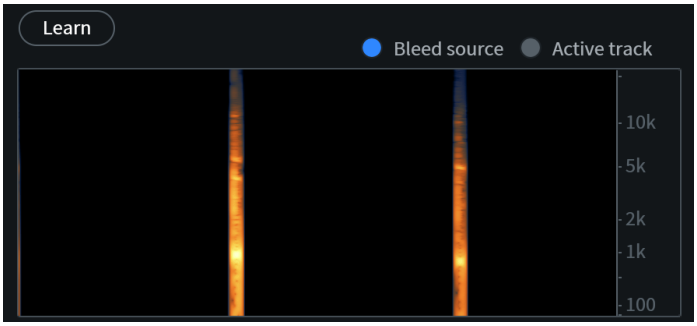
▶ To adjust the length of your file, apply Insert Silence using the Signal Generator module.

Controls

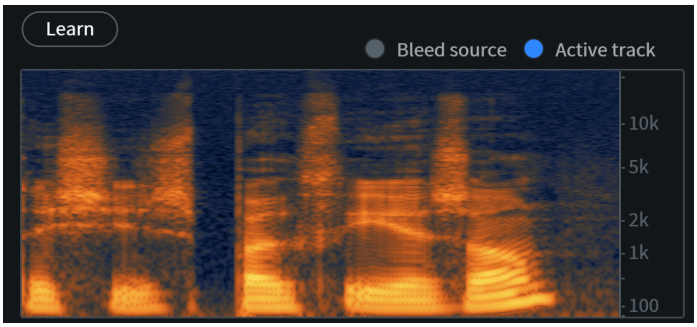


- **BLEED SOURCE TRACK SELECTION MENU:** Select the Bleed Source track from this dropdown.
- **ACTIVE TRACK:** Displays the name of the track tab you are currently viewing.
- **LEARN:** Learns the relationship between the two tracks.
- **BLEED PROFILE DISPLAY:** After learning the bleed profile, this displays a portion of the relationship captured between the two tracks. Toggle between the Bleed Source and Active track displays to check that the bleed present in the bleed source track is present and aligned with the bleed in your active track.

Bleed source display:



Bleed active track display:



- **STRENGTH:** Determines the amount of bleed reduction applied during processing. Higher Strength values may result in the removal of audio you wish to preserve. It is recommended you start with lower strength values and increase the values if needed to achieve the most desirable results.
- **ARTIFACT SMOOTHING:** Helps to reduce or eliminate the "musical noise" that is often characteristic of FFT-based processing. Musical noise can be described as how something may sound underwater. Increase this slider if your output sounds watery, but decrease it when too much smoothing makes your audio sound dull.

> **WHAT IS AN FFT?**
>
> **Fast Fourier Transform:** a procedure for the calculation of a signal frequency spectrum. The greater the FFT size, the greater the frequency resolution, i.e., notes and tonal events will be clearer at larger sizes. However, when using FFT-based processing, the more audio you remove from your source, the more likely you are to create undesirable artifacts.

Workflow

1. Import the bleed source file (**Bleed Source Track**) and the file you are removing bleed from (**Active Track**) into the RX Audio Editor. Ensure that the files are in sync (instructions above).
2. Open the De-bleed module.
3. Ensure that the **Active Track** name that is displayed in the De-bleed module is the file you wish to modify.
4. In the De-bleed module, select the **Source Track** from the **Bleed Source Track** drop-down menu.
5. In the active file tab (**Active Track**), make a selection where the bleed is most obvious.
6. Click the Learn button in the De-bleed module. *Learn analyzes the general relationship between the Bleed Source and the Active Track.*
7. After the Learn pass is complete, select part or all of the Active Track and click Process.

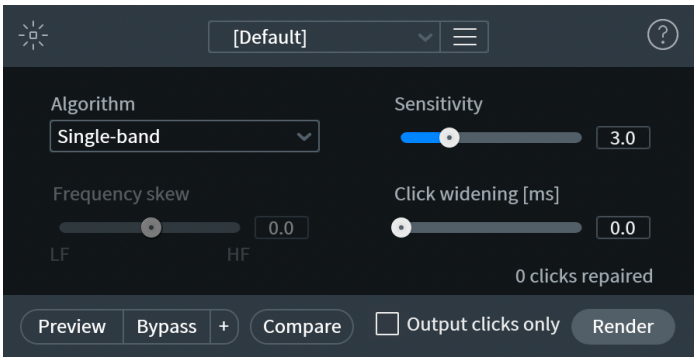
22. De-click

Module & Plug-in

Overview

The De-click module's sophisticated algorithm analyzes audio for amplitude irregularities and smoothes them out. This means that you can use De-click to remove a variety of short impulse noises, such as clicks caused by digital errors, mouth noises, and interference from cell phones.

Controls



- **ALGORITHM:** Provides options for adjusting the configuration and processing quality of click interpolation, depending on the types of clicks and pops present in the audio.
  - **SINGLE BAND:** Processes quickly and works well on very narrow "digital" clicks
  - **MULTI-BAND (PERIODIC CLICKS):** Uses multi-band processing for removing regularly repeating clicks with a wider spectrum, or regular clicks that have concentrated low or high energy (like thumps or optical soundtrack perforation noise)
  - **MULTI-BAND (RANDOM CLICKS):** Uses multi-band processing for wider vinyl clicks and thumps, with a protective algorithm for preserving periodic audio elements characteristic to certain instruments such as brass or vocals
  - **LOW LATENCY:** Works well on mouth clicks and other clicks that cannot be handled by other algorithms. This mode has very low latency and is suitable for real-time work in RX De-click plug-in.
- **SENSITIVITY:** Determines how many clicks are detected in the signal. Increasing sensitivity can impact plosives, which can in turn reduce or damage the original signal.
- **FREQUENCY SKEW:** Adjusts the weighting of click removal toward higher or lower frequency clicks. Negative values are more suitable for generic clicks such as those found on vinyl recordings. A setting of zero or above targets mouth clicks in the middle frequencies.

> **NOTE**
> Frequency skew is not available when using the Single Band Algorithm option.

- **CLICK WIDENING:** Extends the repair area around detected clicks, compensating for mouth sounds such as lip smacks that have a decay.
- **OUTPUT CLICKS ONLY:** Outputs the difference between the original and the processed signals (suppressed clicks).

Instant Process Tool

The Instant Process Tool offers a smart De-click mode, which instantly applies the active settings from the De-click or Interpolate modules. Simply put, you may make any selection, and this mode will automatically remove all clicks present in that selection. This is particularly useful for editing a dialogue file, mismatching sample rate clicks and pops, and vinyl clicks.

If you've made a selection under 4000 samples in length, Instant Process will automatically use the Interpolate algorithm. Interpolate will fill the selection with audio information based on the surrounding audio.

For selections above 4000 samples, Instant Process will use the current settings of the De-click module. De-click is effective on selections above 4000 samples in size, identifying clicks in relation to desirable audio, and intelligently separating and removing the clicks.

For example, if the De-click module is indicating that a preset named 'Remove mouth clicks' is loaded, these settings will be applied every time you use the Instant Process Tool in 'De-click' mode (on selections longer than 4000 samples).

23. De-clip

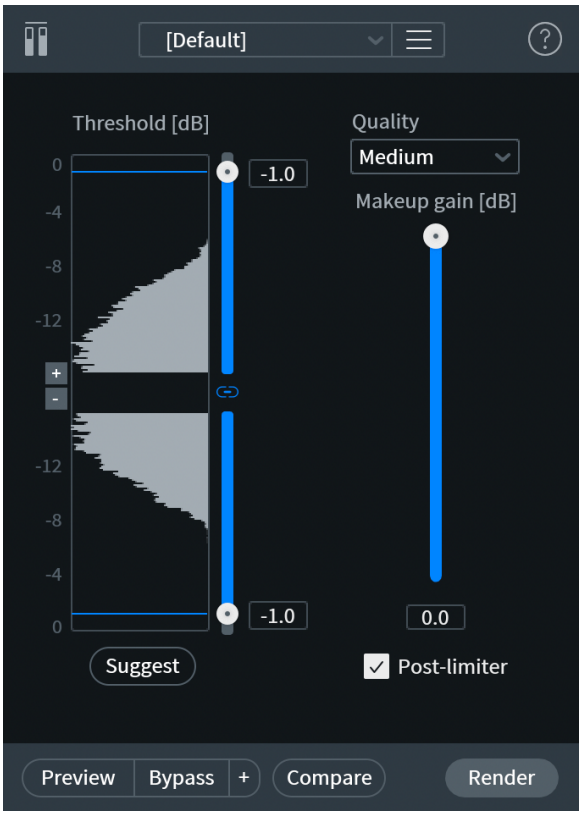
Module & Plug-in

Overview

De-clip repairs digital and analog clipping artifacts that result when A/D converters are pushed too hard or magnetic tape is over-saturated. It can be extremely useful for rescuing recordings that were made in a single pass, such as live concerts or interviews, momentary clipping in “perfect takes”, and any other audio that cannot be re-recorded.

De-clip will process any audio above a given threshold, interpolating the waveform to be more round. Generally, the process is as easy as finding the clipping you want to repair, then setting the threshold just under the level where the signal clips.

Controls



- **HISTOGRAM METER:** Displays waveform levels for the current selection as a histogram. A histogram is an analytical tool that displays how many samples are present at a given signal level over a window in time. The longer the line for the histogram is, the more energy is present at that amplitude. If a lot of energy tends to collect near the top and bottom edges of a waveform, that waveform is probably clipped and distorted. The histogram meter helps you set the Threshold control by displaying the audio level where the waveform’s peaks are concentrated. This usually indicates at what level clipping is present in the file.

- **HISTOGRAM ZOOM CONTROLS:** The histogram's range can be scaled if you need a better view of your signal. Use the (+) and (-) buttons to scale your display and value resolution for the De-clip module. These buttons reduce (+) and/or expand (-) the range of the threshold slider and histogram. You may want to extend the histogram range when the clipping point is lower than what you can see on a histogram or if you don't see anything on the histogram.
  - In the RX Audio Editor, the Histogram meter updates based on selection: Select a section of the recording where clipping is prominent and De-clip will analyze the levels of the program material. If clipping is present in the selection, it will usually appear as a horizontal line in the histogram that extends all the way across the meter.
  - In the De-clip Plug-in, the histogram runs as a real-time meter.
- **THRESHOLD [dB]:** Defines the level used for detection of clipped intervals. Generally, this should be set just below the actual level of clipping. To set the threshold, move the Threshold slider until it lines up with the place in the histogram just below where clipping is concentrated. Adjusting the Clipping Threshold will display a blue line within the histogram and a gray line on the waveform itself (when the De-clip Threshold effect overlay). These lines indicate the audio information that will be considered as "clipping" by the de-clip algorithm.
- **THRESHOLD LINK:** Toggles the ability to adjust positive and negative clipping thresholds independently.



  - When this option is enabled, you can adjust the positive and negative clipping Threshold controls independently. This is useful in cases where more clipping is occurring on one side of your waveform.
  - You can also set asymmetric de-clipping thresholds directly from the waveform by toggling the lock box between the threshold controls on the waveform display.
- **SUGGEST:** Calculates suggested threshold values based on the levels in your current selection.
- **QUALITY:** Controls the interpolation processing quality. There are three quality modes in the De-clip module: Low, Medium, and High.
  - When enabled and the De-clip module is open, the De-clip threshold overly will be displayed in the spectrogram/waveform display.
  - High quality mode processes slowly but is capable of achieving better results.
  - In many cases you will find that Low quality mode gives you great results. To save time, always start by previewing the Low quality modes first. You can also use the Compare feature to try multiple modes and preview the results.
- **MAKEUP GAIN [dB]:** Selects the gain to be applied to the selection after De-clip. Since the De-clip process causes an increase in peak levels, the Makeup gain control can be used to prevent the signal from clipping after processing. It is also useful for matching the level after processing to unprocessed audio outside of the selection.
- **POST-LIMITER:** Applies a true peak limiter after processing to prevent the processed signal from exceeding 0 dBFS.
  - De-clip usually increases signal levels by interpolating signal segments "above" the clipping point, which can make the signal clip again if the waveform format offers no headroom above 0 dBFS.
  - If the post-limiter is disabled, the restored intervals above 0 dBFS can be safely stored even without makeup gain as long as the file is saved as 32-bit float. Intervals above 0 dBFS will clip when played back through a digital-analog converter.

Using the De-clip Threshold Effect overlay in the Spectrogram/waveform view

- By default, De-clip Threshold is enabled in the View > Effect Overlay menu.
- When enabled and the De-clip module is open, the De-clip threshold overly will be displayed in the spectrogram/waveform display.
- You can adjust the Threshold controls in the De-clip module by adjusting the Threshold overlay lines in the spectrogram display.
- You can use the mousewheel on the waveform amplitude ruler to adjust the threshold control values.

More Information

Suggestions for severe distortion

- For certain situations, using the Deconstruct module to extract the noise components of the distortion can help remove additional artifacts beyond the clipped peaks in a waveform.
- In cases where severe distortion is visible on the spectrogram, the Spectral Repair tool can be used to select those problem areas, and attenuate or replace them with undistorted audio.

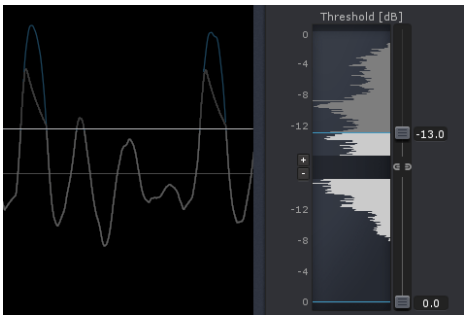
Visual Examples

Before & After Clip Repair

A waveform before and after clip repair. The after example (bottom) shows the original repaired waveform (faded) and the post-limiting waveform (bright). 

Unlinking Threshold Controls To Curb Asymmetric Clipping

This problematic waveform (grey) shears off around −13 dB on only the positive side of the waveform (the histogram on the right shows the extra positive energy of clipped audio). Extra processing of the negative side would be unnecessary, so the Threshold controls can be unlinked to process above −13 dBFS on the positive side only. The resulting waveform is drawn in blue above the grey sheared peak.



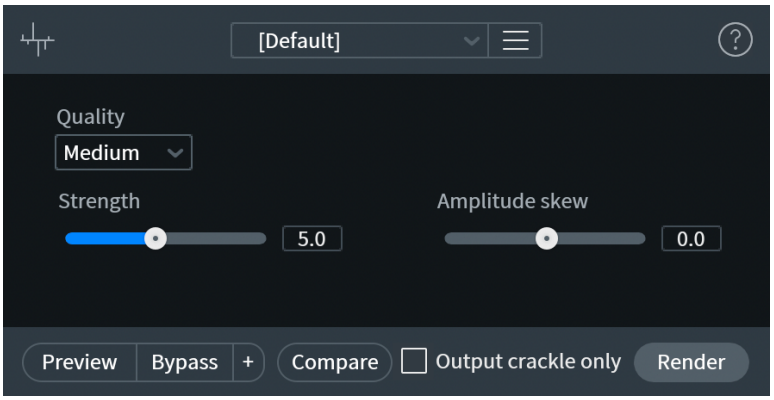
24. De-crackle

Module & Plug-in

Overview

When an audio signal contains many clicks that are close together and low in volume, this is described as crackle. De-crackle is very effective at removing these audio problems and even more effective if it is run after De-click has eliminated the worst offending clicks.

Controls



- **QUALITY:**
  - **LOW** quality offers fast processing
  - **MEDIUM** quality will remove periodic, rapidly repeating clicks
  - **HIGH** quality will help preserve the tonal qualities of a signal
- **STRENGTH:** Controls the amount of crackle that is detected and repaired.
- **AMPLITUDE SKEW:** Biases the processing toward higher or lower volume crackle. If the crackle accompanies transients and other high-level signal passages (such as during clipping), set this control more to the right. If the crackle mostly happens during low-amplitude signal passages, set this control more to the left.
- **OUTPUT CRACKLE ONLY:** Outputs the difference between the original and processed signals (suppressed crackle).

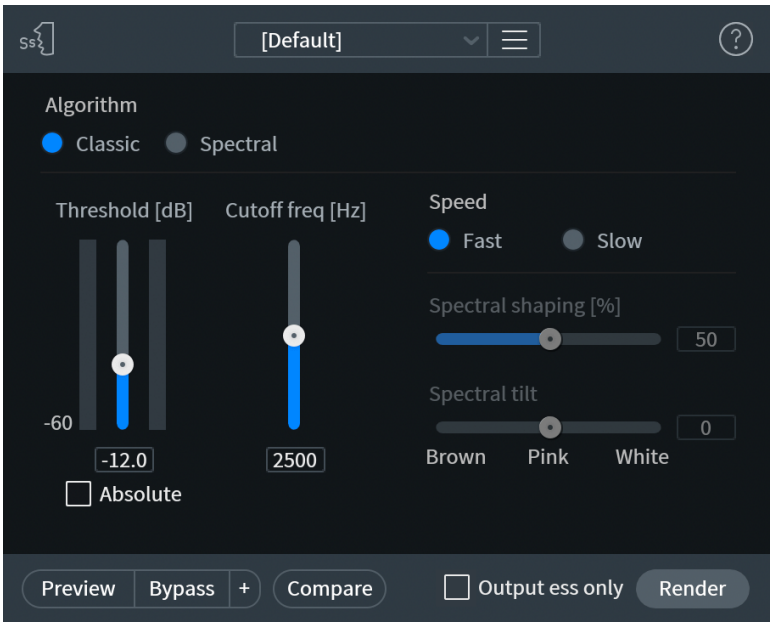
25. De-ess

Module & Plug-in

Overview

De-ess attenuates or reduces sibilance, the harsh high-frequency sounds that come from “S,” “F,” “X,” “SH,” and a soft “C.”

Controls



- **MODES:** The De-ess module offers two processing algorithms:
  - **CLASSIC MODE:** Detects sibilants and attenuates them with a broadband gain envelope. Since attenuation is applied to all frequencies, this mode is less targeted than Spectral De-ess.
  - **SPECTRAL MODE:** Offers a more transparent, intelligent, and frequency-specific type of de-essing than Classic Mode. Spectral Mode only attenuates the high frequencies where sibilance is most active, leaving the lower frequencies untouched. Technically, Spectral De-ess is a multiband compressor with dozens of bands. It’s able to compress the level of sibilants, shape their spectra, and avoid modulation of ambient noise. Each band can operate independently or with a link to adjacent bands (adjusted by the **Spectral Shaping** slider, explained below) and band thresholds can be adjusted for the desired shape of a sibilant (achieved by adjusting the **Spectral Tilt** slider).

- **THRESHOLD:** Determines the level at which the De-ess module begins compressing sibilance. The Threshold control has two modes that determine how it reacts to incoming signal level. It is specified in decibels, relative to speech level (Relative Mode) or full scale (Absolute Mode).
  - **RELATIVE MODE:** Determines the level of speech and sets the threshold relatively to that level. This is the default Threshold mode, Relative mode is active when the "Absolute" checkbox below the Threshold slider is *not* checked.
  - **ABSOLUTE MODE:** Sets the threshold to a decibel level below full scale (dBFS). This mode is enabled by checking the "Absolute" checkbox below the threshold slider.
- **CUTOFF FREQUENCY:** Specifies the crossover point between speech (to be preserved) and sibilance (to be reduced). The Cutoff frequency value functions as the lower boundary for sibilance detection.
- **SPECTRAL FLATTENING:** Spectral Flattening determines how much the spectral shape of the sibilant is changed. Think of Spectral Flattening as a way to fine tune the strength of the sibilant processing. The flatter you go, the more the sibilant is reduced. A setting of 0% leaves the natural shape of the sibilance by applying uniform compression across all bands. A setting of 100% flattens the shape of the sibilant toward a specified noise profile (see Spectral Tilt).
- **SPECTRAL TILT:** Spectral Tilt creates a target noise profile for the sibilance. It gives you the flexibility to determine what the ideal shape of your high frequency signal could be. Moving the shape toward brown noise moves it toward a darker sound. Moving the shape toward white noise provides a brighter result. Together with flattening and threshold, you can determine how much shaping takes place. A setting of 0 creates a natural spectral decay similar to pink noise. Values below or above 0 create a profile that is heavier in low frequencies (like brown noise) or high frequencies (like white noise). Spectral Tilt is most effective when the Spectral Shaping control is set to a non-zero value.
- **SPEED:** Sets the attack and release times for the processing. Attack times are program-dependent in both modes.
  - **FAST:** Uses quicker attack and release times. If De-ess is causing pumping in the high frequencies, try using Fast mode to mitigate this problem. If the settings are too slow, the processor won't recover quickly enough, resulting in too much high frequency attentuation.
  - **SLOW:** Uses longer attack and release times. If De-ess is smoothing out the transients too much, try using Slow mode to mitigate this problem. Settings that are too fast can reduce the high frequency signal too much in the initial transient phase and can introduce a smoothing effect that reduces useful high frequency definition.

26. De-hum

Module & Plug-in

Overview

De-hum is designed to remove persistent tonal noise, like AC noise, electrical buzz or interference introduced by EMF (Electro-magnetic fields). De-hum includes 2 modes: a Static mode that's best for ground noise or simple hums with a few harmonics, and a Dynamic mode for complex hums or buzz.

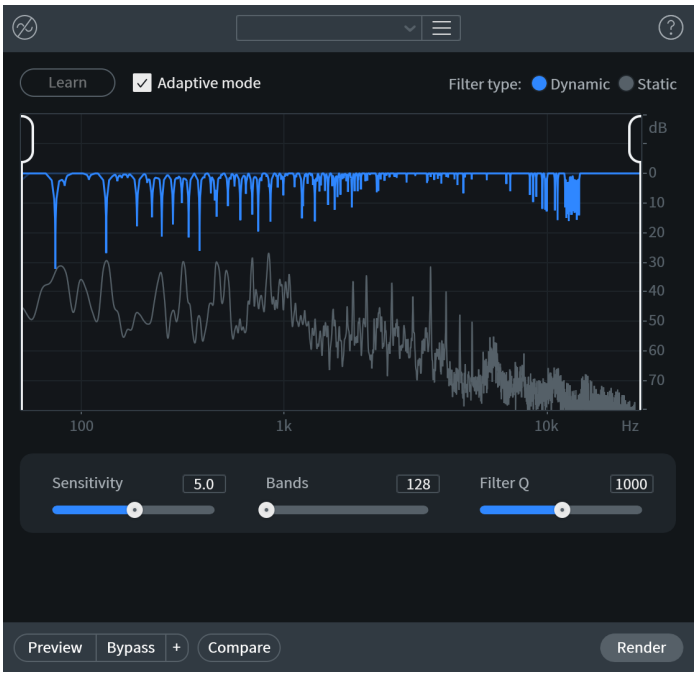
Mode Selection

- **Static:**
  - Ideal for simple tonal noise with only a few harmonics of the fundamental.
  - Hum that doesn't overlap the speech frequencies.
- **Dynamic:**
  - Hum that has more than a dozen harmonics.
  - Hum that is inharmonic (contains multiple unrelated tones).
  - Buzz or EMF that extend into the upper frequencies.
  - Hum/buzz that overlaps speech frequencies.
  - Less Ringing due to gated notching

MORE INFORMATION ABOUT DYNAMIC MODE

For more information about the benefits of the Dynamic Mode in De-hum, read this article by iZotope Principal DSP Engineer Alexey Lukin.

Controls in Dynamic Mode



- **LEARN:** Captures the full tonal profile for removal of complex hum. To learn in the app, make a selection containing the hum in isolation and click the Learn button. In the plug-in, enable learn and play back the section of audio containing the hum. If you cannot find a selection of hum in isolation, RX can analyze any audio with prominent hum, but the results may not be as useful.
- **ADAPTIVE MODE:** Allows De-hum to adjust the notch filters based on changes in the audio over time. In this mode, RX will analyze incoming audio to determine what is hum and what is desired audio material. Adaptive mode will work better with hum that changes in pitch throughout the file.

> **ADAPTIVE DYNAMIC MODE NOTE**
>
> Adaptive Dynamic mode in the De-Hum plug-in is not optimized for real time use. In the plug-in, please use this feature offline.

- **SENSITIVITY:** Adjusts the amount of hum that will be removed.
- **BANDS:** Controls the number of dynamic notch filters, which update their gains based on the input signal to avoid ringing. Adjust to higher values for complex hums, at the expense of transparency.
- **FILTER Q:** Controls the bandwidth of the notch filters.
- **RANGE SELECTORS:** Sets the action region of the hum removal by defining the upper and lower bounds of De-hum processing.

Controls in Static Mode



- **BASE FREQUENCY:** Sets the fundamental frequency of the hum to be removed. With Preview engaged, adjust the slider until you find the point where the hum is appropriately reduced.
  - **FREQUENCY (Hz):** Sets the fundamental frequency for the filter.
  - **FILTER Q:** Controls the bandwidth of the notch filters for the base frequency and all of the harmonics.

    During Preview, you can use the Spectrum Analyzer to help identify the base frequency of the hum.

- **LEARN:** Automatically set the Base Frequency based on the learned profile. To learn in the app, make a selection containing the hum in isolation, and click the Learn button. In the plug-in, enable learn and playback the section of audio containing the hum. If you cannot find a selection of hum in isolation, RX can analyze any audio with prominent hum, but the results may not be as useful.
- **ADAPTIVE MODE:** Allows De-hum to adjust the notch filters based on changes in the audio over time. In this mode, RX will analyze incoming audio to determine what is hum and what is desired audio material. Adaptive mode will work better with hum that changes in pitch throughout the file.

- **LINEAR-PHASE FILTERS:** Linear-phase enables FIR (Finite Impulse Response) filters with a high FFT (Fast Fourier Transform) size. These filters provide very accurate frequency response with no change in phase at the expense of latency and filter pre-ringing.

  DISABLING LINEAR PHASE (FIR) FILTERS

  When Linear Phase is disabled, De-hum will use minimum-phase IIR (Infinite Impulse Response) filters. These are also very accurate, and are only susceptible to post-ringing, which is usually less noticeable than the pre-ringing introduced by FIR filters.

  Latency Consideration: Disabling Linear Phase Filters will reduce the latency used by De-hum when it is being used as a real-time plug-in.

- **HIGH/LOW-PASS FILTERS:** These traditional filters come ahead of the De-hum notch filters, and allow for frequencies to pass above or below a certain cutoff point. These can be useful for tackling extreme hum or buzz.
  - **FREQUENCY (Hz):** sets the cutoff frequency for the filter
  - **FILTER Q:** Sets the bandwidth of the filter (or dB/octave cut). In the default IIR filter mode with a high Q setting, you may notice a resonance at the cutoff frequency characteristic of traditional analog filters. That resonance can be mitigated by engaging the Linear-phase filters.
- **HARMONICS:** Harmonics often accompany the fundamental frequency of a hum. With De-hum you can also attenuate these overtones with additional notch filters. Using the Harmonics control, you can add up to 16 harmonics. The spectrogram display makes it easy to identify the harmonics. After setting the number of harmonics, use the Slope control to set how aggressively the higher harmonics are attenuated.
- **LINK HARMONICS:** Connects the gain controls of the notch filters.
  - **ALL:** presents a single node on the display for controlling the gain of all the notch filters. This is the default setting.
  - **ODD/EVEN:** presents two nodes on the display, one for controlling the gain of the fundamental frequency and even harmonics, and another for controlling the 1st harmonic and any following odd harmonics.
  - **NONE:** presents individual gain nodes for the fundamental and each harmonic.
- **SLOPE:** When harmonics are linked, this controls the harmonic slope of the gain nodes for each overtone. As the harmonic order increases, the gain level resolves closer to 0 dB. When the Link Harmonics control is set to Odd/Even, a separate control appears that affords independent control over the slope for both odd and even harmonics.
- **HARMONIC PANEL [dB]:**



- **ENABLE:** This control allows you to choose which notch filters are active. Disabling unnecessary notches can help preserve the original sound.
- **GAIN [dB]:** This controls the level for each notch filter. You can also manually enter gain settings for the fundamental, or any of the harmonics if Link Harmonics is set to None.
- **FILTER DC OFFSET:** This checkbox will engage a filter to remove any DC (direct current) offset that sometimes occurs in A/D converters or analog circuits used in the recording process.
- **OUTPUT HUM ONLY:** Selecting this check box will isolate the hum that is being removed. This is useful for fine-tuning your settings. Identify a section of your file where the hum is mixed with other material, select this mode, and click Preview. Now adjust parameters like Filter Q and Slope control to maximize hum removal, thus minimizing the effect on the program material.

More Information

Ringing

Every notch filter can be characterized by its frequency response and its impulse response. The *frequency response* determines how the filter is changing amplitudes of different frequency components in the signal. The *impulse response* determines what the filter does to the signal waveform and how it handles transients.

Increasing the Q of a notch filter makes it narrower, but adds more “tails” to the impulse response of the filter. These tails are called *ringing* because they add a reverberant, ringing character to the perfect click of the impulse. This ringing matches in frequency with the bends in the filter’s frequency response.

Minimum Phase EQ approximates the behavior and sound of analog EQs and places all ringing behind (after) the signal transient, where it has a higher chance of being psychoacoustically masked by the signal itself. Since narrow notching in De-hum is likely to create significant ringing, minimum-phase filters are usually the preferred type when addressing ringing issues after notch filtering with De-hum.

Alternative Modules to use for Complex Hum Issues

The Dynamic mode of De-hum is a powerful tool for attenuating tonal noises prior to other denoising algorithms that target random components of the noise: Voice De-noise or Spectral De-noise. There are cases, however, when other modules may successfully replace or complement De-hum.

When there is no time to run De-hum and then follow up with De-noise (for hiss), just using RX Spectral De-noise can also take care of the hum rather efficiently. Spectral De-noise has separate controls for reduction of tonal and broadband components of the noise. A manually drawn reduction curve provides an additional way to fine-tune the amount of attenuation.

A quick fix for an occasional single-tone buzz is provided by Spectral Repair: just select the offending frequency on a spectrogram and apply the Attenuate Vertically mode with a reasonably high number of Bands (1024–4096).

Some other alternatives to De-hum:

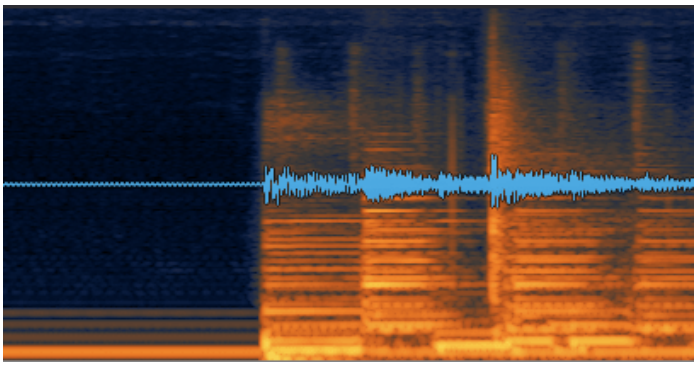
- Guitar De-noise For guitar tracks containing buzz, try using Guitar De-noise. This module can help reduce amp buzz along with string squeaks and picking that is too aggressive.
- Spectral De-noise For hum that has many harmonics that extend into higher frequencies (often described as "buzz"), try using Spectral De-noise. Spectral De-noise features tonal noise reduction controls that can make short work of harmonic hum and buzz across the entire spectrum.
- De-click When the buzz is very harmonically rich, sometimes using De-click in either *Single-band* or *Multiband (periodic clicks)* mode can treat the clicks that comprise buzz.

MORE INFORMATION ABOUT DYNAMIC MODE

For further information about RX De-hum and when to use it, refer to this article authored by iZotope Principal DSP Engineer Alexey Lukin.

Visual Example

This image shows the spectrogram of a file with 3 harmonics of a 60 Hz Hum:



27. De-plosive

Module & Plug-in

Overview

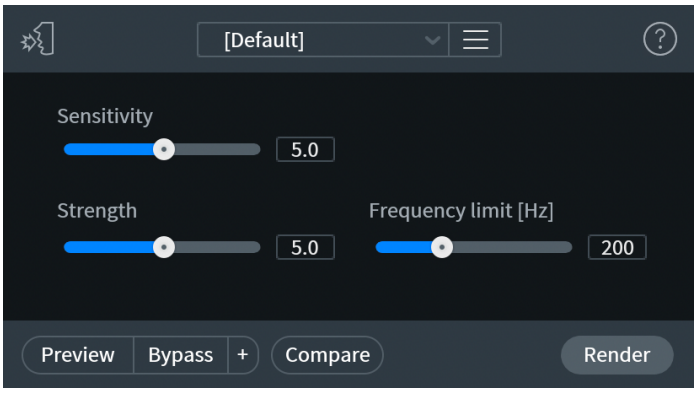
A *plosive* is a consonant speech sound that occurs when the vocal tract briefly blocks the flow of air during speech. When the block in the vocal tract is resolved, an audible release of pressure occurs, called a plosive. Pop filters (in a studio) or wind filters (on location) are commonly used when recording vocals or dialogue, in order to reduce the popping effect that can occur when the higher pressure plosive signal comes into contact with the diaphragm of a microphone.

De-plosive can intelligently identify, separate and reduce plosives in an input signal while still preserving the fundamental frequency content and harmonics of the dialogue.

Controls

The following section describes the controls available for refining plosive reduction.

> ⚠ **RECOMMENDATION: USE DE-PLOSIVE BEFORE APPLYING A HIGH PASS FILTER**
>
> De-plosive attempts to detect plosives between 20 to 80 Hz. If the input file has already been filtered with a high-pass filter, it is possible that plosive detection will not function correctly. To achieve the best results, it is recommended that De-plosive be used *before* applying a high-pass filter.



Sensitivity

Adjusts how much of the input audio will be categorized as a plosive. Sensitivity has a greater effect on the overall effectiveness of plosive reduction than simply increasing the Strength control.

- **Lower values** will instruct the algorithm to narrowly define what it will detect as a plosive. This may result in less overall reduction of plosives with the benefit of maintaining the quality of the speech signal.
- **Higher values** will instruct the algorithm to more broadly detect plosives in the input signal, which may result in too much of the speech signal being identified as, and subsequently reduced with, the plosive signal content.

Strength

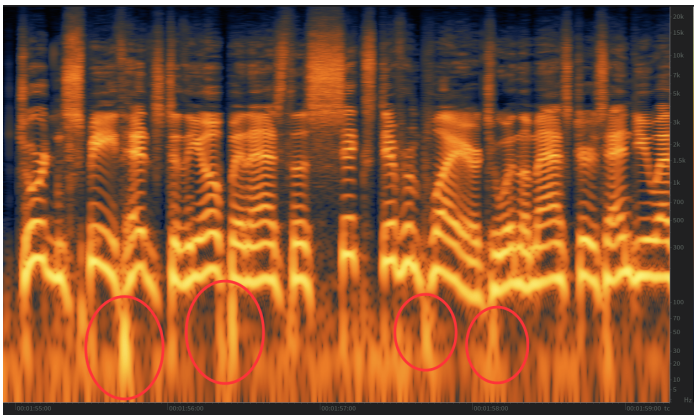
Adjusts the amount of reduction applied to the detected plosive signal content. Higher values can result in significant plosive reduction at the cost of decreasing the quality and clarity of the speech signal. To maintain the speech signal while still reducing plosives, try increasing the Sensitivity control and decreasing the Strength control.

Frequency Limit [Hz]

Sets the upper frequency boundary for plosive reduction. Plosives typically manifest between 20Hz to 300Hz, but may sometimes reach as high as 500Hz. Setting this control just above where the plosives are occurring in the input signal can help reduce unwanted reduction of frequency above where the plosives are occurring.

Identifying plosives in the spectrogram display

The spectrogram can be used to help determine the highest plosive frequency in the input signal. Plosives are caused by small bursts of air pressure hitting the diaphragm of a microphone which results in an increase in amplitude. The colors used to represent signals in the spectrogram are tied to the amplitude of a signal. Darker colors represent lower amplitudes and brighter, more vibrant colors represent higher amplitudes. Plosives will generally appear brighter than the signal content that surrounds them.



28. De-reverb

Module & Plug-in

Overview

De-reverb gives you control over the amount of ambient space captured in a recording. It can make large cathedrals sound like small halls and make roomy vocals sound like they were recorded in a treated space. De-reverb processes audio according to the reverberant/direct ratio (also known as wet/dry ratio) detected in the signal. It can learn your audio to suggest some frequency and decay time settings, or you can estimate these yourself.

Controls



- **LEARN:** Teaches De-reverb how much reverb is in your signal.
  - The Learn feature analyzes the signal and determines the wet/dry ratio per frequency of your signal, as well as the overall rate of decay of reverb.
  - When the Learn operation completes, the Reverb Profile and Tail Length controls will be set to their suggested values.
  - The Learn operation can be performed on any reverberant audio.

> **RECOMMENDATION**
>
> We strongly recommend using the Learn feature on a selection of audio that starts with some noise floor (or room tone), is several seconds long, and includes both the direct signal and its reverberant tails.

- **METERING**
  - The top meter shows a comparison between the input and output signal energy over the past five seconds of playback.
  - The bottom meter shows the amount of reverb reduction over time. It is the difference between input and output plotted on a flat line.
  - Both of the meters together give you an idea of what De-reverb considers reverb and help you refine your settings.
- **REDUCTION:** Controls the applied amount of De-reverb.
  - Larger amounts mean more reverb is removed.
  - Smaller amounts perform less processing.
  - This control represents the target wet/dry ratio for processing. In other words, if it is set very high, it will treat the signal as though it has more reverb and process it more.

> **NOTE: NEGATIVE REDUCTION VALUES**
>
> Negative Reduction values will increase the amount of reverb in the signal.

- **REVERB PROFILE:** Controls the amount of De-reverb effect applied per band.
  - These controls are set automatically by the Learn feature.
  - If a group of reverberant tones are more prominent in a signal, increase the control for that band.
  - Generally you want to set these controls to match the reverb originally present in your signal. For example, if reverb takes longer to decay (or is more present) in a particular band, set that control higher.
  - These controls can also be used to address more prominent ringing or resonant groups in signals. **For example,** increasing the profile control for low frequencies can remove muddiness from a resonant bass guitar, while increasing the high band control can curtail ringing sibilance in live vocal recordings.

- **TAIL LENGTH:** Controls the decay of De-reverb processing. This control is an approximation of RT-60, the rate of time it takes for a reverberant signal to decrease in amplitude by 60 dB. This is automatically set by the Learn feature.
  - **Increase** this control if reverb tails reappear after processing, or if early reflections are too apparent.
  - **Decrease** this control if reverb tails and noise floors sound over-processed, or if the processed audio sounds dull.
  - Setting this control to the minimum value is effective at processing early reflections.
- **ARTIFACT SMOOTHING:** Controls the frequency accuracy of De-reverb processing.

> **ARTIFACT SMOOTHING DEFAULT VALUE NOTE**
>
> Because reverb is generally smooth across the frequency spectrum, the default value of this control is very high. However, if you need more accuracy to address issues like **resonant tones in a room**, you can decrease this control. The tradeoff is generally more artifacts from strong processing, so you may have to balance adjusting this and the Reduction control.

- **ENHANCE DRY SIGNAL:** Increases the level of the direct signal.
  - Boosting the direct signal can help create more dynamic range in the signal, and is a good option to try when working with voice or transient material.
  - Enabling this option can also help prepare material for later de-noising.
- **OUTPUT REVERB ONLY:** Changes the output of De-reverb from the processed signal to the wet reverberant signal.
  - This is useful for monitoring the processing to get better results. Hearing just the reverb helps you understand the impact of controls like Reduction, Reverb Profile, Tail Length, and Artifact Smoothing.
  - When this option is enabled, the output may not sound much like reverb because it is the difference of processing against the original signal, with some enhancement to expose more of the reverb apparent in the recording.

More Information

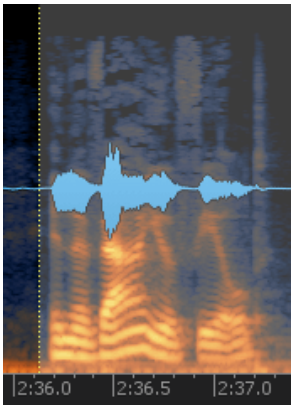
Early Reflections

Early reflections are the rapid echoes of a direct sound from a nearby surface. They are often distinct from the rest of a reverberant tail because they have a lot of energy but end quickly. Early reflections typically comprise the first 5 to 100 milliseconds of a reverb tail.

Using De-reverb as a real-time plug-in

- De-reverb is available as a VST/AU/RTAS/AAX real-time plug-in.
- However, due to the complexity of this processing, it can be resource intensive.
- To achieve high-quality results, it is always best to bring the audio file in question into RX Audio Editor (via RX Connect or by opening it directly), applying De-reverb, and then returning the file back to your original session.

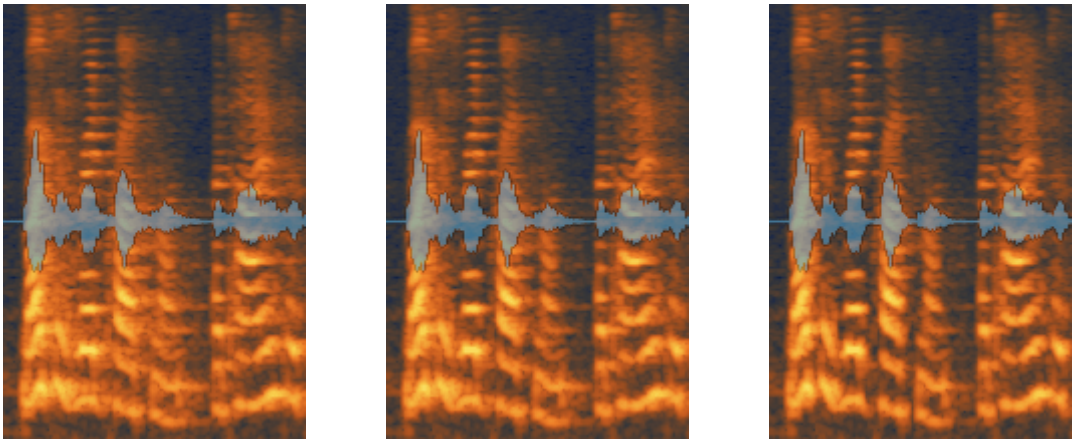
Tips for Learning a Reverb Profile



- To find the best settings for your signal quickly, find about five seconds of audio that starts with noise and has both direct signal and reverberant tails.
  - If you can find enough direct signal and reverberant tails to fill the De-reverb signal trace meter while using Learn, you will probably get a good reverb profile.
- Direct signal, reverberant tail, and noise are all important to help De-reverb understand your audio and set its controls appropriately. It needs to understand the ratio of dry signal to reverberant signal, how long reverb tails last, and where the noise floor of your signal is (to avoid excessive processing).
- If you have trouble getting good results from the Learn feature, you can try:
  - learning on transient broadband audio, like drums, claps, or coughs
  - learning on any audio that is obviously reverberant
  - learning for a longer amount of time. Most reverb can be analyzed in a few seconds, but some reverb profiles can require up to ten seconds of analysis.
- Does De-reverb processing sound unnatural?
  - If the output of De-reverb sounds unnatural after Learning, try slowly decreasing the reduction control.

Visual Example of De-reverb processing

De-reverb has the effect of sharpening a signal in time. You can see this transition in the spectrogram: reverberant audio looks blurred, and cleaned audio appears more focused. Here, a recording of a distant speaker (left) has had its long tails processed (center) before another De-reverb pass with shorter tail lengths to tackle the early reflections (right):



Tips for dealing with complicated reverbs

If the audio you are working with has a very complex reverb, such as a reverb with apparent early reflections, you may get better results after trying a few passes of De-reverb.

1. First, start by training the De-reverb and set the Reduction amount to a value that yields good results on the long reverberant tail.
2. After processing, Learn a new reverb profile and try reducing the level of early reflections: set the Tail Length control to 0.5, Artifact Smoothing around 3.0, and increase Reduction.
3. A combination of De-reverb and Spectral De-noise can be used to tame very reverberant signals. It does not matter whether you process with De-reverb or Spectral De-noise first.

29. De-rustle

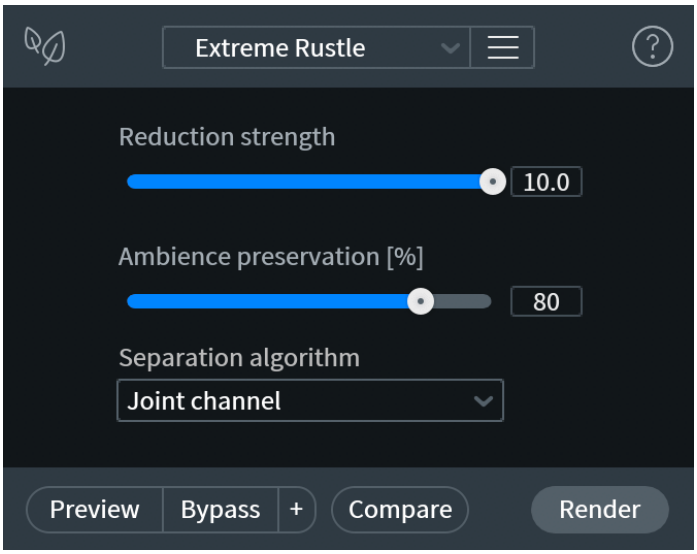
ADV | Module & Plug-in (Audiosuite Only)

Overview

De-rustle is designed to reduce the noise or "rustle" that is often the result of a lavalier microphone rubbing or brushing against clothing during a recording. This type of noise can vary unpredictably over time and exhibit a wide variety of sonic characteristics from higher frequency "brushing" to lower frequency "thuds," making it a challenging issue to resolve.

The De-rustle module uses a machine learning algorithm trained on isolated rustle samples, clean dialogue samples and dialogue samples with rustle. When processing, De-rustle leverages the trained data to identify and separate the rustle from the dialogue.

Controls



Reduction Strength

Adjusts how sensitive the separation algorithm is to identifying dialogue in the input signal.

- **Lower values** allow for broader detection of dialogue in the input signal. This can result in more rustle in the processed signal, but can increase clarity of dialogue in the processed output signal.
- **Higher values** allow for stricter detection of dialogue in the input signal. This can result in more significant reduction of rustle or lav noise, but can decrease the clarity of dialogue in the processed output signal.

Ambience Preservation

Adjusts the amount of background noise to retain in the processed output. The separation algorithm may include background noise and ambience in the signal it has detected as rustle. This can result in unwanted reduction of unrelated background noise in the processed signal.

- **Lower values** allow for broader detection and separation of rustle against other background noise in the input signal. This can result in *more* background noise being included and reduced with the separated rustle component.
- **Higher values** allow for stricter detection and separation of rustle against other background noise in the input signal. This can result in *less* background noise being included and reduced with the separated rustle component.

Separation Algorithm

The following separation algorithm modes are available in the De-rustle module.

Channel Independent

When this mode is selected, the separation algorithm is applied to the input audio channels independently. Channel independent mode is the fastest Separation algorithm option. It offers the most efficient real-time preview performance and processing speeds when working with the De-rustle module in the RX Audio Editor.

Joint Channel

When this mode is selected, joint channel processing is applied to the input audio before determining the separation between dialogue and rustle signal components. Joint Channel mode offers higher quality separation results than Channel Independent mode, especially when processing stereo files with similar content on both channels (correlated signals, strong stereo image).

Advanced Joint Channel

When this mode is selected, joint channel and additional advanced processing is applied to the input audio before determining separation between dialogue and rustle. Advanced Joint Channel mode offers the highest quality separation results, especially when processing files with high sampling rates. This mode requires longer processing times than the other two modes. If processing time is a concern, Channel Independent mode can be used as a faster, lower quality alternative.

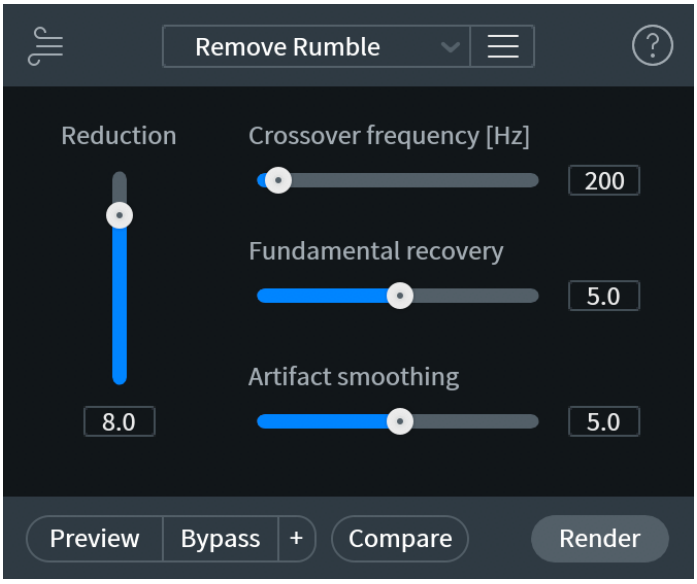
30. De-wind

ADV

Overview

De-wind removes intermittent low frequency rumble that can occur when light to moderate bursts of wind come into contact with a microphone diaphragm. De-wind is not designed to remove heavy wind bursts that blow out and distort a microphone signal.

Controls



Reduction

Determines the balance between the depth of wind reduction and the preservation of the original signal.

Crossover Frequency

Sets the upper frequency limit for the De-wind processing algorithm.

Fundamental Recovery

Re-synthesizes lower voice harmonics that may be lost or obscured by wind.

Artifact Smoothing

Eliminates “musical noise” that is often characteristic of FFT-based processing. Musical noise can be described as how something may sound underwater. Increase this slider if your output sounds watery, but decrease it when too much smoothing makes speech sound muffled.

WHAT IS AN FFT?

Fast Fourier Transform: a procedure for the calculation of a signal frequency spectrum. The greater the FFT size, the greater the frequency resolution, i.e., notes and tonal events will be clearer at larger sizes. However, when using FFT-based processing, the more audio you remove from your source, the more likely you are to create undesirable artifacts.

Tips

When to use the De-wind Module

The De-wind module is best suited to remove and reduce **intermittent** wind noise in the foreground (for example, a recording with periodic wind gusts that come into direct contact with the diaphragm of the microphone) as opposed to constant wind noise in the background.

The De-wind algorithm attempts to preserve the noise floor of a recording by tracking how it changes over time. Specifically, De-wind is looking for noise floor changes that are characteristic of the noise introduced when a burst of wind comes in contact with the diaphragm of a microphone. If the wind in your recording is constant or in the background, De-wind will treat it as part of the desirable noise floor and will preserve it rather than remove it.

Alternative Modules

For recordings with constant or background wind noise, you may achieve better results by:

- Using the Spectral De-noise module.
- Running De-wind processing more than once.

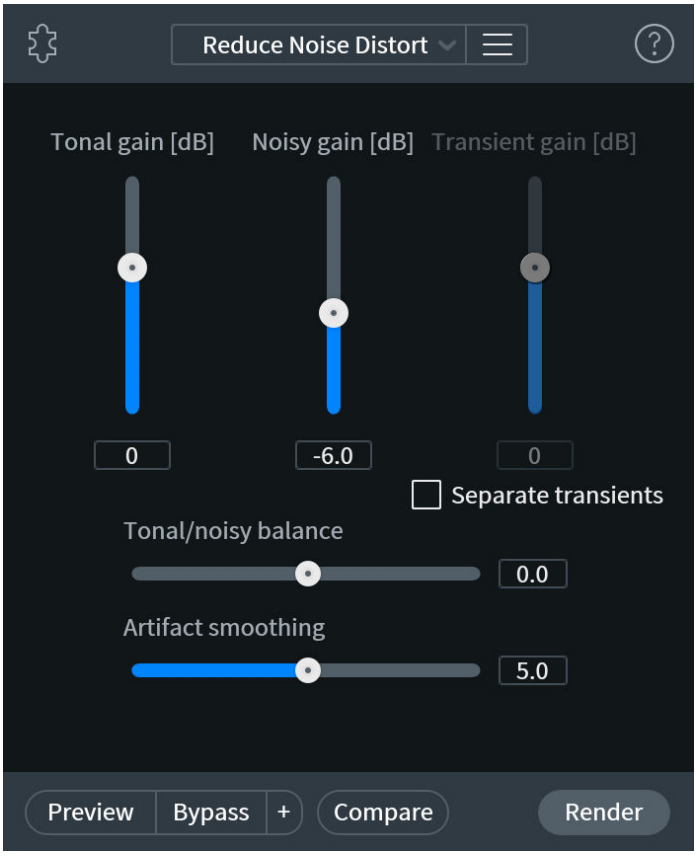
31. Deconstruct

ADV

Overview

Deconstruct analyzes your audio selection and separates the signal into Tonal, Noisy, and Transient (optionally) audio components. The separate components of the signal can then be cut or boosted individually using their associated Gain control.

Controls



Tonal Gain

Adjusts the level of the tonal components of the signal. Boosting a tonal signal (voice or instrumental) can help lift it out of a noise floor.

Noisy Gain

Adjusts the level of noisy components of the signal. This can be very useful for highlighting areas of raspiness or distortion only, and then attenuating the noisy gain to reduce overall distortion.

Separate Transients

Enables transient separation processing and activates the **Transient Gain** control.

Transient Gain

Adjusts the level of transient components of the signal. This can work as a transient shaper or a declicker, allowing you to attenuate or boost clicks and attacks.

> **TRANSIENT SEPARATION PERFORMANCE**
>
> The Separate Transients option allows you to control the level of transients, but increases CPU load. This may impact the performance of Preview. In this case, using the Compare functionality instead of Preview is recommended.

Tonal/Noisy Balance

Modifies the default weighting of the separation algorithm used by Deconstruct to categorize components of a signal as either "noisy" or "tonal."

- Negative values (*Tonal weighting*) will classify more of the "noisy" components of a signal as "tonal" components and apply Tonal Gain to them during processing.
- Positive values (*Noisy weighting*) will classify more of the "tonal" components of a signal as "noisy" components and apply Noisy Gain to them during processing.

Artifact Smoothing

Reduces "musical noise" artifacts that are often characteristic of FFT-based processing. Fast Fourier Transform (FFT) is a procedure for the calculation of a signal frequency spectrum. The greater the FFT size, the greater the frequency resolution, i.e. notes and tonal events will be clearer at larger sizes. However, when using FFT-based processing, the more audio you remove from your source, the more likely you are to create undesirable artifacts. Increase this slider if Deconstruct's output sounds watery, but decrease it when too much smoothing reduces the separation between signal components.

More Information

- Deconstruct can be useful for a variety of audio files and applications, particularly when attempting to remove noise that varies throughout the length of a file.
- Deconstruct differs from the Spectral De-noise and Voice De-noise modules, which separate signal from noise based purely on amplitude. Deconstruct analyzes the harmonic structure of a signal independently of level. It does not matter if a tonal signal like hum is quiet or prominent. Deconstruct will treat it as a tonal component and adjust its gain accordingly.
- Deconstruct can be effective in removing residual vinyl noise that may be present after applying De-click or De-crackle processing. Using Deconstruct in this situation may produce better results than using the Spectral De-noise or Voice De-noise modules.

32. Dialogue Contour

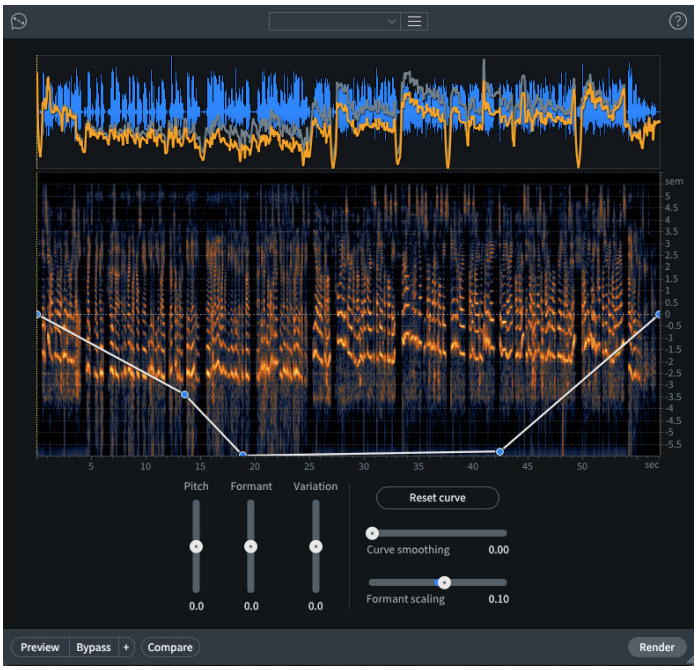
ADV

Overview

Dialogue Contour allows for the manipulation of the pitch envelope of a dialogue selection. It features pitch correction processing that is tailored to speech. It is useful for adjusting the inflection and/or expressiveness of words that may not flow or fit with the rest of the dialogue in a clip. It can also be used to make a series of dialogue edits sound more cohesive.

Displays

Dialogue Contour features a waveform panel and a spectrogram panel that each display information about the current selection in the active file tab. These panels will dynamically update when the selection is changed. If no selection is made in the active file tab, no information will be displayed in the spectrogram or waveform panels.

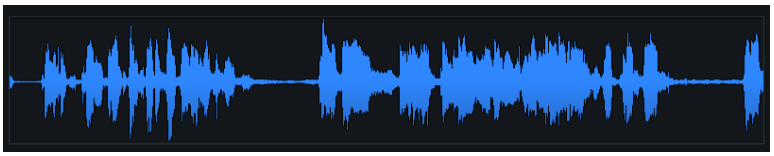


WINDOW RESIZING

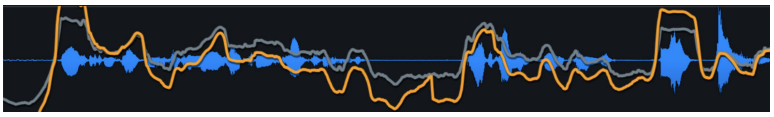
Click and drag on the bottom right-hand corner of the module window to customize the window size.

Waveform and Pitch Curve Display

The single waveform drawn in this panel represents a sum of all enabled channels in the current selection. The waveform drawing is normalized to allow for consistent vertical resolution when working with selections of varying amplitude.

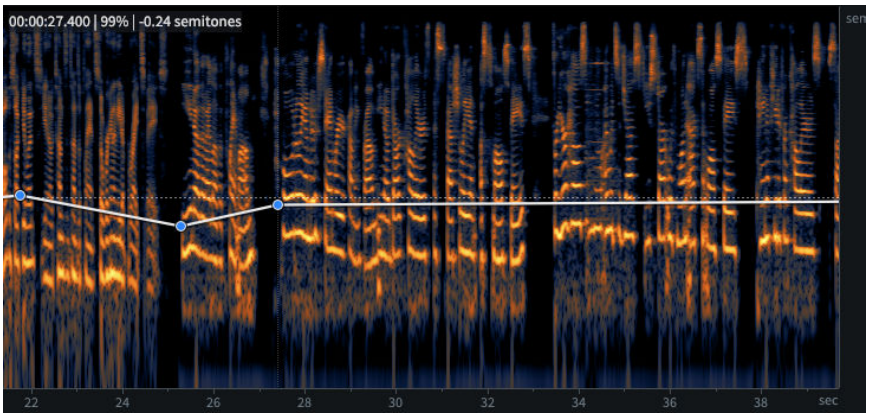


Dialogue Contour now has two curves displayed over the waveform view. The grayed out curve is the original pitch of the selected audio. The orange curve is the projected pitch based on parameters set by the user and will update to reflect any subsequent changes.



Spectrogram Display

The spectrogram drawn in this panel represents a sum of all enabled channels in the current selection.

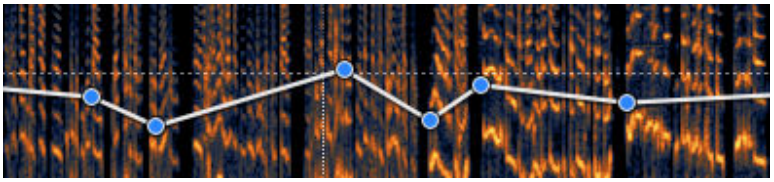


Playhead Indicators

The solid white vertical line and dotted yellow vertical line overlaid on the waveform and spectrogram panels indicate the current playhead position (white) and the playhead anchor position (yellow).

Contour Curve

The white line overlaid on the spectrogram panel represents the pitch contour curve. Nodes can be added to this curve and adjusted to make changes to pitch over the course of the active selection.



Contour Curve Axes

The contour curve allows for adjustments along two axes: Pitch and Time.

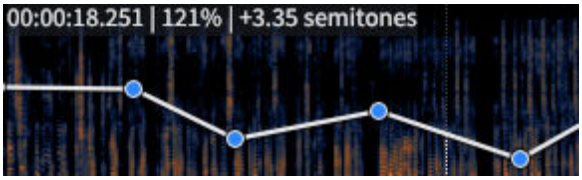
- **PITCH**: The vertical y-axis of the contour curve represents pitch in semitones.
  - The Pitch axis ranges from -6 (bottom) to +6 (top) semitones.
  - The center of the Pitch axis equates to 0 semitones.
- **TIME**: The horizontal x-axis represents time.
  - The time format used here is determined by the time format display selection in the transport section of the main editor window.
  - The range of the time ruler matches the length of the current selection.

To zoom the ruler in and out:

1. Hover over the ruler and use a mousewheel or trackpad to zoom in and out.
2. Click and drag left or right on the ruler when zoomed in to change the ruler position.
3. Double-click on a ruler display to reset the zoom level to default.

Contour Curve Readout

When the cursor is positioned over the spectrogram panel, a text readout will appear in the upper left hand corner of the panel. This readout displays information about the processing that will be applied when the contour curve is rendered.



The readout displays the following information about the cursor position, from left to right:

- **TIME**: Current time position of the cursor within the spectrogram panel.
- **PITCH SHIFT (%)**: Percentage of pitch shift that will be applied at the cursor's current time position.
- **PITCH SHIFT (Semitones)**: Amount of pitch shift that will be applied at the cursor's current time position.

Contour Curve Editing

The following section describes the methods and controls available for editing the contour curve.

Add Nodes

Click in the spectrogram panel to add a new node to the contour curve.

> ⚠ **CONTOUR CURVE NODE LIMIT**
> The contour curve supports adding up to 25 nodes.

Semitone Adjustments

Click and drag a node up or down to adjust its semitone value.

Time Adjustments

Click and drag a node left or right to move it earlier or later in time.

- Nodes cannot be moved outside of the time bounds of the current selection.
- The contour curve shape will be maintained when the selection changes.
- The contour curve shape will be maintained after rendering.

Remove Nodes

Individual nodes can be deleted from the curve using the following methods:

▶ Click and drag a node past the top or bottom edge of the contour curve display to quickly remove it from the curve.

▶ Control-click (Mac) or right click (Windows) on a node to remove it from the curve.

Reset Individual Nodes

Double-click or Alt-click (Option-click) on a node to reset it to default (0 semitones).

Controls

The following section describes the controls available for refining formant scaling and applying a global pitch offset to the entire selection.

Pitch

Applies a global semitone offset value to the current selection. This value is added to or subtracted from the processing applied by the contour curve. Adjusting the Pitch amount will not update the contour curve display.



The Pitch slider can be useful for adjusting a single word by a static amount. For example, setting Pitch to +2 semitones with the contour curve set to default will shift the pitch of the current selection up by 2 semitones.

Formant

Applies a global formant shift (in semitones) to the current selection. Formants are resonances of a vocal tract that define the timbre of a voice. They are independent of pitch, so adjusting formants will not update the pitch curve display.

Shifting formants up makes the voice sound thinner, while shifting them down makes the voice sound heavier and bigger.

Variation

Changes the amount of expressivity in a voice by scaling the existing pitch variations. When the Variation is set to -100%, the pitch becomes flat (monotone). 0% corresponds to no modification. +100% corresponds to 2X scaling of existing variations for maximum expressivity and exaggerated pitch.

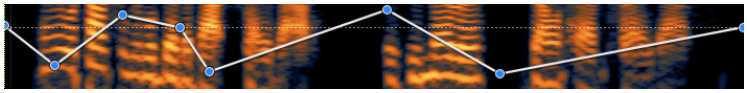
Reset Curve

Removes all custom nodes from the curve, resetting it to default. Two nodes are present in the default curve, one at the start and one at the end of the current selection. The default nodes are set to 0 semitones (no pitch adjustment).

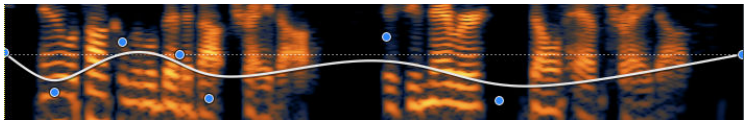
Curve Smoothing

Adjusts the amount of smoothing applied between nodes on the contour curve. Smoothing is a global control and is applied to all nodes on the curve.

- Lower smoothing values: Applies little to no smoothing between nodes on the curve. Allows for stricter transitions between nodes.



- Higher smoothing values: Applies more smoothing between nodes on the curve, resulting in a gradual, rounded slope between points on the curve. Allows for more gradual transitions between nodes.



Formant Scaling

Adjusts the amount of formant shift applied when adjusting pitch. The formants are shifted together with pitch in the same or opposite direction. When the Formant Scaling is 0, the formants stay in place. When it is 1, the formants move together with the pitch shift, i.e. the pitch shift does not preserve formants.

This control can be helpful for maintaining or correcting the timbre and quality of the dialogue after processing. In some situations, the formants may sound unnaturally high or low after processing. Formant scaling can be used to correct for these unnatural sounding results.

Alternative Modules

For more generalized, non-dialogue specific pitch envelope editing, try using the Variable Pitch module. The Variable Pitch module features the ability to shift pitch over the course of a selection with or without preserving timing and is suitable for use on a wide range of input material.

33. Dialogue Isolate

STD & ADV | Module & Plug-in

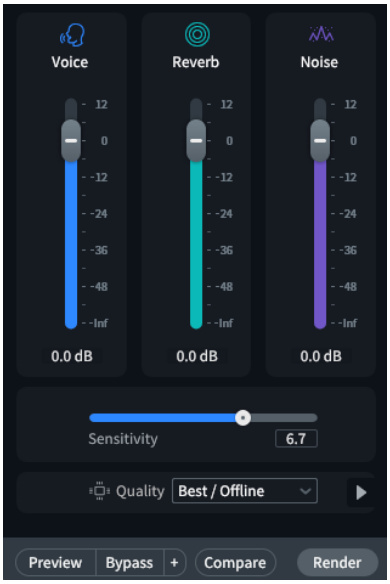
Overview

Dialogue Isolate is designed to separate spoken dialogue from a constant or non-stationary background noise, such as hiss, crowds, traffic, footsteps, weather, or other noise with highly variable characteristics. It can be particularly effective at increasing the level of dialogue in challenging low signal-to-noise ratio recordings.

New in RX 11, Dialogue Isolate now includes reverb removal technology, which separates dialogue from the characteristic echo and reverberation generated by different spaces, such as a tiled room, a lecture hall, or a cathedral.

Dialogue Isolate uses a machine learning algorithm trained on a large library of speech, noise, and reverberant data. It automatically detects and separates dialogue, noise, and reverberation into three distinct signal components. The levels of the separated components can be independently adjusted using the Voice, Reverb, and Noise gain controls.

Controls (STD and ADV)



Voice

Adjusts the amount of gain (in dB) applied to the components identified as voice.

Reverb

Adjusts the amount of gain (in dB) applied to the components identified as reverb.

Noise

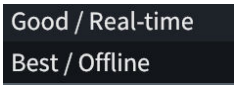
Adjusts the amount of gain (in dB) applied to the components identified as noise.

Sensitivity

Determines how much of the noise and reverb will be extracted from dialogue by the separation algorithm.

- **Higher values** will instruct the separation algorithm to broadly define what it categorizes as noise and reverb in the input signal. This can result in more significant noise reduction, at the cost of introducing artifacts and potentially reducing dialogue clarity.
- **Lower values** will instruct the separation algorithm to narrowly define what it categorizes as noise and reverb in the input signal. This can result in more noise being included in the rendered signal, but can help to maintain dialogue clarity.

Quality



The following quality modes are available in the Dialogue Isolate module.

- **Good / Real-time:** A fast, high-quality processing mode that is optimized for speed (low-latency) and for use in the plugin.
- **Best / Offline:** Our highest-quality processing mode with fewer artifacts and better isolation at the cost of a longer processing time (higher latency). For use in the RX Audio Editor and for offline bounces/renders in DAWs. (Advanced only)

> The RX 10 Dialogue Isolate legacy processing modes do not exist in RX 11.

Controls (ADV only)

The following controls are only available in the Advanced version of the Dialogue Isolate module and plugin (ADV):



An expand button shows or hides the metering and multi-band interface by making the window larger or smaller.

If the reduction of any band is set to a value other than 100%, the expand button will be highlighted in blue to indicate that the multi-band controls will have an effect on the output.

Multi-band Interface

The multi-band interface in the Advanced module and the plugin adjusts the amount of processing across four different frequency bands (Low, Low Mid, High Mid, High). Each band is individually set from 0 to 100% of the processing parameter. (For example, if the Noise gain slider is set to -10dB, and a band is set to 0%, the gain of the detected noise in that band will be 0dB. If the band is set to 100%, the gain of the detected noise in that band will be -10dB.) The three crossover points are adjustable.

Difference metering provides visual feedback as to how much gain adjustment is happening across the four bands.

Alternative Modules

For stationary noise, such as hiss, buzz, line noise, etc., Dialogue Isolate should produce good results, but we also suggest trying the Spectral De-noise module if you are having trouble achieving acceptable results with Dialogue Isolate. Spectral De-noise is not focused on dialogue and can be used for a variety of source signals, like singing voice, instrumental tracks, or full mixes.

For more general, non-dialogue specific reverb reduction, try using the De-reverb module.

34. Guitar De-noise

Module & Plug-in

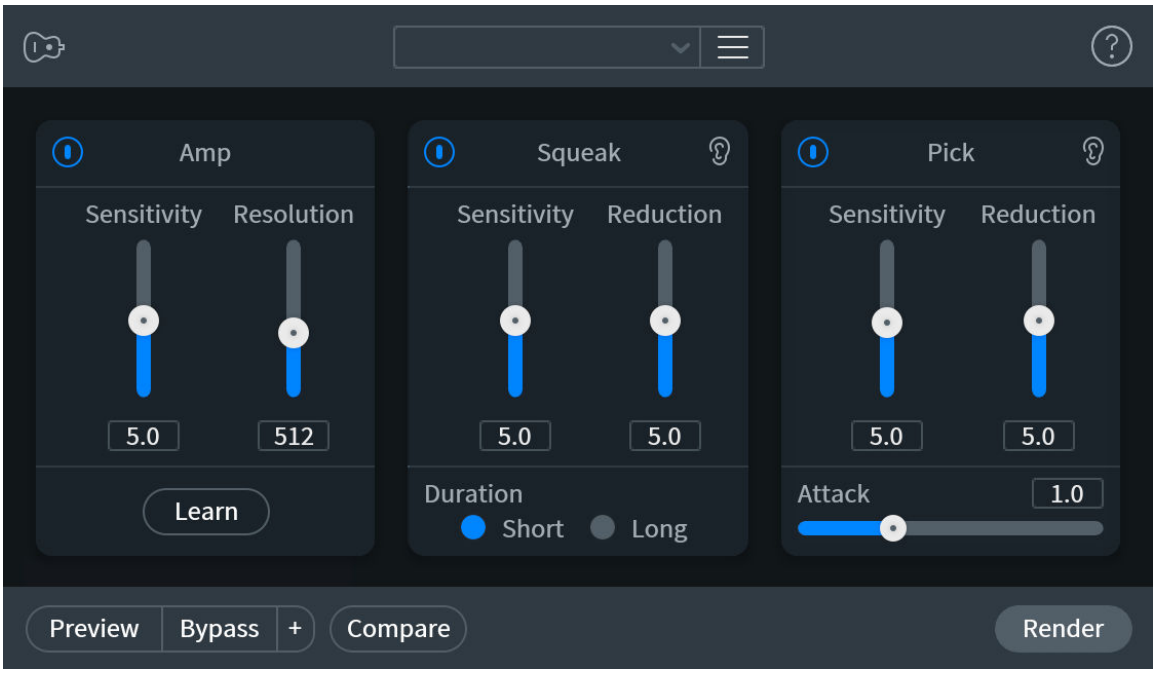
Overview

Guitar De-noise is designed to provide control over acoustic and electric guitar performance-related noises such as squeaks, picking sounds, and the hum or buzz that comes from guitar pickups or amplifiers. With Guitar De-noise, you'll be able to easily adjust these noises or eliminate them completely, depending on the needs of the production.

Guitar De-noise consists of three sections, each targeting its own type of artifact. The Amp section eliminates hum or buzz of amplifiers. The Squeak section attenuates fret noises. The Pick section tames excessively sharp attacks of plucked strings.

Controls

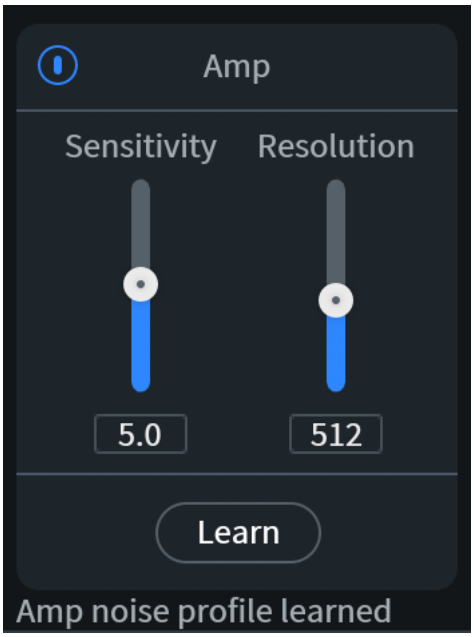
The Guitar De-noise module is divided into three main sections: Amp , Squeak , and Pick.



Amp

Amp section is designed to deal with the hum or buzz that comes from guitar pickups or amplifiers. Unlike the De-hum module, it can handle noises consisting of hundreds of harmonics that reach well into the high-frequency range. It can also deal with buzzes that do not consist of a single harmonic series: the Learn feature captures a snapshot of tonal noises at the arbitrary set of frequencies.

The noise for the Amp section needs to be tonal and static. If noise contains both tonal and broadband components (like hum + hiss), the broadband components can be addressed with a separate pass of Spectral De-noise or Voice De-noise. Static noise means that the frequencies of hum or buzz do not change in time.



- **LEARN:** will intelligently identify the noise profile of the hum or buzz in your audio
  - In the **plug-in**:
    - Loop a section of audio that contains only the hum or buzz that you want to reduce.
    - Click Learn to start learning.
    - After a moment click Learn again to finish the process.
  - In the **RX Application**:
    - Select the section of audio that contains only the hum or buzz that you want to reduce.
    - Click Learn.
- **Sensitivity:** determines how much hum or buzz will be removed from the signal. Use higher sensitivity values when the amplitude of noise varies in time: it will allow the algorithm to attenuate noises that are louder than the learned snapshot. If the noise level is constant, using lower sensitivity will allow for better preservation of the tone of the guitar.
- **Resolution**: sets the maximum number of harmonics that will be removed from the amp noise. The default of 128 will be enough for most cases, but greater resolution is available at the cost of increased CPU usage.

Squeak

The Squeak section is for controlling string noises (fret/finger squeaks) that occur when sliding up or down the guitar neck between notes or chords.



- **Sensitivity:** determines how many string squeaks are detected in the signal. Increasing sensitivity to higher amounts can attenuate some high-pitched guitar notes.
- **Reduction:** adjusts the amount of attenuation that is applied to the detected squeaks in the signal.
- **Duration**:
  - **Short** is good for the majority of shorter squeaks up to 200 ms in length.
  - **Long** can handle both short and long squeaks up to 1000 ms in length. This mode has high latency.

Pick

The Pick section allows you to reduce the attack of picked or plucked guitar strings for situations where they might stand out too much in music or production sound.



- **Sensitivity:** determines how picks or plucks are detected in the signal. Increasing sensitivity to higher amounts will impact the tone.

- **Reduction:** adjusts the amount of attenuation that is applied to the detected picks in the signal.
- **Attack:** determines how quickly the reduction is applied by altering the compressor's attack time.

Processing Tips

- Before learning the noise profile in the Amp section, identify and select the longest section (ideally a few seconds in length) of the recording that contains only the hum or buzz you wish to remove or reduce.
- Use the **Ear** button available on each section to set the sensitivity so that only the desired noises are removed. This will help to preserve the original tone.
- Use the reduction amount sparingly at first and increase it to higher settings to remove more of the noise if needed.

Alternative Modules

While the **Amp** section of Guitar De-noise successfully attenuates tonal noises, like hum or buzz, it may not work well on broadband noises like hiss or rumble. We suggest trying RX's Spectral De-noise or Voice De-noise modules for efficient attenuation of broadband noises. Voice De-noise has been specifically designed to provide high efficiency, zero-latency adaptive noise removal when inserted on a track in your DAW or NLE. The Spectral De-noise plug-in is far more resource intensive and has higher latency.

Although the **Squeak** section of Guitar De-noise is able to automatically identify and remove most of the fret noises, sometimes a manual correction of a few outstanding squeaks is required. In such cases, running Squeak with higher Sensitivity on a few isolated selections may be helpful. An alternative way employs RX's Spectral Repair module for precise editing of the squeaks in manual selections.

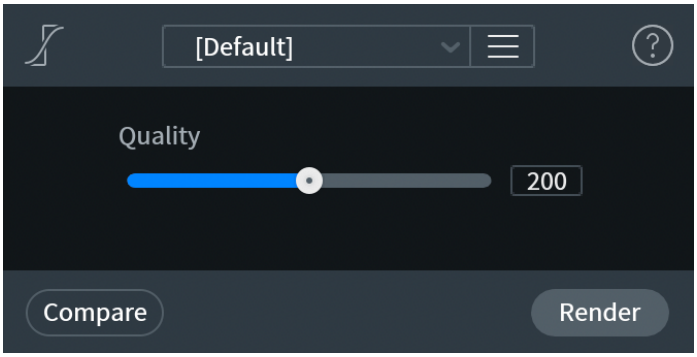
When the flexibility of the **Pick** section is not enough, RX's modules like De-ess (especially, the Spectral algorithm) or De-click provide alternative automatic ways of attenuating sharp attacks. Of course, in the most difficult situations, manual tools like Spectral Repair or Gain provide the ultimate editing precision.

35. Interpolate

Overview

The Interpolate module is used for repairing individual clicks below 4000 samples in length. This module corrects clicks by synthesizing a replacement signal based on the content in your selection. It can be used as a substitute for the 'pencil' editing tool found in many audio editing applications.

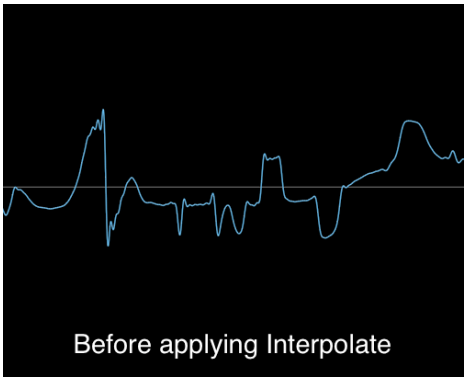
Controls

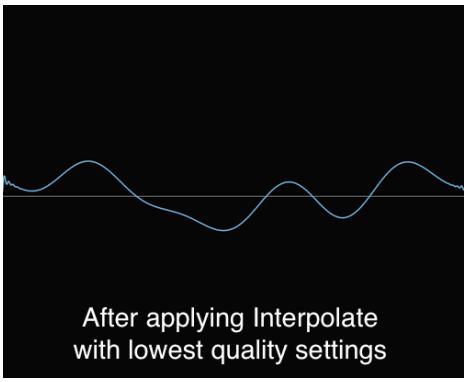


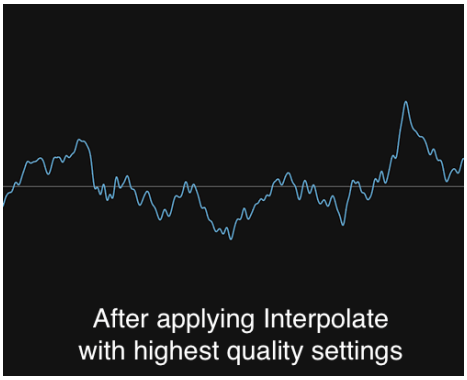
Quality

Adjusts the complexity of the sythesized replacement signal by adjusting the interpolation order. This allows you to tailor the results of the processing to better fit the surrounding audio.

The following images illustrate the effect of Interpolate processing on a selection that is 214 samples in length. The images are zoomed in to illustrate the effect of adjusting the Quality parameter.




After applying Interpolate
with lowest quality settings


After applying Interpolate
with highest quality settings

36. Loudness Optimize

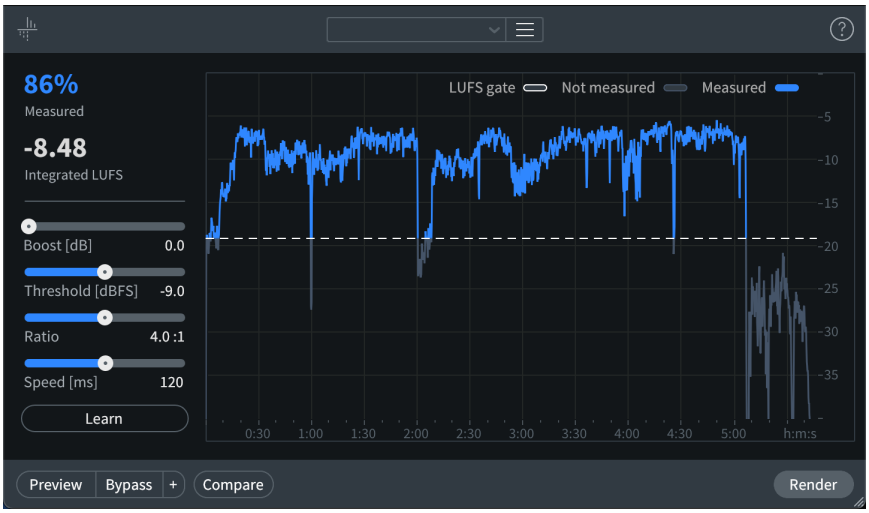
Overview

The Loudness Optimize module gives you information and control over your audio's loudness measurement and the relative gate in the BS.1770 loudness algorithm. When momentary loudness falls below this gate, it does not contribute to the file's overall integrated loudness measurement.

Applying upward compression can bring quiet content above the gate, causing it to contribute to the measurement and reduce your audio's integrated loudness. As a result, your audio can playback at a higher level on loudness-normalized streaming platforms.

There may also be cases where your audio's momentary loudness is already above the gate, and in this case we hope that Loudness Optimize will be a valuable check to ensure that your audio will sound optimally loud on streaming platforms.

Metering and Readouts



- **LOUDNESS TRACE:** Represents the momentary loudness of your audio over time. When audio is above the LUFS Gate and contributes to the integrated loudness measurement, it is colored blue. Audio below the LUFS Gate that does not contribute to the measurement is colored gray. This chart can be scrolled and zoomed by dragging and using the mouse wheel on either axis. The chart will focus on whatever region in time is highlighted in the main RX window. You can also hover over the chart to see the audio's momentary loudness at any point in time.
- **LUFS GATE LINE:** The dashed white line represents the relative gate in the BS.1770 Integrated Loudness measurement. This line will re-position dynamically depending on the audio and cannot be directly adjusted by the user. Its position is 10 LU below the average momentary loudness measurement.
- **MEASURED %:** Displays the amount of audio in the file that is above the LUFS Gate. Low Measured % values indicate that there is potential for loudness optimization because there is significant content that can be raised above the gate and therefore reduce integrated loudness. As the Measured % approaches 100%, it is less likely that the audio's integrated loudness measurement can be reduced.

- **INTEGRATED LUFS:** Displays the integrated loudness of the audio file after applying the LUFS gate. This is the measurement that streaming platforms will use when calculating loudness normalization.

Controls

Boost

Controls the maximum amount of gain boost in dB that the processing will apply to quiet content.

Threshold

Determines the level in dBFS where the upward compression will begin to be applied. Audio above this threshold level will be completely unaffected, so the audio's original peak values and limiting will be maintained. Audio below this threshold will be upward-compressed with the specified Ratio, subject to the maximum Boost. Threshold is also displayed in the main RX window (when the waveform is visible) and can be adjusted via the connected handles.

Ratio

Affects how much compression is applied below the threshold. It works similarly to a typical downward compressor's ratio control, but rather than affecting how much gain is reduced above a threshold it affects how much gain is boosted below a threshold.

For example, a Ratio of 4:1 on a typical downward compressor means that for every 4 dB that the signal exceeds the Threshold, the output gain will only be 1 dB above the Threshold. Applying Loudness Optimize with a 4:1 Ratio means that for every 4 dB the signal falls below the Threshold, the output will only be 1 dB below the Threshold.

Speed

Sets the ballistics of the upward compression, in milliseconds. It primarily determines how quickly the processor will begin to boost when audio falls below the threshold. It also sets the lookahead time for the upward compression. Longer speeds will set more lookahead time and be less prone to distortion.

Learn

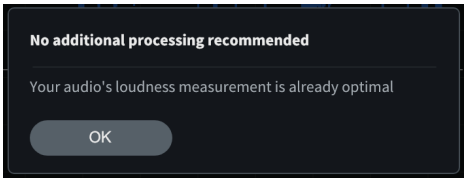
The Learn function will quickly apply many combinations of Boost and Threshold, measure the resulting integrated loudness, and select a Boost and Threshold amount that will result in the most optimal reduction in Integrated Loudness.

Learn will not change Ratio and Speed. Instead, your selected Ratio and Speed controls will be accounted for as Learn is running. The following workflow is recommended:

1. Set your desired **Ratio** and **Speed**. The default Ratio and Speed of 4:1 and 120 ms are a good starting point (achieve higher LUFS reduction with higher **Ratios** and faster **Speeds**).
2. Click **Learn**.

Learn will attempt to reduce the loudness of the highlighted audio region. In almost all cases, it is recommended to run Learn on the entire audio file. Even if you only intend to apply the Loudness Optimize processing to a specific section, the recommended workflow is to Learn from the entire audio file and then highlight and render the learned settings to the specific section.

It is possible that no reduction in integrated loudness can be achieved with any combinations of Boost and Threshold, given the selected Ratio and Speed. In this case, the Loudness Optimize module will show this message after running Learn:



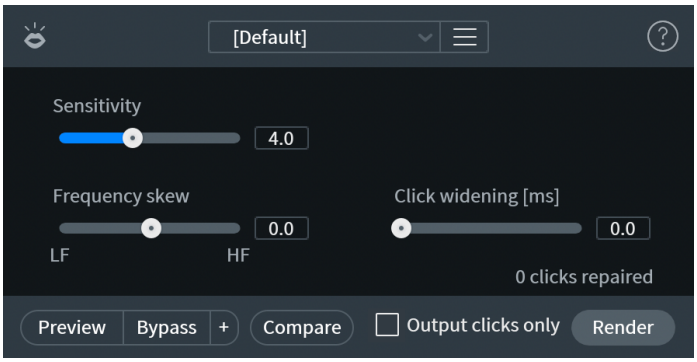
37. Mouth De-click

Module & Plug-in

Overview

Mouth De-click detects and reduces mouth noises such as clicks and lip smacks. It's designed for use on longer audio selections, but it can also be used to remove individual clicks.

Controls



- **SENSITIVITY:** Determines how many mouth clicks are detected in the signal. Increasing sensitivity can impact plosives, reducing or damaging the original signal.
- **FREQUENCY SKEW:** Targets the detection and removal of clicks to lower or higher frequencies. Negative values are more suitable for generic clicks such as those found on vinyl recordings. A setting of zero or above targets mouth clicks in the middle frequencies.
- **CLICK WIDENING:** Click Widening extends the repair area around detected clicks, compensating for mouth sounds such as lip smacks that have a decay.

More Information

Running Mouth De-click twice sometimes produces a better result than a single pass. This is because an initial run can miss the quieter clicks that were masked by louder clicks.

38. Music Rebalance

Module & Plug-in

Overview

Music Rebalance leverages a machine learning algorithm trained to identify and separate the following elements in a mix: Vocal, Bass, and Drums. Any content that is not otherwise identified as Vocal, Bass, or Drums will be categorized as Other. This typically includes all the melodic instruments. The level of each mix element can be independently adjusted after separation.

Music Rebalance can be useful for adjusting the level of a particular mix element when the original tracks or stems are not available as an alternative solution. In some cases, it can also be used to isolate a single mix element (e.g. the lead vocal) by reducing the level of the other three mix elements. It is available as a module in the RX 11 Audio Editor, a VST-3 and AU ARA plug-in, and as an Audiosuite and AAX real time plug-in in Pro Tools.

Controls



- **VOCAL:** Adjusts the level (in dB) of vocal signals.
- **BASS:** Adjusts the level (in dB) of bass signals.
- **DRUMS:** Adjusts the level (in dB) of drums and percussion signals.
- **OTHER:** Adjusts the level (in dB) of melodic instruments.
- **MUTE BUTTON:** Beside each of the separation sliders is a mute button for removing the separated stem from what is being auditioned.
- **SOLO BUTTON:** Beside each of the separation sliders is a solo button for auditioning only that stem.
- **QUALITY:** There are three separation algorithms to choose from in Music Rebalance: Good, Better, and Best. Good provides separation results the fastest. Better improves separation quality with only a minor decrease in processing speed. Best outputs the highest separation quality, including anti-phasing processing, with the slowest speed.

- **SENSITIVITY:** Individually adjusts the attention of the stem splitting algorithm toward a particular stem. When Sensitivity is higher, more of the subtle signals will get classified as this stem. For example, increasing the Sensitivity of vocals will preserve more of the breath sounds in the vocal stem. However higher sensitivity may also cause some leakage of other instruments into this stem. Note: raising the Sensitivity of all stems to 100% simultaneously does not change the results of stem separation (as the algorithm cannot be more attentive to everything) - Sensitivity only operates on relative values.

The RX 11 Music Rebalance Audiosuite plug-in **does not** include the ability to preview processing before rendering.

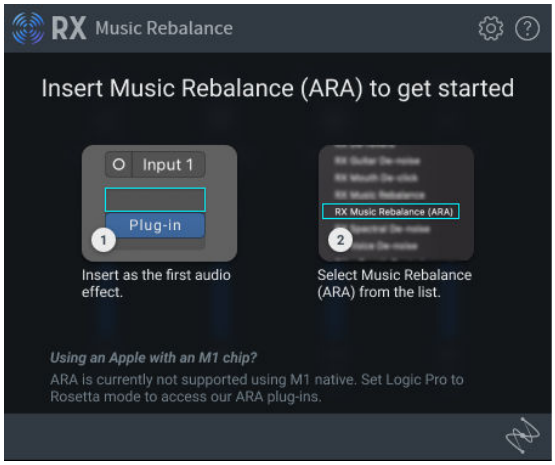
Stem Split

Stem Split automatically splits your selected audio into four distinct stems (Vocal, Bass, Percussion, Other), delivering each stem to its own tab. This process ignores the Gain settings for individual stems and always assumes 0 dB of gain. Once the stems are separated, you can process and export each stem as an independent audio file.

Getting started with Music Rebalance in Logic

Music Rebalance can be used in Logic as an ARA 2.0 plug-in (Rosetta only). In order to initialize the ARA version of Music Rebalance, please follow the following steps:

1. Update to the latest version of Logic Pro for full compatibility with the ARA plug-in format. There are known issues in versions prior to Logic Pro 10.7
2. Add the **RX 11 Music Rebalance** plug-in to a track.



3. Close and reopen Logic.
4. Add RX 11 Music Rebalance (ARA) as the first insert.

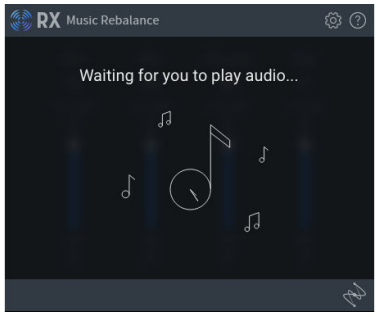
Still not seeing RX 11 Music Rebalance (ARA) appear after restarting? Make sure you are adding it as the first insert on your track. Then try closing Logic and deleting your AudioUnits cache here: ~/Library/Caches/AudioUnitCache, and repeat the steps above.

Using Music Rebalance (ARA) in Logic

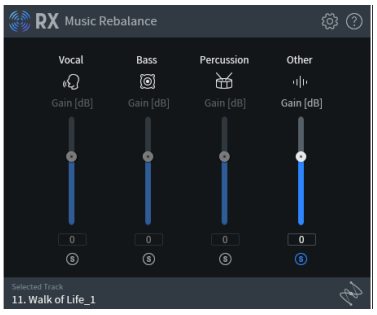
Now that you have Music Rebalance (ARA) on your track, you're ready to go.

1. Start playback to have Music Rebalance analyze your audio.
2. Once playback is started, you can pause playback and analysis will continue.

> ARA plug-ins are updated when playback is started. If you make a change to your audio, it will not be reflected until you've stopped and started playback again.



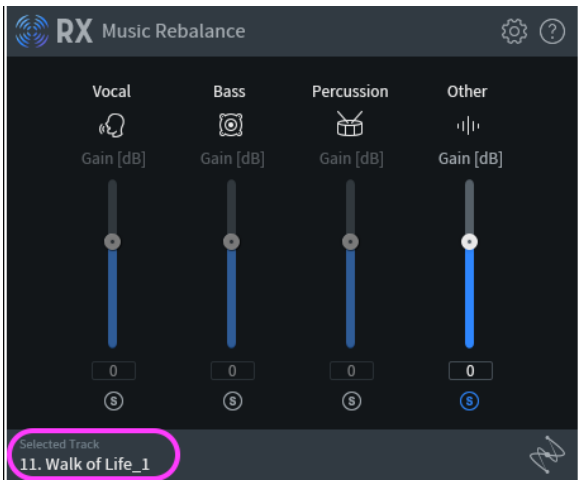
3. Once you have analyzed you're audio, you'll have access to the controls:



4. Move the sliders to rebalance your music in real time.The ARA plug-in uses the Good quality mode. If you need even higher quality, use Music Rebalance in the RX 11 standalone editor.

Using Multiple Instances

When using multiple instances of the plug-in, the Music Rebalance window will follow the selected track. The track name is shown at the bottom of the window.



If you select a track that doesn't contain Music Rebalance ARA, the plug-in will remember the last selected track that contains an instance of Music Rebalance ARA.

If you have multiple Music Rebalance ARA plug-in windows open, they will all show the controls and settings for the last selected track with Music Rebalance ARA.

Music Rebalance (ARA) Requirements & Limitations

- Music Rebalance (ARA) is supported in Rosetta mode only, as a limitation of Logic Pro
- Music Rebalance (ARA) must be the first plug-in on a track.
- Music Rebalance (ARA) cannot be added to a software instrument track or bus.
- Music Rebalance (ARA) is not supported by Selection Based Processing.
- Host presets are not supported by Music Rebalance (ARA).
- Automation is not supported by Music Rebalance (ARA).
- Playback must be started to gather audio data.
  - After editing, playback must be stopped and started again to update data.
- Music Rebalance (ARA) cannot process Apple Loops, compressed audio, flex audio, or reversed regions.
  - These regions can be bounced in place in order to be processed by Music Rebalance (ARA).

39. Spectral De-noise

STD & ADV | **Module & Plug-in**

Overview

Spectral De-noise is designed to remove stationary or slowly changing tonal noise and broadband hiss by learning a profile of the offending noise and then subtracting it from the signal. It can be useful for tape hiss, HVAC systems, outdoor environments, line noise, ground loops, camera motors, fans, wind, and complex buzz with many harmonics.

Spectral De-noise learns a profile of the background noise, then subtracts that noise when a signal's amplitude drops below the specified threshold. It is a flexible tool that can be used to quickly achieve accurate, high-quality noise reduction. It also provides separate controls for tonal and broadband noise, management of denoising artifacts, and an editing interface for controlling reduction across the frequency spectrum.

Controls



Learn

When Learn is enabled, Spectral De-noise will capture a noise profile from your selection. After a noise profile is captured using Learn, it remains fixed for the duration of processing. Manually learned noise profiles are best suited to removing or reducing noise that is constant and continuous throughout the duration of the file.

To learn a noise profile in Spectral De-noise:

1. Make a selection of the longest section of noise you can find in your file (ideally a few seconds in length).

2. Click the Learn button to capture a noise profile:
   - To capture a noise profile in the RX Audio Editor Spectral De-noise module, make a selection and click “Learn”
   - To capture a noise profile in the RX Spectral De-noise plug-in, engage the Learn button and playback audio, OR choose “Preview” in Audiosuite to capture the noise profile from your current selection.

MORE INFORMATION ABOUT LEARNING NOISE PROFILES

See the More Information section below to learn more about getting the best results when capturing a noise profile and Learning a noise profile from multiple selections in the RX 11 Audio Editor.

Adaptive Mode

When Adaptive Mode is enabled, the noise profile used for Spectral De-noise processing will change based on the incoming audio. Adaptive mode can work well with noise sources that are constantly changing, like recordings in outdoor environments, traffic noise, or ocean waves.

SPECTRAL DE-NOISE ADAPTIVE MODE PERFORMANCE NOTE

Adaptive mode in Spectral De-noise uses a significant amount of memory and computational power. For a more efficient form of adaptive noise reduction, try the Adaptive mode in Voice De-noise, which is designed to be highly efficient and zero-latency.

Learning Time [s]

Determines the amount of lookahead time used by Adaptive mode when learning noise profiles that change over time.

Threshold (Noisy/Tonal)

Controls the amplitude separation of noise and useful signal levels.

- Higher threshold settings reduce more noise, but also suppress low-level signal components.
- Lower threshold preserves low-level signal details, but can result in noise being modulated by the signal. Threshold elevation can be done separately for tonal and random noise parts. A good default is 0 dB.

TIP

If background noise changes in amplitude over time (like traffic noise or record surface noise), raise the Threshold to accommodate for the changes.

Reduction (Noisy/Tonal)

Controls the desired amount of noise suppression in decibels. Spectral De-noise can automatically separate noise into tonal parts (such as hum, buzz or interference) and random parts (such as hiss). You can specify the amount of suppression for these parts separately (e.g. in some situations it can be desirable to reduce only unpleasant buzz while leaving unobjectionable constant hiss).

> **NOTE**
>
> Strong suppression of noise can also degrade low-level signals, so it is recommended to apply only as much suppression as needed for reducing the noise to levels where it becomes less objectionable.

Quality

Affects the quality and computational complexity of the noise reduction. This selection directly affects CPU usage. RX's Spectral De-noise module offers four algorithms that vary in processing time.

- **A:** is the least CPU intensive process and has the lowest latency, which makes it the most suitable for real-time operation. It reduces musical noise artifacts by time smoothing of the signal spectrum.
- **B:** achieves more advanced musical noise suppression by using adaptive 2D smoothing (both time and frequency). It is more CPU intensive and has more latency, but can still run in real-time on some machines.
- **C:** adds multiresolution operation for better handling of signal transients and even fewer musical noise artifacts. It is a very CPU intensive algorithm and is not recommended for real-time operation.
- **D:** adds high-frequency synthesis for reconstruction of signal details buried in noise. The speed of algorithm D is similar to algorithm C. This algorithm is not recommended for real-time operation.

Artifact Control

Determines how much noise reduction will depend upon either spectral subtraction or wide band gating.

- When using lower values, noise reduction will rely on spectral subtraction. This can more accurately separate noise from the desired audio signal, but can produce musical noise artifacts, resulting in a "chirpy" or "watery" sound during heavy processing.
- When using higher values, the noise reduction will rely more heavily upon wider band gating which will have fewer musical noise artifacts, but sound more like broadband gating, resulting in bursts of noise right after the signal falls below the threshold.

Noise Spectrum Display

The Noise Spectrum display shows useful information during both playback and when the noise reduction process is being applied.

- **Noise Spectrum Color Legend**
  - **Input (Gray):** spectrum of input audio signal
  - **Output (White):** spectrum of the denoised output audio signal
  - **Noise Profile (Orange):** the learned noise profile plus offset from the Threshold control
  - **Residual Noise (Yellow):** desired noise floor after denoising, can be controlled by modifying the Reduction Curve
  - **Reduction Curve (Blue):** manual weighting of the noise reduction across the spectrum

Smoothing

When the Reduction Curve is enabled, this controls the amount of interpolation between your reduction curve points, allowing for sharper or more gradual slopes between edit curve points.

Reduction Curve

When enabled, allows for fine tuning of the reduction spectrum with up to 25 edit points. This enables you to customize the amount of noise reduction being applied across different frequency regions.

- Higher edit point values result in less noise reduction in the associated frequency region.
- Lower edit point values result in more noise reduction in the associated frequency region.
- **For example**, if you wanted to reduce some low HVAC rumble but preserve some energy in higher frequencies, you could drag the curve's leftmost point down a little bit, then create a point around 5 kHz and drag it up a bit.

You can interact with the Reduction Curve Edit points in the following ways:

▶ Left-click to add an end point, displayed as gray box along envelope curve.

▶ Right-click or drag it outside the screen to emove an edit point.

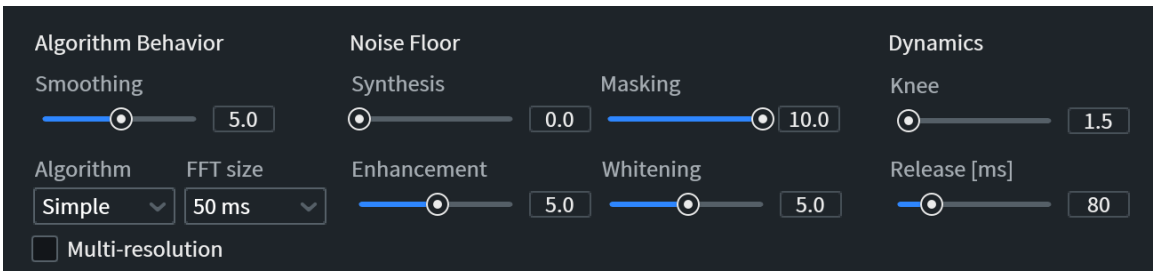
▶ Hold Shift while dragging reduction curve points to lock them on an axis.

▶ Hold Control/Command to get very fine control over positioning.

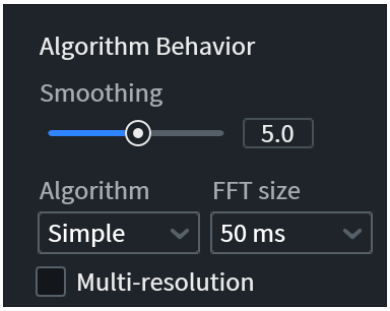
Reset

Returns the Noise Reduction Curve to its default setting of 0dB.

Advanced Settings



Algorithm Behavior (Advanced Settings)



Smoothing - Advanced

Controls the reduction of musical noise artifacts which can be a result of heavy denoising.

What is musical noise?

Musical noise is caused by random statistical variations of noise spectrum that cause random triggering of sub-band gates. These artifacts are sometimes described as “chirpy” or “watery” sounds left behind during the noise reduction process.

Algorithm

Selects the smoothing algorithm for the removal of random ripples (“musical noise” artifacts) that can occur in the spectrogram when processing your audio. The strength of smoothing is controlled by the Smoothing slider.

- **SIMPLE:** Performs independent noise gating in every frequency channel of FFT. Release time of sub-band gates is controlled by the Release slider. This is a fast algorithm with low latency that is suitable for real-time operation.
- **ADVANCED & EXTREME:** Perform joint time-frequency analysis of the audio signal which results in better quality and fewer “musical noise” artifacts. These algorithms have higher latency and computational complexity.

FFT Size (ms)

Selects the time and frequency resolution of the processing.

- Higher FFT sizes give you more frequency bands allowing you to cut noise between closely spaced signal harmonics, or cut steady-state noise harmonics without affecting adjacent signals.
- Lower FFT sizes allow for faster response to changes in the signal and produce fewer noisy echoes around transient events.

> ⚠ **RE-LEARN YOUR NOISE PROFILE IF YOU CHANGE FFT SIZE**
>
> If the FFT size is changed, it is recommended that you run the De-noise module’s Learn feature again because the old noise profile was taken at a different FFT size and therefore becomes inaccurate.

Multi-Res

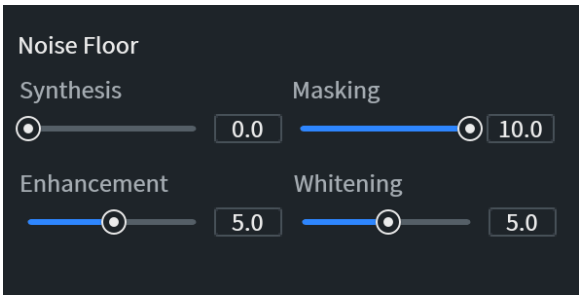
Enables multi-resolution processing for the selected algorithm type. When you select the Multi-res checkbox, the signal is analyzed in real-time and the most appropriate FFT size is chosen for each segment of the signal. This is done to minimize the smearing of transients and at the same time achieve high frequency resolution where it is needed.

> **NOTE**
>
> The FFT size control does not have any effect in multi-resolution mode as the FFT resolution is selected automatically. The noise profile does not need to be re-learned when switching to multi-resolution mode.

Fast Fourier Transform (FFT)

A procedure for the calculation of a signal frequency spectrum. The greater the FFT size, the greater the frequency resolution, i.e., notes and tonal events will be clearer at larger sizes. However, when using FFT-based processing, the more audio you remove from your source, the more likely you are to create undesirable artifacts.

Noise Floor (Advanced Settings)



- **Synthesis**: Synthesizes high frequency material after denoising.
  - When Synthesis is set to a value greater than zero, signal harmonics are synthesized after denoising. The synthesized harmonics remain at the level of the noise floor, and serve to fill in gaps in high frequencies caused by processing.
  - Increasing Synthesis can increase the sense of life and air in processed audio. Too much Synthesis may cause apparent distortion in the signal.
- **ENHANCEMENT:** Enhances signal harmonics that fall below the noise floor.
  - Enhancement predicts a signal's harmonic structure and places less noise reduction in areas where possible signal harmonics could be buried in noise. This aids in preserving high-frequency signal harmonics that may be buried and not detected otherwise.
  - Enhancement can make the resulting signal brighter and more natural sounding, but high values of harmonic enhancement can also result in high-frequency noise being modulated by the signal.

- **MASKING:** Reduces the depth of noise reduction where you wouldn't perceive any effect from it.
  - Masking enables a psychoacoustic model that dynamically controls suppression amount to facilitate the use of softer suppression where noise is subjectively inaudible. When noise in certain regions is calculated to be inaudible, this feature prevents any signal processing in these regions. This potentially reduces the amount of processing done to the signal and may positively affect overall signal integrity. The position of the slider controls the influence of psychoacoustic model on suppression levels.
  - If you need to cut very high, inaudible frequencies, set this to 0. Otherwise, leave this at 10.

    NOTE

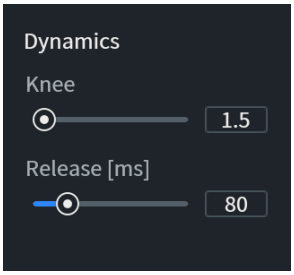
    When the Masking slider is set to 0, the feature is turned off, and the amount of noise suppression is uniformly governed to the yellow curve in spectrum analyzer (more precisely — by the difference between the yellow curve and orange curve).

- **WHITENING:** Shapes the noise floor after processing to be more like white noise. Whitening modifies the amount of noise reduction (shown by the yellow curve) applied at different frequencies to shape the spectrum of the residual noise.
  - When Whitening is set to zero, the suppression is uniform at all frequencies, as controlled by Reduction (tonal/broadband) sliders, and the suppressed noise has a similar spectral shape to the original noise.
  - When Whitening is set to the maximum value, the desired shape of suppressed noise floor is made close to white noise, so that residual noise has more neutral sound.

    UNDERSTANDING THE EFFECT OF THE WHITENING CONTROL

    Changing the noise floor balance with Whitening can help prevent gaps from over-processing, but an unnaturally white noise floor can introduce problems like noise modulation when editing or mixing with other noises from a unique space (like a set location.)

Dynamics (Advanced Settings)



- **KNEE:** Controls how surgical the algorithm's differentiation is between the signal and noise. This slider controls the sharpness of the gating knee in the denoising process.
  - At higher values, transitions in the De-noise are more abrupt and can become prone to errors in the detection of the signal with respect to the noise.
  - At lower values, the denoising becomes more forgiving around the knee, and applies less attenuation to signals that are only slightly below the threshold. This may result in a lower depth of noise reduction, but can also have fewer artifacts.

- **RELEASE [ms]:** Selects the release time of sub-band noise gates in milliseconds. Longer release times can result in less musical noise, but may also reduce or soften the signal's initial transients or reverb tails after the signal's decay.

> **NOTE**
> The Release control is only available when the **Simple** algorithm is selected.

More Information

Tips for getting the best results when learning noise profiles manually

- Before learning a noise profile, identify and select the longest section (ideally a few seconds in length) of the recording that contains *only* the noise you wish to remove or reduce.
- To ensure the best results, your selection should not contain any content that you wish to preserve (for example, do not include any audio you do not consider to be "noise" in your selection.)
- Usually you can find noise only sections at the beginning or end of a file, or during a pause or break in the recording (for example, a pause between words in dialogue recordings.)

Learning a Noise Profile From Multiple Selections

In the RX standalone application, it is possible to create a spectral profile from multiple isolated selections. This is useful when you have a file where it's impossible to find enough isolated noise to build the profile.

For example, if you are trying to restore a file where someone is speaking over noise, you can select noise in frequencies where none of the voice is present at a given time. If you select enough of this noise with the Lasso or Brush selection tools, you can create an accurate noise profile that will give you good results with Spectral De-noise. You can create more than one selection at a time by holding Shift while making a selection.

Select noise anywhere you can to build a better noise profile.

This feature is not available in the Spectral De-noise plug-in because it requires using RX's spectral selection tools as well as accurate calculation of the time and frequency of the selected areas. If you are unable to create a full noise profile with multiple selections, RX can try to build a reasonable noise profile out of your existing profile. If you have an incomplete noise profile, RX will ask you if you want it to complete the profile.

For example, if you can only capture a low frequency rumble below 100 Hz, some broadband noise between 200 Hz and 5000 Hz, and all the noise above 8000 Hz, RX can fill in the gaps for you. Building a profile from multiple selections gives you some flexibility, and RX will guess any noise you missed.

40. Spectral Repair

STD & ADV

Overview

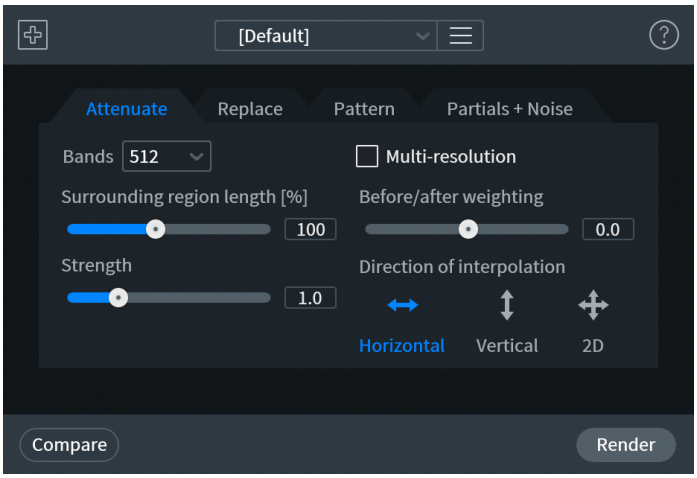
Spectral Repair intelligently removes undesired sounds from a file with natural-sounding results. This tool treats selections within the spectrogram/waveform display as corrupted audio that will be repaired using information from outside of the selection. Select the noise you want to repair and Spectral Repair will reduce it to the level of the noise floor, replace it with audio from around the selection, or generate entirely new audio to fit the selection.

Attenuate

This mode removes sounds by comparing the content inside of a selection to the content outside of the selection. Attenuate reduces spectrogram magnitudes in the selected area to match magnitudes from the surrounding area, resulting in the removal of the sound without leaving an audible gap behind. Attenuate does not resynthesize any audio. It modifies dissimilar audio in your selection to be more similar to the surrounding audio.

Attenuate is suitable for recordings with background noise or where noise is the essential part of music (drums, percussion) and should be accurately preserved. It's also good when unwanted events are not obscuring the desired signal completely. For example, Attenuate can be used to bring noises like door slams or chair squeaks down to a level where they are inaudible and blend into background noise.

Controls



- **BANDS:** Selects the number of frequency bands used for interpolation.
  - A higher number of bands can provide better frequency resolution, but also requires wider surrounding area to be analyzed for interpolation.
  - A lower number of bands is ideal for processing short selections or transient signals.
- **SURROUNDING REGION LENGTH:** Defines how much of the surrounding content will be used for interpolation.

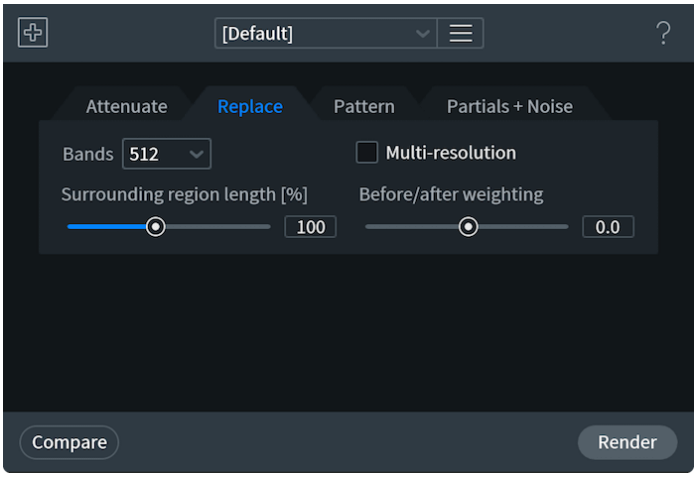
- **STRENGTH:** Adjusts strength of attenuation.
- **MULTI-RESOLUTION:** The multi-resolution mode allows for better frequency resolution for the interpolation of low-frequency content and better time resolution for the interpolation of high-frequency content.
- **BEFORE/AFTER WEIGHTING:** Gives more weight to the surrounding audio before or after the selection.
- **DIRECTION OF INTERPOLATION:** Determines where the material used in the repair process is located in relation to the current selection.
  - **HORIZONTAL:** Signal to the left and right of the current selection will be used for interpolation.
  - **VERTICAL:** Signal above and below the current selection will be used for interpolation.
  - **2D:** Signal above, below, to the left and to the right of the current selection will be used for interpolation.

> **NOTE: DIRECTION OF INTERPOLATION CONTROL AVAILABILITY**
>
> Replace, Pattern and Partials + Noise tabs only utilize Horizontal mode and do not display this option.

Replace

The Replace tab can be used to replace badly damaged sections (such as gaps) in tonal audio. It completely replaces the selected content with audio interpolated from the surrounding data.

Replace Controls



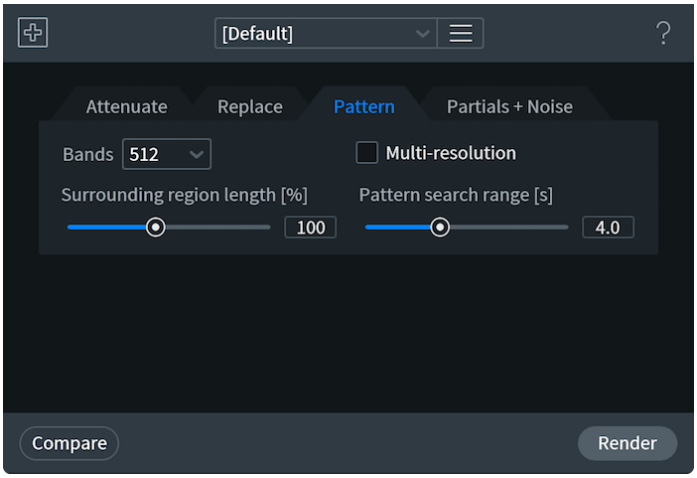
- **BANDS:** Selects the number of frequency bands used for interpolation.
  - A higher number of bands can provide better frequency resolution, but also requires wider surrounding area to be analyzed for interpolation.
  - A lower number of bands is ideal for processing short selections or transient signals.
- **SURROUNDING REGION LENGTH:** Defines how much of the surrounding content will be used for interpolation.
- **MULTI-RESOLUTION:** The multi-resolution mode allows for better frequency resolution for the interpolation of low-frequency content and better time resolution for the interpolation of high-frequency content.

- **BEFORE/AFTER WEIGHTING:** Gives more weight to the surrounding audio before or after the selection.

Pattern

This mode finds the most similar portion of the surrounding audio and uses this to replace the corrupted audio. Pattern mode is suitable for badly damaged audio with background noise or for audio with repeating parts.

Pattern Controls



- **BANDS:** Selects the number of frequency bands used for interpolation.
  - A higher number of bands can provide better frequency resolution, but also requires wider surrounding area to be analyzed for interpolation.
  - A lower number of bands is ideal for processing short selections or transient signals.
- **SURROUNDING REGION LENGTH:** Defines how much of the surrounding content will be used for interpolation.
- **MULTI-RESOLUTION:** The multi-resolution mode allows for better frequency resolution for the interpolation of low-frequency content and better time resolution for the interpolation of high-frequency content.
- **PATTERN SEARCH RANGE:** Selects the length of the audio segment used in a search for a suitable replacement interval. For example, setting it to 5 seconds will allow search within ±5 second range from the selection.

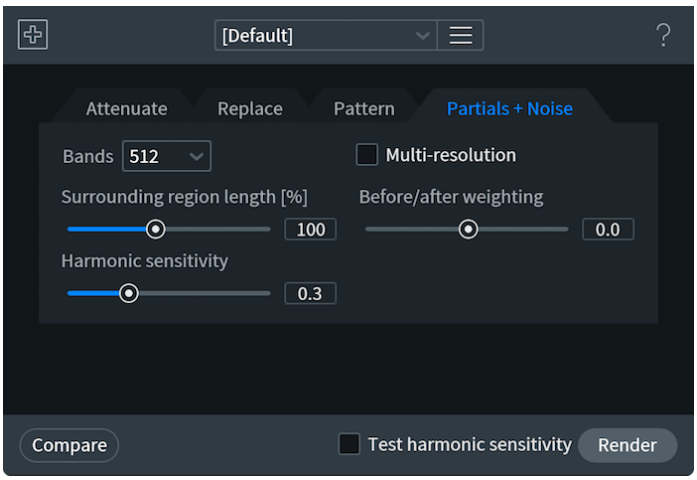
Partials and Noise

The advanced version of Replace mode. It restores harmonics of the audio more accurately with control over the Harmonic sensitivity parameter.

This mode allows for higher-quality interpolation by explicit location of signal harmonics from the 2 sides of the corrupted interval and linking them together by synthesis.

Partials + Noise is able to correctly interpolate cases of pitch modulation, including vibrato. The remaining of non-harmonic material (“residual”) is interpolated using Replace method.

Partials and Noise Controls



- **BANDS:** Selects the number of frequency bands used for interpolation.
  - A higher number of bands can provide better frequency resolution, but also requires wider surrounding area to be analyzed for interpolation.
  - A lower number of bands is ideal for processing short selections or transient signals.
- **SURROUNDING REGION LENGTH:** Defines how much of the surrounding content will be used for interpolation.
- **HARMONIC SENSITIVITY:** Adjusts amount of detected and linked harmonics.
  - Lower values will detect fewer harmonics
  - Higher values will detect more harmonics and can introduce some unnatural pitch modulations in the interpolated result.
- **MULTI-RESOLUTION:** The multi-resolution mode allows for better frequency resolution for the interpolation of low-frequency content and better time resolution for the interpolation of high-frequency content.
- **BEFORE/AFTER WEIGHTING:** Gives more weight to the surrounding audio before or after the selection.

Surrounding Region Display

- **SURROUNDING REGION SHADING:** When using the Spectral Repair module, your selections will be shown with a dotted line surrounding your selected region. This dotted line is directly controlled by the Surrounding Region and Before/After Weighting controls inside of your Spectral Repair modules, and provides a visual representation of your set values.
  - The surrounding region is the region that RX uses for interpolation of the selected region. The data from the surrounding region is used to restore the selected region.

Workflow

Applying Spectral Repair

1. To start working with Spectral Repair, switch to the spectrogram view by dragging the spectrogram/waveform transparency balance slider to the right.

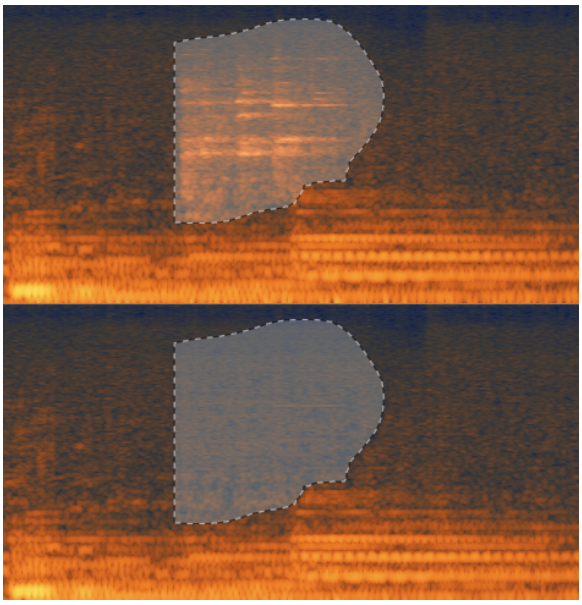
2. Next, identify the unwanted event on a spectrogram and select it using a selection tool (use the time-frequency, brush, lasso or magic wand tools to fit your selection to the corrupted audio)
3. You can audition just the signal in your selection by pressing the Play Selection button in the RX transport.
4. Once you've found the event(s) to repair, select the appropriate Spectral Repair mode from the tabs at the top of the Spectral Repair settings window.
5. You can use the Compare Settings functionality to audition the processing before applying it, or click the process to apply the active Spectral Repair tab's settings.

More information

This section contains useful information, examples and tips for getting the most out of the Spectral Repair module.

Visual example

The following image shows the selection before processing on the top and the selection after processing with Spectral Repair on the bottom.



Processing limitations

Depending on the mode and settings, Spectral Repair will have varying limits to the amount of audio that can be processed in your selection.

- Unlimited — Attenuate when in Vertical mode only.
- 10 seconds — Attenuate Horizontal or 2D; Replace modes.
- 4 seconds — Pattern, Partials + Noise modes.
- Longer selections will automatically adjust processing to use the correct mode.

Spectral Repair as an alternative to De-click processing

When used with a time selection, Spectral Repair is able to provide higher quality processing than De-click for long corrupted segments of audio (above 10 ms).

Example use cases for Spectral Repair

When used with a time/frequency, lasso, brush, or wand selection, it can be used to remove (or attenuate) unwanted sounds from recordings, such as: squeaky chairs, coughs, wheezes, burps, whistles, dropped objects, mic stand bumps, clattering dishes, mobile phones ringing, metronomes, click tracks, door slams, sniffles, laughter, background chitchat, digital artifacts from bad hardware, dropouts from broken audio cables, rustle sounds from microphone movements, fret and string noise from guitars, ringing tones from rooms or drum kits, squeaky wheels, dog barks, jingling change, or just about anything else you could imagine. Spectral Repair can also effectively repair gaps or replace audio intelligently by using advanced resynthesis techniques.

Increase efficiency with the Find Similar Event Tool

Some unwanted events consist of several separate regions on a spectrogram. In some cases, it's possible to achieve more accurate results by repairing several smaller selections one by one, instead of one large selection. You can use the Find Similar Event tool to save time when searching for and fixing many similar events in large files.

Use the Compare Settings window to experiment with Spectral Repair processing

Sometimes it's worth trying several different methods or number of bands to achieve the desired result. A higher number of bands doesn't necessarily mean higher quality! We encourage you to use the Compare Settings window to experiment and find the best settings for the project at hand.

Adjust Surrounding Region Length and Before/after weighting to perfect your processing

Common parameters for many modes include Surrounding region length which determines how far around the selection Spectral Repair will look for a good signal. Before/after weighting allows you to use more information from either before or after the selection for interpolation. For example, if your unwanted event is just before a transient (such as a drum hit) in the audio, you may want to set this parameter to use more of the audio before the selection to prevent smearing of the transient into the selection.

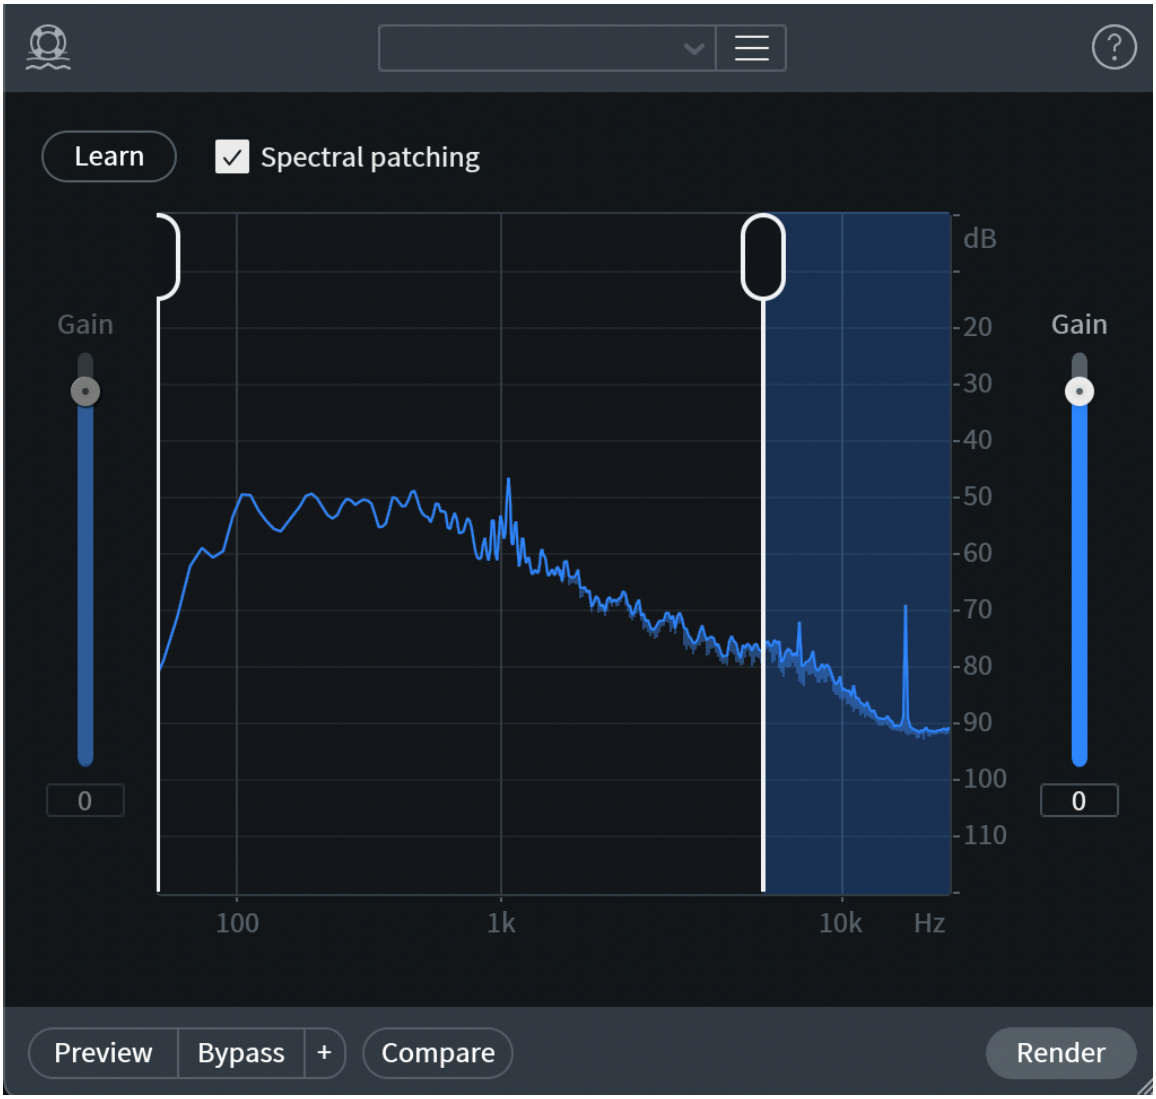
41. Spectral Recovery

ADV

Overview

Spectral Recovery automatically adds missing frequencies to bandwidth-limited speech content and patches holes in the frequency spectrum due to compression artifacts. Spectral Recovery is particularly useful in handling VoIP recordings, from Skype or Zoom, where there is a hard cutoff above which no audio exists and the natural low end may be filtered or reduced.

Controls



Learn

When Learn is pressed, Spectral Recovery analyzes your selection to determine the suggested values for the high and low cutoff frequencies.

Spectral Patching

When selected, this automatically fills in holes in the spectrogram between the high and low cutoff frequencies by sampling the area around the missing audio. This control is very subtle audibly, but you can see the effect in the Spectrogram.

Low Gain

This sets the level of synthesized low frequency signals below the cutoff.

Low Cutoff

Spectral Recovery synthesizes audio below this frequency and fills the spectral holes above this frequency.

TIP

Try adjusting the low cutoff to just above the fundamental frequency of your speaker for Voice-over-IP material to make it sound more natural.

High Cutoff

Spectral Recovery synthesizes audio above this frequency and fills the spectral holes below this frequency.

High Gain

This sets the level of synthesized high frequency signals above the cutoff.

Render

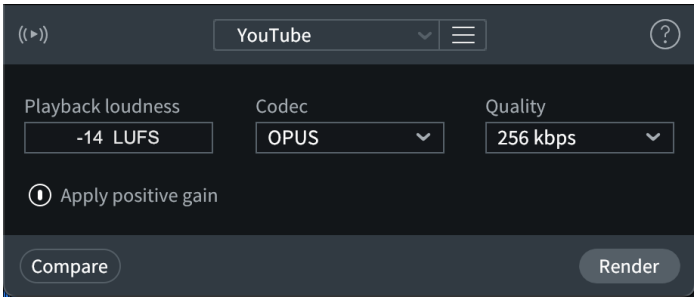
To apply Spectral Recovery processing to the current selection, click the Render button in the bottom right hand corner of the window.

42. Streaming Preview

Overview

The Streaming Preview module applies the same loudness normalization and lossy encoding that streaming platforms will apply. This means you can use Streaming Preview to hear how your song will sound on those platforms and compare it with reference songs. We hope this process helps you uncover improvements for your mixes and masters and gives you more confidence to release your music.

Controls



- **PLAYBACK LOUDNESS:** Represents the loudness normalization target for the streaming platform in Loudness Units Full Scale (LUFS). When Streaming Preview is applied, gain will be applied to your audio to match this loudness. The loudness measurement and normalization occurs after lossy encoding has been applied.
- **APPLY POSITIVE GAIN:** When this option is enabled and your audio is quieter than the Playback Loudness you're targeting, positive gain will be applied to bring your audio up to the target. However, streaming platforms like Spotify will not boost your audio to be greater than -1 dB True Peak, so the same is true in the Streaming Preview module. For example, if an audio file's loudness is -20 LUFS, and its True Peak maximum is -5 dB FS, Streaming Preview will only lift the track up to -16 LUFS.
- **CODEC:** Selects the lossy encoding format that streaming platforms apply to reduce the amount of data required to stream and store audio.
  - **LOSSLESS (NONE):** Does not apply any lossy compression to simulate streaming platforms offering HiFi lossless audio like Tidal and Apple Music.
  - **OGG:** A free, open container format maintained by the Xiph.Org Foundation. It is primarily used by Spotify.
  - **MP3:** (MPEG Audio Layer III) is a coding format for digital audio developed largely by the Fraunhofer Society in Germany.
  - **AAC:** (Advanced Audio Coding) is designed to be the successor of the MP3 format and generally achieves higher sound quality than MP3 at the same bit rate. It is used when streaming lossy audio through Apple Music and Tidal.
  - **OPUS:** Opus is a lossy audio coding format developed by the Xiph.Org Foundation and standardized by the Internet Engineering Task Force. It's designed to efficiently code speech and general audio in a single format, while remaining low-latency enough for real-time interactive communication and low-complexity enough for low-end embedded processors. It is the primary codec used by YouTube.

- **QUALITY:** Selects the bitrate for the lossy encoding. This represents the amount of information, or detail, that is stored per unit of time of a recording. Lower bit rates will result in more audio loss and a reduction in sonic quality. When the Lossless codec option is selected, the Quality control is not accessible.

Workflow and Presets

The recommended workflow with the Streaming Preview module is to use the provided presets. We have researched the major streaming platforms and created these presets to represent their Playback loudness levels and codecs. Some streaming platforms enable you to make changes that will affect these settings and these options are also represented in the presets. Streaming platform companies may change their settings and we will do our best to keep the presets up-to-date if the settings change.

Streaming Preview is designed **only** for auditioning codec processing/loudness normalization. We **do not** recommend exporting or uploading your audio to a streaming service after applying Streaming Preview processing. This would effectively bake in any lossy encoding you hear in the Preview.

43. Voice De-noise

Module & Plug-in

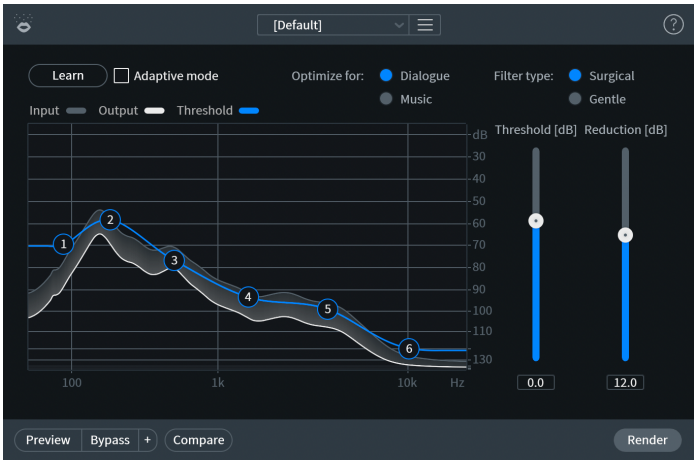
Overview

Voice De-noise is an intuitive, zero latency de-noiser that offers high quality results on a variety of material. Voice De-noise can intelligently analyze speech signals and determine the best noise threshold for your signal. In a DAW, this feature can be used to write automation in case you need to override the automatic settings and correct the noise threshold by hand.

How does Voice De-noise processing work?

- Under the hood is a series of 64 psychoacoustically spaced bandpass filters which act as a multiband gate to pass or stop a signal based on user-defined threshold values.
- If a signal component is above the threshold for the filter, it will be passed (not processed).
- If a signal component is below the threshold for the filter, it will be attenuated (processed).

Controls



- **ADAPTIVE MODE:** Analyzes the incoming signal and adjust the noise threshold automatically to compensate for changes in the noise floor. This can be useful for removing noise from recordings with variable noise floor and continual noisy sections, and works well for almost any recording of dialogue and spoken word.
  - The noise threshold settings in Adaptive Mode may be different from the settings achieved by running Learn to set the noise threshold manually.
  - Because the adaptive noise threshold is continually being adjusted, it is set lower to prevent artifacts from occurring.
- **LEARN:** When using Manual mode, you can use the Learn button to set the noise threshold to a noise reference.
  - Find a passage of pure noise in your audio and use Learn to analyze it.
  - Longer selections of noise will set the Threshold Nodes to more ideal locations.
  - We recommend finding at least one second of pure noise to Learn your noise profile from.

- **OPTIMIZE FOR DIALOGUE OR MUSIC:** Because dialogue tends to be in short bursts and vocals tend to have sustained notes, we've added modes to provide better results when applying Voice De-noise processing.
  - **Optimize for DIALOGUE** reacts to noise changes faster and isn't meant to handle the noise found in sung vocals.
  - **Optimize for MUSIC** does not attenuate sustained notes and is more transparent when applied to sung vocals.
- **FILTER TYPE:** Choosing the filter type changes the bandwidth of the noise reduction filters.
  - **SURGICAL Mode:** Removes more unwanted noise than Gentle Mode, but makes some sacrifices in terms of timbre and can lead to musical noise artifacts. Musical noise is caused by random statistical variations of noise spectrum that cause random triggering of sub-band gates. These artifacts are sometimes described as "chirpy" or "watery" sounds left behind during the noise reduction process.
  - **GENTLE Mode:** provides more transparent noise reduction than Surgical, but removes less high end "sizzle."
- **THRESHOLD NODES:** The Threshold Node controls on the frequency spectrum display allow you to modify the noise threshold curve, which can be thought of as the "noise profile." These six points can be adjusted manually to suit the noise currently in your signal. These controls can be automated to compensate for shifts in the audio's noise floor.
  - **In ADAPTIVE Mode**, the Threshold Nodes are adjusted automatically in real-time.
  - **In MANUAL mode**, more than one Threshold Node can be selected at a time for manual adjustment by clicking and dragging anywhere on the interface.
- **THRESHOLD:** The master Threshold control allows you to offset all Threshold Node values by the same amount. If you find that processing is too aggressive or processing is affecting audio you want to leave unprocessed, try adjusting this control.
- **REDUCTION:** Provides control over the maximal depth of noise reduction (in dB) that will occur per frequency band while a signal component is below its threshold. If you have your thresholds set properly and don't like the results you're getting, try adjusting this control.
- **METERING**
  - The **Input Spectrum** meter shows the level of the signal at the input of the denoiser filters.
  - The **Output Spectrum** meter shows the level of the signal at the output of the denoiser filters.
  - The **Gain Reduction Region** is the area between the Input and Output Spectra. This shows the amount of noise reduction processing being applied to your signal.

Voice De-noise Plug-in

Voice De-noise has been specifically designed to provide high efficiency, zero latency adaptive noise removal when inserted on a track in your DAW or NLE. The Spectral De-noise plug-in is far more resource intensive and uses higher latency.

44. Wow & Flutter

ADV

Overview

The Wow & Flutter module allows for transparent correction of pitch variations in your audio due to inconsistent speed of the recording medium. Wow & Flutter can be tailored to different rates of wow, and allows for global pitch correction to compensate for static pitch offsets. The detected wow or flutter pitch contour can also be written to the spectrogram for diagnostic purposes.

> **NOTE**
>
> When applied to a bandwidth-limited selection, Wow & Flutter will only use the selected frequency band for detection of pitch variations. When applied to a full-band selection, the algorithm will automatically determine the most suitable frequency range for analysis. Multiple selections will be combined into one time selection and processed as a whole.

Controls



Wow

Addresses slower pitch variations (typically 0 to 5 Hz).



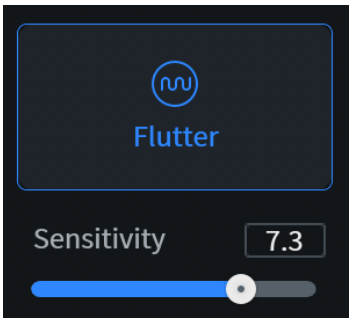
You can choose from the following **Wow Rate** options:

- **Fast:** targets wow in the 2 to 8 Hz range.
- **Medium:** targets wow in the 0.5 to 2 Hz range.
- **Slow:** targets wow in the 0 to 0.5 Hz range.

When this mode is selected, a **Display wow** checkbox is available in the module footer. See the Display Wow/Display Flutter section below for more information.

Flutter

Addresses quicker pitch variations (8 to 40 Hz).



When this mode is selected, a **Display Flutter** checkbox is available in the module footer. See the Display Wow/Display Flutter section below for more information.

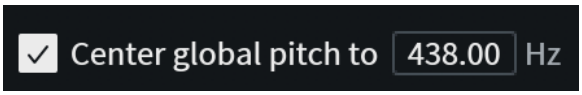
Sensitivity

Determines how aggressively the pitch variations will be corrected. A lower sensitivity can prevent desirable pitch fluctuations like vibrato from being impacted.

Center Global Pitch

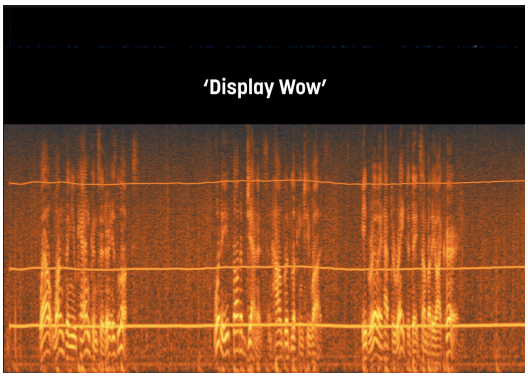
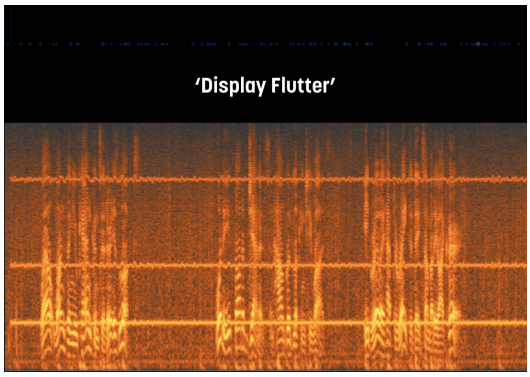
When checked, this control centers the audio around the specified concert pitch. If a static pitch offset is present in the material, it will be corrected toward a specific tuning of the musical scale. The default pitch value is 440 Hz.

When unchecked, processing will preserve the length of the selection and its average pitch intact.



Display Wow/Display Flutter

Enabling this checkbox will render the detected pitch contour of the wow or flutter onto the audio file. No correction will be applied, but the detected pitch variations will be printed as a diagnostic tool.

Rendered with Display Wow checked	Rendered with Display Flutter checked
	

NOTE

Wow & Flutter is not available in the Batch Processor, Module Chain or Composite View.

45. Azimuth

ADV

Overview

The Azimuth module provides control over left and right channel gain and delay. Azimuth adjustment can help repair stereo imbalances and phase issues that can occur as a result of issues introduced by speed inconsistencies.

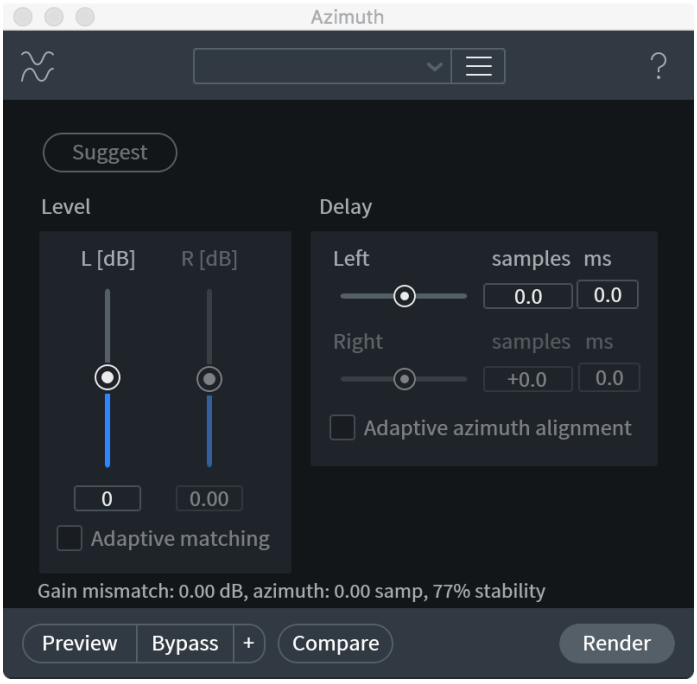
Controls



- **LEVEL (dB):** Adjusts the gain of the left and right audio channels
- **ADAPTIVE MATCHING:** Enables automatic gain adjustment of the right channel in order to match the level of the left channel over time.
- **DELAY (samples/ms):** Allows for manual adjustment over the delay in samples or milliseconds of the left and right audio channels. For very accurate azimuth correction, RX uses oversampling to achieve sub-sample delays.
- **ADAPTIVE AZIMUTH ALIGNMENT:** Enables automatic time-variable adjustment of the right channel's sample delay in order to align the waveform with the left channel.
- **SUGGEST:** Analyzes the selection and determines the appropriate amounts of fixed gain and delay to apply in order to align the two channels.

Suggest and Adaptive modes are only available on stereo files

The Suggest function and Adaptive Matching modes are meant to function on stereo files, these controls will be disabled in the Azimuth module interface when a mono file is selected.



More Information

- **Use Azimuth to fix gain and delay alignment issues**:
  - Azimuth adjustment can be useful to repair inconsistent gain or delay alignment between left and right channels.
  - For example, gain and delay alignment inconsistencies can be introduced by improper tape head alignment.
- **Use Azimuth before processing with Center Extract**: Azimuth adjustment is recommended to be applied before Center Extract processing to achieve the best results.

46. Dither

Overview

Dithering is a necessary process when converting audio from a higher bit resolution to a lower bit resolution. Dithering is used to tame the quantization distortion that happens when converting between bit depths due to requantization. Dither also preserves more of the dynamic range of a signal when converting to a lower bit depth. The Dither module applies iZotope's MBIT+ dithering and noise shaping technology to maintain the highest audio quality possible when you are converting to 24, 20, 16, 12, or 8 bits. MBIT+ uses psychoacoustic methods to distribute dithering noise into less audible ranges. The result is a more pleasing sound and smoother fades.

Controls



New Bit Depth

This sets the target resolution (bit depth) of the audio file.

Noise Shaping

Sets the aggressiveness of dither noise shaping. It is possible to provide more effective and transparent dithering by shaping the dithered noise spectrum so less noise is in the audible range and more noise is in the inaudible range. You can control the aggressiveness of this shaping, ranging from None (no shaping, plain dither) through Ultra (roughly 14 dB of audible noise suppression). More noise shaping can cause slightly higher peaks in your signal, even at high bit depths.

Dither Amount

The dithering amount can be varied from None (noise shaping only) to High.

- In general, the Normal dither amount is a good choice.

- No dithering or Low dither amount can leave some non-linear quantization distortion or dither noise modulation, while higher settings completely eliminate the non-linear distortion at the expense of a slightly increased noise floor.
- A dither amount setting of High with no noise shaping produces a standard TPDF dither: a common white noise generation method with a triangular distribution of amplitudes between −1 and +1 of the Least Significant Bit (LSB).

Auto-blanking

Automatically mutes dither output (i.e. dither noise) based on the selected mode and characteristics of the input signal. Auto-blanking mode options include:

- **IF QUANTIZED:** Mutes dither output when already dithered or quantized signal is detected. When existing quantization is detected, the auto-blanking feature will mute dither until a signal with no existing dither is detected.
- **ON SILENCE:** Mutes dither output when silence is detected.
- **OFF:** Disables auto-blanking.

Limit Noise Peaks

Dither noise is random in nature and has a very low amplitude. However, after noise shaping, especially in aggressive dithering modes like Ultra, the high-frequency dither noise is significantly amplified, and the overall dither signal can show spurious peaks up to −60 dBFS during 16-bit quantization. If such high peaks are undesirable, you can enable this option to effectively suppress the spurious peaks in the noise-shaped dither.

Suppress harmonics

If, for some reason, any dithering noise is undesirable, simple truncation remains the only choice. Truncation results in harmonic quantization distortion that adds overtones to the signal and distorts the timbre. In this case you can enable Suppress Harmonics option to slightly alter the truncation rules, moving the harmonic quantization distortion away from overtones of audible frequencies. This option doesn't create any random dithering noise floor. Instead it works more like truncation, but with better tonal quality in the resulting signal. This option is applicable only in the modes without dithering noise and without aggressive noise shaping.

47. EQ

Overview

RX includes an eight-band parametric EQ module with six adjustable notch/shelving filters and two adjustable passband filters. The EQ module is useful for manually shaping the overall sound of a file or selection, for both corrective and enhancement purposes. Beyond a traditional equalizer, the EQ also offers remarkably high Q values for precision filtering.

EQ can often be a simple first step to preparing a file for restoration, and can be used for cutting harsh high frequencies, removing rumble from dialogue, steeply high passing out wind noise from a location recording, or cutting distortion overtones to increase the intelligibility of a voice.

Controls



EQ TYPE

- **Analog (IIR):** a minimum-phase EQ with analog-like EQ curves.
- **Digital Linear Phase (FIR):** a linear-phase EQ with surgical EQ curves.

FREQUENCY PRECISION

This controls the number of bands, i.e. frequency resolution vs. time resolution, for the FIR filter. A lower value will increase the frequency resolution, for tighter cuts, at the expense of potentially introducing filter ringing.

FREQUENCY and GAIN

You can adjust an EQ band by clicking on a node and dragging it to change the frequency and gain of the band. You can click and drag to select multiple filter nodes and move them as a group.

Q/BANDWIDTH

If you move the mouse over the bracket handles on the side of the band, you can adjust the Q or bandwidth of the EQ by dragging with the mouse. You can also adjust Q of the selected filter with the mouse wheel.

More Information

Interacting with the EQ nodes

To adjust the EQ curves, grab an EQ node and drag it to a new point on the grid. When a node is selected, handles on either side of it appear which can be dragged together or apart to control the bandwidth of the node. Alternately, you can enter precise EQ settings by typing values into the table below the main EQ grid.

Choosing between Analog and Digital EQ Types

There are reasons why you may want to use one EQ mode over another. Analog mode uses shapes based on analog equalizers, implemented as digital IIR filters. One reason to use this mode is that analog bells are narrower than our digital shapes at high Q values. Since these shapes are minimum-phase, they cause no pre-echos. Also, since they match analog designs, they can be used to emulate what an analog equalizer would do, or reduce damage caused by an analog equalizer.

Digital mode uses special shapes designed to correct audio problems as surgically as possible, implemented as FIR filters. These shapes, unlike the Analog shapes, affect only precisely defined frequency ranges. In Digital mode the EQ is linear-phase, meaning no phase shift will occur.

EQ curve display & EQ Type

You may notice a slight difference in the visual shape of the bell filter when moving between Analog and Digital modes. Analog mode follows the traditional shapes of analog circuitry, while Digital mode uses special filters designed by iZotope. While these shapes are similar between the two modes, they will not be identical.

Composite EQ Curve Display

As you adjust a band you will see two EQ curves. The white curve is the composite of all EQ bands while the curves colored the same shade as the node shows the EQ curve of the selected band.

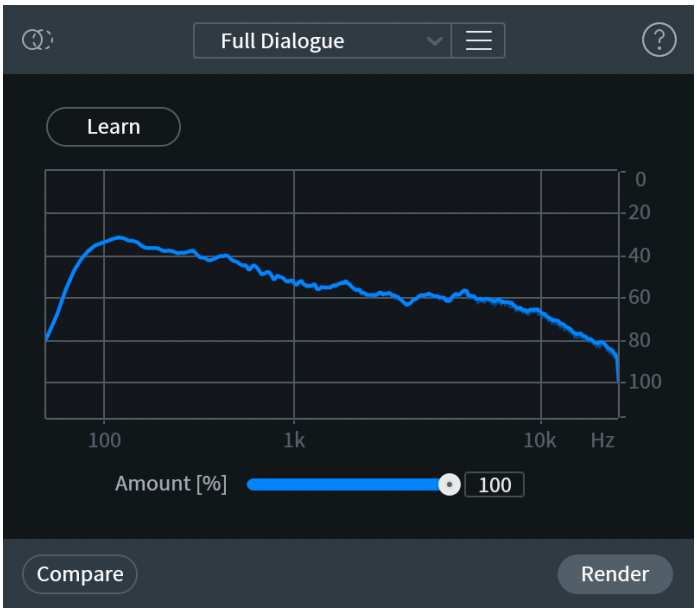
48. EQ Match

ADV

Overview

The EQ Match module lets you match the EQ profile of a selection with the profile of a different selection. This is useful if you're ever tasked with matching a lav mic with a boom mic, matching location dialogue to ADR or vice versa, or perhaps you had multiple mics on an audio source that you'd like more closely aligned in terms of frequency response.

Controls



AMOUNT

Sets how closely the frequency spectrum you're processing will be matched to the learned spectrum. Many times, using a 100% match could sound unnatural, and lower values, between 10-40% do enough to make the two audio signals more closely aligned.

Workflow

How to Apply EQ Match to different selections

1. Open the EQ Match module
2. Make a selection in a file.
3. Click Learn.
4. Make another selection.
5. Click Process.

How to save a captured spectrum as an EQ Match Preset

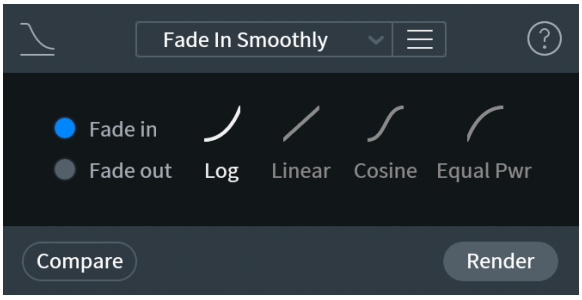
1. In EQ Match, click the preset menu to the right of the preset drop-down menu.
2. Select Add Preset.
3. Enter the name for the new preset.
4. Press Enter.

49. Fade

Overview

The Fade module can be used to apply a gradual increase or decrease to gain throughout your selection. Different amplitude curve options are available in the Fade Type menu. For in-phase tonal material, a Cosine fade curve will work well, while the Equal power fade type can be more effective on noisy material.

Controls



Fade Type

- **Log:** Logarithmic fade for general use
- **Linear:** Linear fade for general use
- **Cosine:** Cosine fade for crossfades of in-phase tonal audio material
- **Equal power:** Equal power fade for crossfades of noisy audio material

Instant Process Tool

STD & ADV

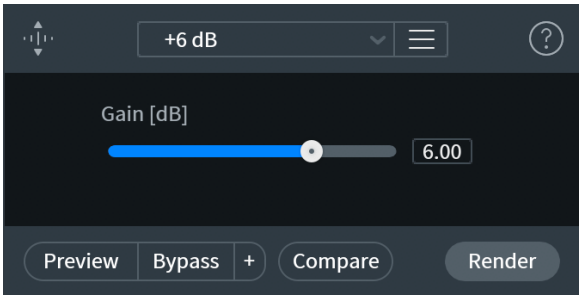
Fade is available as an option in the Instant Process menu. When Fade is selected and Instant Process is enabled, the active settings in the Fade module will be applied to your selection. **For example:** if the Fade module is set to **Fade in: Log** (as shown in the image above) the **Fade in: Log** type of fade will be applied every time you use the Instant Process Tool in 'Fade' mode.

50. Gain

Overview

The gain module is useful for bringing the level of your audio up or down. Gain can also be applied to a specific time-frequency selection, allowing you to manually attenuate or boost selections in the Spectrogram window.

Controls



Gain

Boosts or cuts the level of the signal by the designated decibel amount.

Instant Process Tool

STD & ADV

Gain is available as an option in the Instant Process menu. When Gain is selected and Instant Process is enabled, the active settings in the Gain module will be applied to your selection.

51. Leveler

ADV

Overview

The Leveler module automatically rides the gain in your file to even out the variations of the signal level. The algorithm consists of a compressor with a makeup gain to achieve a smooth signal that's aiming towards (though may not exactly hit) a desired Target RMS level. The compressor has the ability to prevent pumping on speech pauses or breathing sounds, using the Optimization mode, for either Dialogue or Music, in addition to the Ess and Breath parameters.

The level detector stage includes the K-weighting filter that helps equalize the audible loudness, not just RMS level. However, the Leveler module is designed for the smoothing of overall audio signals, rather than taking an entire signal and using a fixed gain to ensure it hits a loudness compliant LKFS level, which is the goal of the Loudness module.

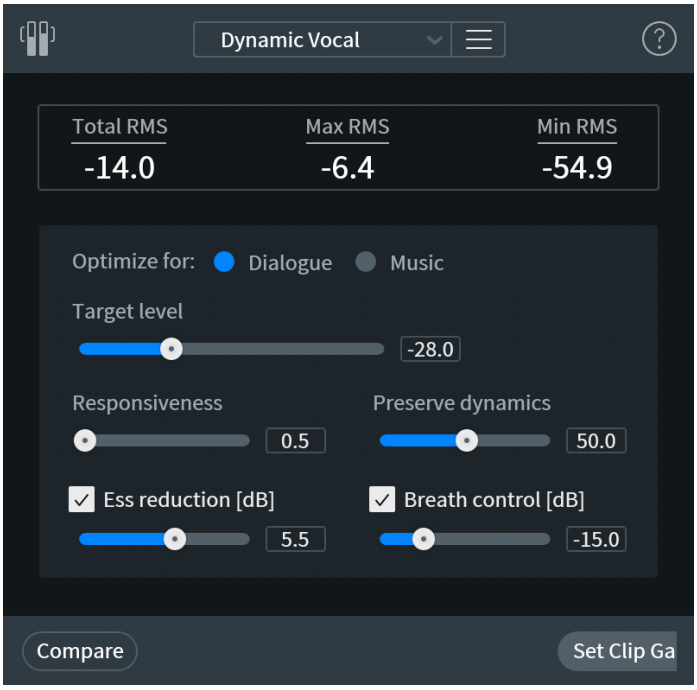
This all combines to create a transparent, non-destructive Clip Gain curve, without the color or artifacts of a traditional compressor.

Unlike the Loudness module, which applies a constant gain based on Loudness compliant analysis to the whole file, Leveler applies a time-variable gain. For convenience of RX users, the time-variable gain is applied as a Clip Gain envelope, which can be viewed and edited by the user.

PLEASE NOTE

If additional processing is applied from another module after running Leveler, the clip gain values assigned by Leveler will be destructively written to the file and the clip gain nodes will return to zero. However, the clip gain settings from Leveler will be saved in the Undo History list).

Controls



Numerical Readouts

Numerical Readouts provide you with the Total, Maximum and Minimum readouts for RMS. The total value is the overall RMS of your audio signal, which may inform where you choose to set the Target RMS level parameter.

Optimize For

Optimize For switches between two modes, Dialogue and Music. Each mode utilizes a slightly different handling of the noise floor.

- Dialogue tends to be audibly juxtaposed against the noise floor, as it's typically very transient, whereas music often tends to fade into the noise floor, with chords, notes, and other instrumental decays.
- Switching between these two modes will affect the behavior of the Leveler, and prevent pumping.

Target Level

Target Level sets the desired average RMS level of the recording.

- Note that Leveler uses K-weighted RMS to better level perceived loudness, but that it is not a loudness compliant leveling tool. It uses the Target Level as a guide, but with the goal of smoothing out variations in an audio signal much more transparently than a compressor typically would. As such, it is not unusual to see the resulting output of Leveler not be an exact 1:1 with the defined Target Level.
- At high target levels, the leveler may not be able to hit the target without clipping, so the target level will not be reached.

Responsiveness

Sets the integration time for RMS level detection and is similar to the attack/release setting on a compressor.

- Lower settings will result in more aggressive Leveling, useful if a signal has a lot of sudden variations.
- Higher settings will result in smoother behavior, leveling words or phrases rather than individual syllables.
- If you find the Leveler is responding to any sudden unwanted sounds, such as a cough, and boosting it, increase the slider to a higher value to see if this results in less aggressive jumps.

Preserve Dynamics

This can be thought of the maximal amount of gain applied by the Leveler. The wider the range of gain adjustments allowed, the further away from the original dynamic range the audio signal will be.

- At lower values, the Leveler will preserve fewer of the original dynamics in the audio signal.
- At higher values, the Leveler will preserve more of the original dynamics in the audio signal.

Ess Reduction

Ess Reduction is aimed at anyone using the Leveler on dialogue or vocals, and utilizes a smart algorithm, inspired by the DBX 902 De-esser, to detect when ess is present in a signal, and then attenuate it accordingly. This avoids adding any boost to esses, which may otherwise be seen as quiet sounds requiring a boost. The slider sets the amount of ess reduction, applied in dB.

Breath Control

Breath Control will automatically detect breaths in your vocal takes and attenuate them. This saves time when editing dialogue or vocal tracks, and streamlines a task that is typically done manually.

- Breath Control automatically analyzes the incoming audio and distinguishes breaths based on their harmonic structure. If any piece of the incoming audio matches a harmonic profile similar to a breath, the Leveler will apply a Clip Gain adjustment.
- Different from a 'Threshold' based process in which the module is only engaged once the audio has risen to a certain volume, this feature will perform its analysis regardless of level.
- This allows for accurate breath recognition with a multitude of quiet or loud dialogue / vocal styles with minimal adjustment of the module's controls.
- The slider represents the desired level, in dB, that you wish all detected breaths to be reduced to. This can result in much more natural sounding breath reduction as the detected breaths in your audio are only reduced when necessary.
- Loud and abrasive breaths will be reduced heavily, and quiet, natural sounding breaths will be left at the same volume. The volume level specified by this slider is a guide, but may not result in exact values.

Limiter

The Leveler has a built-in Limiter in order to avoid introducing any clipping to the audio signal once the Clip Gain envelope has been applied. This cannot be adjusted, but you'll see the Clip Gain envelope smooth off an audio signal if you're pushing peaks close to 0 dB.

52. Loudness Control

Overview

The Loudness Control module adjusts the gain of the signal in order to meet the specified loudness standard, such as BS.1770. Unlike the Leveler module, the gain in the Loudness Control module stays fixed in time. However a post-limiter will be applied to the signal, if it is required to meet the True peak specification.

Controls



Integrated loudness

Sets the desired integrated loudness of the clip, in LKFS (which is the same as LUFS).

Tolerance

Sets the +/- margin of error for the integrated loudness, in LU (loudness units). Each standard determines the tolerance, which is typically set between 0.5 - 2.

Short-term/Momentary Loudness

Short term loudness is calculated over a moving window of 3 seconds. Some loudness standards like EBU R128 s1 Short Term define a maximum short-term value. Momentary loudness is calculated over a moving window of 400ms. The Short-term/Momentary power button enables or disables this processing, and the drop-down allows you to choose between short term or momentary loudness.

True peak

Sets the maximum true peak value allowed in a clip.

Program loudness gate

When computing Integrated loudness, certain low-level signals are excluded from the computation in order to only measure the loudness of "typical" signal levels. This is called gating and is a part of the BS.1770-2, 3 standards (and also R128). Firstly, all signals with a Momentary loudness below −70 LKFS are excluded, so that the noise floor in pauses does not bias the computed loudness. Secondly, all signals with a lower Momentary loudness than 10 dB (or LU) below the average Momentary loudness are also excluded. This allows the computation of the Integrated loudness only from the typical and loudest parts of the program, while excluding silence and quieter parts, which have less effect on the listener's impression of overall loudness of the program . Most of the loudness standards shipping with RX Loudness Control use gating, an exception being BS.1770-1. This parameter allows you to toggle gating on or off, if you're working to a unique requirement.

Workflow

To use the Loudness Control module:

1. Make a selection in your file. If you want to apply Loudness Control to the entire clip, use ctrl/cmd + A or click Edit > Select All.
2. Select a preset, or make adjustments to the controls manually.
3. Click Render.

Loudness Standards

Loudness Standard	Integrated	Integrated Tolerance (+/-)	Short-term	Momentary	True Peak	Gating	Notes
AES AGOTTVS TD1006.1.17-10	-16	2.0	Off	Off	-1	On	
AGCOM 219/09/CSP	-24	0.5	Off	Off	-2	On	Italy
ARIB TR-B32 A/85	-24	2	Off	Off	-1	On	Japan
ATSC A/85	-24	2	Off	Off	-2	On	USA, Canada, Puerto Rico
BS.1770-1	-24	2	Off	Off	-2	Off	International
BS.1770-2/3/4	-24	2	Off	Off	-2	On	International
EBU R128	-23	0.5	Off	Off	-1	On	Europe
EBU R128 DPP	-23	0.5	Off	Off	-3	On	Europe
EBU R128 (South Africa)	-23	1	Off	Off	-2	On	South Africa
OP-59	-24	1	Off	Off	-2	On	Australia/New Zealand
Portaria 354	-23	0.5	Off	Off	-2	On	Brazil

53. Mixing

Overview

Provides specific control over both left and right signal and balance levels. This simple operation can be used to downmix stereo material into mono, invert waveforms, transcode left/right stereo into mid/side, subtract a center channel, and much more.

Controls



Left Output Mix (%)

Allows you to define how much of the current selection's left and right channel signal will be present in the new target left channel.

Right Output Mix (%)

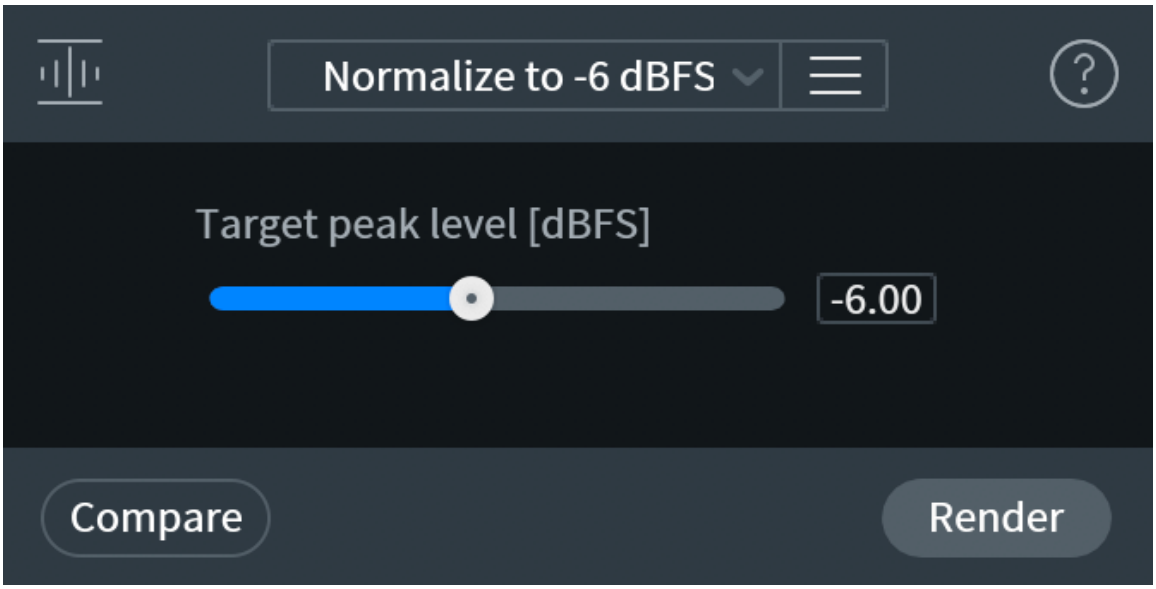
Allows you to define how much of the current selection's left and right signal will be present in the new target right channel.

54. Normalize

Overview

The Normalize module applies enough gain to set the sample peak level of your signal to the specified Target Peak Level.

Target Peak Level



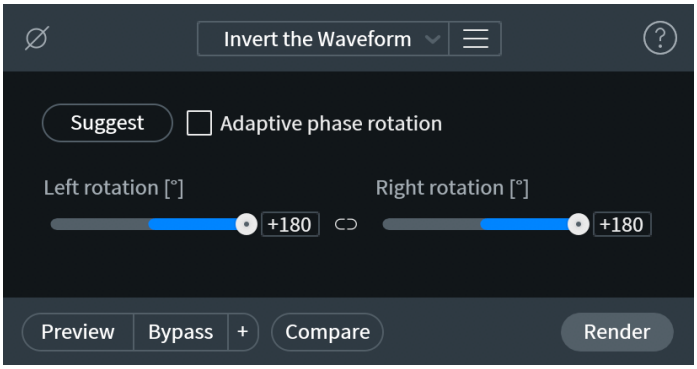
Target Peak Level [dBFS]: Determines the maximum peak level of a signal as a result of normalization.

55. Phase

Overview

The Phase module balances asymmetric waveforms by rotating signal phase. Rotating the phase of a signal changes its peak values but doesn't change its loudness, and otherwise has no audible effect on the signal.

Controls



Adaptive Phase Rotation

Continuously analyzes the audio selection and applies the time-variable phase rotation to both left and right channels, resulting in a symmetrical waveform with minimal signal peak levels.

Adaptive phase rotation is best used on vocal material, as it can occasionally yield pitch artifacts on musical material.

Rotation (deg)

Rotates the channel's phase by the specified degree.

When a waveform's phase is rotated, every frequency is rotated equally. Rotating phase by 180 degrees inverts the waveform.

Suggest

Analyzes the selection for the ideal channel-linked fixed phase for reducing overall peak levels of the signal.

More Information

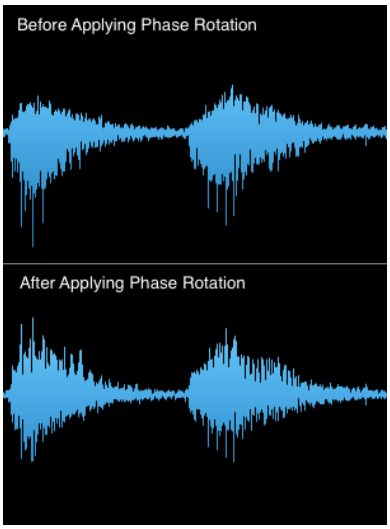
Asymmetric waveforms can occasionally occur in audio such as dialogue, voice, and brass instruments.

- Making the waveform more symmetrical gives the signal more headroom.
- Rotating the phase of a waveform, will change its amplitude characteristics. Phase rotation does not result in a time shift.

- Because the range of rotation is from −180 to +180 degrees, the Phase tool can be used for simpler purposes, such as inverting signal polarity.

Visual Example of Phase Rotation

The top waveform in the following image is a trumpet signal with higher peak values on one side of its waveform (meaning the waveform is asymmetrical.) The bottom waveform in the following image has been processed by the Phase module. The processig rotated the phase of the waveform by −72 degrees to distribute its peak samples more evenly (making the waveform more symmetrical.)



56. Plug-in Hosting

Overview

On Windows, RX supports VST3 and VST2 plug-ins. On Intel-based macs and Apple silicon-based macs running RX with Rosetta, RX supports AU, VST3, and VST2 plug-ins. For macs running RX natively on Apple silicon, AU and VST3 are supported. VST2 is not supported when RX is running natively on Apple silicon.

UPDATE YOUR PLUG-INS

Please make sure your third-party plug-ins are updated to their latest version in order to ensure that RX will be able to scan them correctly. Contact your plug-in manufacturers for updated installers, if necessary.



Selection Based Processing

With a plug-in loaded in the Plug-in module, you can make use of the same audio selection tools and Preview and Compare options that are available for other RX modules.

This can allow for very detailed processing and greater accuracy when working with your existing plug-ins, giving you audio selection options unavailable in a traditional DAW setup.

Plug-in Presets

Only one plug-in may be loaded at a time. However, with the use of presets, multiple settings and presets may be recalled quickly in order to move between plug-in instances. When your plug-in is configured the way you would like it, select Add Preset from the small Preset drop down arrow.

Assign Presets to Keyboard Shortcuts

Once your preset is named and saved, you can then assign that plug-in and preset to a keyboard shortcut with the Set Preset Shortcut feature.

This keyboard shortcut will recall not only the plug-in settings, but the plug-in instance itself. As such, if a preset is saved with Plug-in 1, and Plug-in 2 is currently loaded into RX's plug-in window, pressing the preset keyboard shortcut will re-instantiate Plug-in 1 and recall the exact settings when the preset was made.

This can allow for very quick editing, processing, and recall of plug-in instances and settings, providing a quicker workflow than traditional DAW track/mixer environments.

57. Resample

Module Only

Overview

The Resample module allows you to convert an audio file from one sampling rate to another.

Sample Rate Conversion (SRC) is a necessary process when converting material from one sampling rate (such as studio-quality 96 kHz or 192 kHz) to another rate (such as 44.1 kHz for CD or 48 kHz for video).

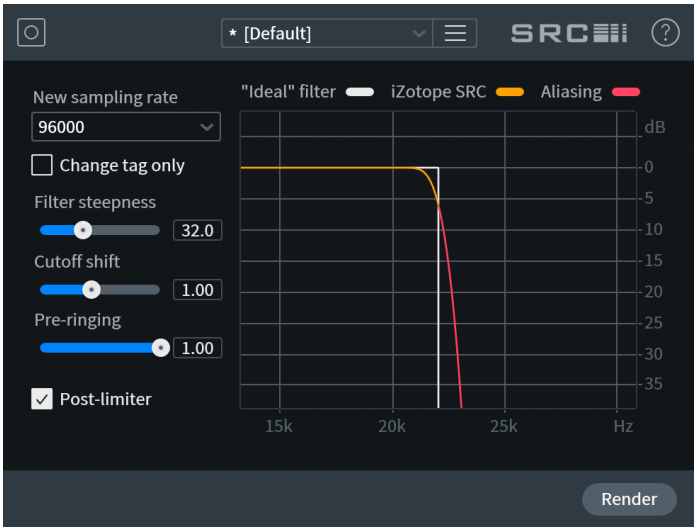
It is common to record and edit in high sampling rates since higher rates allow higher frequencies to be represented. For example, a 192 kHz audio sample can represent frequencies up to 96 kHz whereas a 44.1 kHz audio sample can only represent frequencies up to 22.05 kHz. The highest frequency that can be represented accurately by a sampling rate is half of the sampling rate, and is known as the Nyquist frequency.

When reducing the sampling rate, or downsampling, it is crucial to remove the frequencies that cannot be represented at the lower sampling rate. Leaving frequencies above this point causes aliasing. Aliasing can be heard as the frequencies in an inaudible range are shifted into an audible range, causing distortion and noise. With iZotope SRC's steep low-pass filter, users can completely avoid the common aliasing artifacts while maintaining the maximum frequency content.

Comparison of iZotope's SRC process

A comparison of iZotope's SRC process versus other sample rate convertors can be viewed at: http://src.infinitewave.ca/.

Controls



NOTE ABOUT THE RED ALIASING CURVE DISPLAY

The Aliasing portion of the curve displayed in red shows the reflected frequencies during downsampling or imaged frequencies during upsampling — both due to aliasing.

New sampling rate

This setting chooses the sampling rate you want to convert to. Choose a sampling rate from the drop-down list, or click on the field to type in a custom sampling rate.

Change tag only

Changes the declared sampling rate of the file in the file's properties without resampling the file, effectively changing the playback rate and pitch of the file.

WHEN TO USE CHANGE TAG ONLY OPTION

This feature is useful if the sampling rate tag was damaged by a previous audio editing process and the file is playing back incorrectly.

Filter steepness

This allows you to control the steepness of the SRC filter cutoff. The white line is representative of an ideal low-pass filter.

NOTE ABOUT HIGHER FILTER STEEPNESS VALUES

Higher filter steepness means better frequency performance of the filter: wider passband retains more useful signal, while stronger stopband attenuation provides better rejection of aliasing. At the same time, higher steepness of the frequency response requires a longer filter, which produces more ringing in time domain and energy smearing near the cutoff frequency.

Cutoff shift

SRC filter cutoff frequency shift (scaling multiplier). Allows shifting the filter cutoff frequency up or down, to balance the width of a passband vs. amount of aliasing.

Pre-ringing

SRC filter pre-ringing amount in time domain (0 for minimum phase, 1 for linear phase, or anywhere in between). Adjusts the phase response of the filter, which affects its time-domain ringing characteristic. The value of 0 produces a minimum-phase filter, which has no pre-ringing, but maximal post-ringing. The value of 1 produces a linear-phase filter with a symmetric impulse response: the amount of pre-ringing is equal to the amount of post-ringing. Intermediate values between 0 and 1 produce so-called intermediate-phase filters that balance pre- and post-ringing while maintaining linear-phase response across a possibly wider range of frequencies.

Post-limiter

Keeps true peak levels of the output signal below 0 dBTP to prevent any clipping from occurring.

When to use the post-limiter:

- This option is important when resampling signals that are very close to 0 dB, because filtering during resampling can change peak levels of a signal.

- Engage the Post-limiter option in order to limit the output levels of your signal to prevent any clipping from occurring.

58. Signal Generator

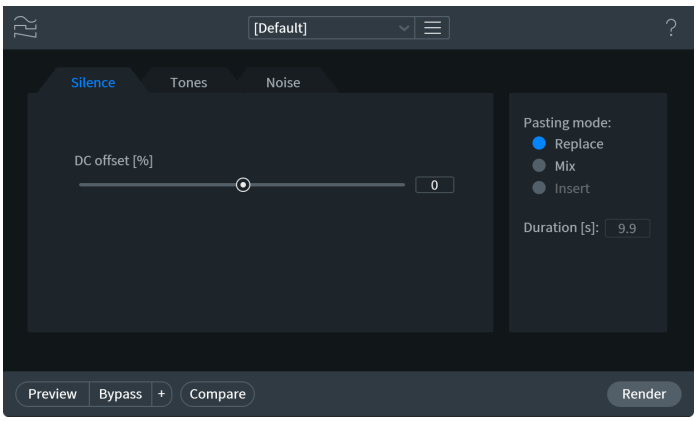
Overview

The Signal Generator is able to synthesize silence, tones, and noise. This is useful for creating test tones, calibration tones for post production delivery specs, repairing DC offset, and even using it as a “bleep” module to eliminate obscenities in a dialogue edit.

The Signal Generator module is capable of generating extremely accurate test tones for research and testing purposes.

Silence

Silence generates digital silence, which can be used to adjust spacing or duration of your file.

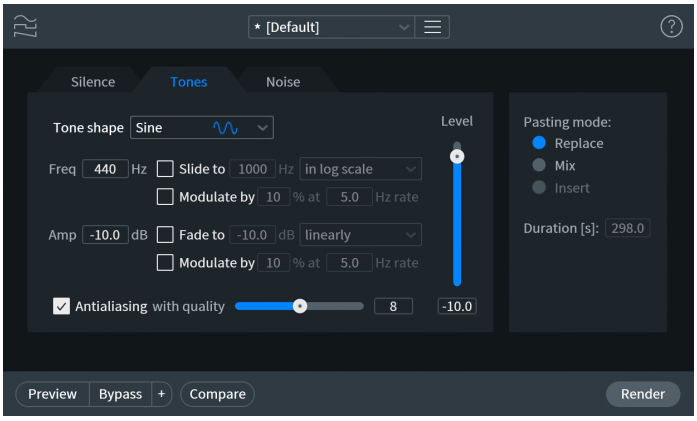


DC Offset

Creates or removes positive or negative amounts of DC offset.

Tones

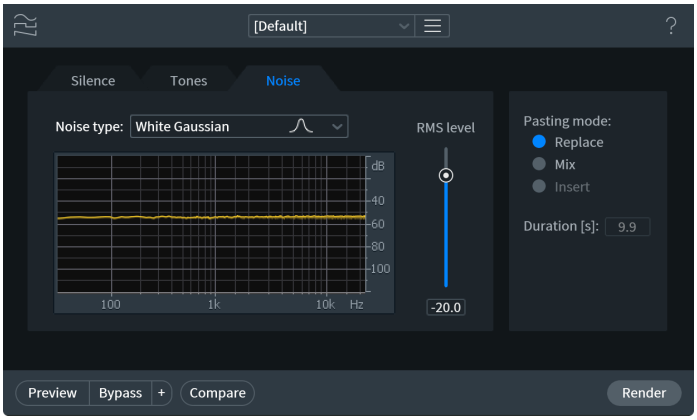
The Tones tab allows you to generate signals.



- **TONE SHAPE:** Tone shape allows you to choose the waveform for the signal to be generated, with a choice of: Sine, Triangle, Sawtooth, Square.

- **FREQ:** Sets the frequency, in Hertz (Hz), of the signal to be generated.
- **AMP:** Sets the amplitude, in Decibels (dB), of the signal to be generated.
- **ANTIALIASING:** All test tone shapes, except the sinusoid, contain an infinite set of harmonics. When Antialiasing is not selected, these tones are generated in a “naive” way in the time domain. It leads to aliasing (folding) of higher harmonics into the audible range, which often sounds undesirable.
  - When Antialiasing option is enabled, these higher harmonics are filtered out with a linear-phase low-pass filter. It prevents aliasing, but introduces some ringing into the waveform, which is usually not a problem, because it is concentrated near half the sampling rate.
  - Higher quality of antialiasing produces a filter with a sharper frequency cut, which retains more of the useful harmonics and rejects more of the aliased harmonics, at the expense of more time-domain ringing and slower processing.
- **LEVEL:** Specifies the sample peak level of the synthesized waveform in dBFS. When Antialiasing is used, actual sample and true peak levels may exceed the specified peak level due to filter ringing.

Noise



Allows you to generate noise of many different types, with a user definable RMS level. The noise types include: White Gaussian, White triangular, White uniform, White binary, Pink, Brown.

Noise color defines the spectral shape: white noise has flat power spectrum, pink noise decays at 3 dB/octave rate, while brown noise decays at 6 dB/octave rate.

Gaussian, triangular, uniform, or binary relates to the p.d.f., or probability distribution function, of the noise. It describes how often samples of different amplitude are encountered in the signal. For example, uniform p.d.f. means that all amplitudes below noise peak level are equally popular.

- **PASTING MODE:** The pasting modes affect how generated signals are added to your file.
  - **REPLACE:** Replace will completely replace the audio in your selection with the generated signal.
  - **Mix:** Mix will retain the audio in your selection, and mix in the generated signal.
  - **Insert:** Insert will allow you to set a custom duration, in seconds, and the Signal Generator will then insert this audio, increasing the length of your audio file by the amount prescribed.

59. Time & Pitch

Overview

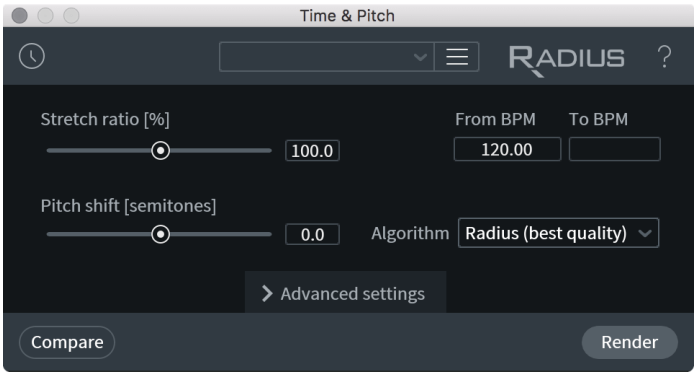
The Time & Pitch module uses the iZotope Radius™ processing algorithm to allow for independent control over the length and pitch of your audio. It is useful for retuning audio to fit in a mix better, or adjusting the length of a selection to account for changes in BPM or timecode.

- The Variable Pitch module can be used for faster pitch shifting with the ability to correct variations in pitch over time.
- The Variable Time module can be used to adjust the time stretch ratio envelope over time.

iZotope Radius

iZotope Radius™ is a world-class time-stretching and pitch-shifting algorithm. You can easily change the pitch of a single instrument, voice, or entire ensemble while preserving the timing and acoustic space of the original recording. iZotope Radius is designed to preserve the natural timbral qualities of the original file, even when applying extreme pitch shifts.

Controls



Algorithm

The Algorithm drop-down menu has three options:

Radius

Designed to work well with polyphonic material such as mixes with more than one instrument, as well as non-harmonic material such as drum loops or rhythmic audio. This is the highest-quality option for most sources.

Radius RT

Good quality, polyphonic, but faster than Radius. If processing speed is important, use the Radius RT algorithm.

Solo Instrument

Designed for monophonic pitched material such as a stringed instrument or human voice. You should use Solo mode only when processing a single instrument with a clearly defined pitch. The human voice is a good candidate for solo mode, as are most stringed instruments, brass instruments, and woodwinds.

In Solo mode, the adaptive window size can significantly affect the quality of Radius's output.

- If the adaptive window size is too small, you will hear a squeaking noise which sounds like the pitch of the audio is changing very rapidly.
- If the adaptive window size is too large then the sound will become grainy as you will begin to hear portions of it being repeated.
- A good approach is to start with the default window size of 37 ms. If the results are unsatisfactory, increase the window size until the squeaking noise described above does not occur. If you cannot get the distortion to disappear, switch to Radius mode for processing.
- Lower pitched instruments and voices may require a longer adaptive window size than the default, but very long adaptive window sizes can cause audible repeating slices of audio.

Formant Correction

Formants are the resonant frequency components of voice that tend to be perceived as characteristics like age and gender. You can shift formants independently of pitch and time by enabling Shift Formants.

- Typically you will leave the Formant Shift Strength set to 1 (full strength) and the Formant Shift Semitones set to 0.
- If you hear what sounds like an EQ adjustment to your audio, you can try lowering the strength to reduce this artifact.
- To achieve special effects, for example to change the perceived gender of a human voice, try adjusting the semitones to a value other than 0.

Stretch Ratio

Determines how much the resulting audio will be stretched in time.

- Values between 12.5% and 100% will cause the audio to speed up without affecting pitch, resulting in a shorter audio file.
- Values between 100% and 800% will cause the audio to slow down without affecting pitch, giving you a longer audio file.

BPM Calculator

If you are using Radius to process audio for a tempo change, you can also adjust the stretch ratio with the BPM Calculator.

Pitch Shift

Controls the amount of pitch shifting up or down that will be applied to the audio.

Transient Sensitivity

Determines the algorithm's handling of transient material. Higher values will result in better preservation of individual transients after processing.

- When stretching percussive material, you usually want transient sensitivity set to its default value of 1.
- If transients in your audio are being “smeared”, a higher value of 2 will tighten up transience at the expense of incurring heavier processing on non-transient audio.
- Bowed instruments such as the violin and cello are especially affected by the transient sensitivity setting. If you hear a stuttering artifact, lower the transient sensitivity to eliminate it.

Noise Generation (Radius Mode Only)

Helps noisy material (like sibilance or snare drums) sound more natural when processed.

NOTE

This control will generate noise instead of stretching the noise that is already present in the signal and creating new tones. Higher values of the noise generation parameter will cause Radius to generate noise more often, but can cause some phase artifacts.

Pitch Coherence (Radius Mode Only)

The Pitch coherence control in the Radius control panel helps preserve the natural timbre for pitched solo voices, such as human speech, saxophone or vocals. While traditional vocoders can smear these signals in time and randomize phase, the pitch coherence parameter of Radius preserves phase coherence for these signals.

High values of pitch coherence will avoid phasiness in Radius’s output at the expense of roughness (modulation) in processed polyphonic recordings. Try turning this up for better results if you’re processing a solo voice or a small group of related instruments.

Phase Coherence (Mix Mode Only)

Preserves the coherence of phase between the left and right channels of the processed audio.

This should be increased if there’s any change in the perceived stereo image after using Radius. It can be decreased when processing a multichannel signal where different channels contain completely different instruments.

Adaptive Window Size (ms) (Solo Mode Only)

Adjusts the window size in milliseconds of Radius’ Solo algorithm.

- If the adaptive window size is too small, you will hear a squeaking noise which sounds like the pitch of the audio is changing very rapidly.
- If the adaptive window size is too large then the sound will become grainy as you will begin to hear portions of it being repeated.
- Increase this if you have trouble getting good results pitching or stretching low-pitched instruments or voices.

Shift Formants

Processes formant frequencies independently of other pitch and time processing.

- When this option is enabled, formant frequencies can be shifted independently of other pitch shifting performed by Radius.
- When Radius performs pitch-shifting without Formant Correction, it will shift these resonant frequencies along with the rest of the audio.

Strength

Adjusts the amplitude strength of the formant correction filter.

Shift

How much formant frequencies are shifted. Typically this control can be set to 0, which leaves the formant frequencies unshifted during processing. Adjust this control to fine-tune the formant correction algorithm or for special effects.

Width

Controls the bandwidth of the formant detection filter.

- Smaller values of this control will offer more precise formant correction in the processed audio.
- Higher values will include a wider band of formant frequencies.

More Information

Pitch shifting single instruments (especially bass instruments) can benefit from some adjustments to formant correction. Try enabling formant correction and moving the strength between 0.1 and 0.2. Move the Formant Correction semitones part of the way towards your pitch shift amount. For example, if you're pitch shifting +4 semitones, move the Formant Correction Semitones between 2 and 3. This can help bring back subtle percussive elements in the original source material.

The formant frequencies of the human voice can actually shift slightly when we sing. You can use the Formant Correction Semitones control to compensate for this. For example, if pitch shifting a human voice by +7 semitones, try setting the Formant Correction semitones between 0 and +2 for more natural results.

60. Variable Pitch

Overview

Variable Pitch can be used to quickly fix small pitch variations or to correct gradual pitch drifts over time. With the addition of Preserve Time mode, Variable Pitch replaces and expands on the processing options available in the Pitch Contour module available in previous versions of RX.

Displays

Variable Pitch features a waveform panel and a spectrogram panel that each display information about the current selection in the active file tab. These panels will dynamically update when the selection is changed. If no selection is made in the active file tab, no information will be displayed in the spectrogram or waveform panels.



WINDOW RESIZING

Click and drag on the bottom right-hand corner of the module window to customize the window size.

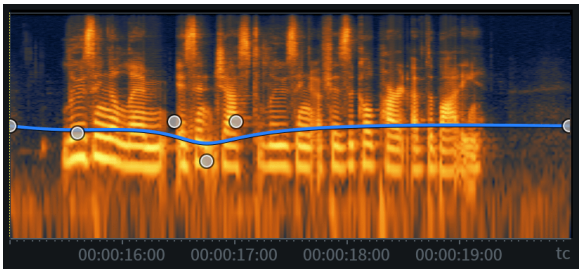
Waveform Display

The single waveform drawn in this panel represents the sum of all enabled channels in the current selection. The waveform drawing is normalized to allow for consistent vertical resolution when working with selections of varying amplitude.



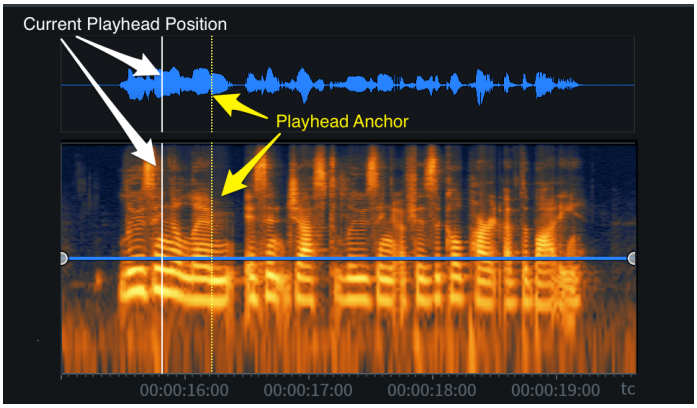
Spectrogram Display

The spectrogram drawn in this panel represents the sum of all enabled channels in the current selection.



Playhead Indicators

The solid white vertical line and dotted yellow vertical line overlaid on the waveform and spectrogram panels indicate the current playhead position (white) and the playhead anchor position (yellow).



Current Playhead Position

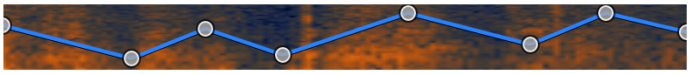
The solid white vertical line overlaid on the waveform and spectrogram panels indicates the current playhead position. This indicator line updates to follow the current playhead position during playback. The playhead position indicator will only appear in the module window when it is within the bounds of the current selection.

Playhead Anchor

The dotted yellow vertical line overlaid on the waveform and spectrogram panels indicates the playhead anchor position in the main editor window. If the playhead anchor position is outside of the current selection bounds, the indicator will not be displayed in the module window.

Contour Curve Display

The blue line overlaid on the spectrogram panel represents the pitch contour curve. Nodes can be added to this curve and adjusted to make changes to pitch over the course of the active selection.



Contour Curve Axes

The contour curve allows for adjustments along two axes: Pitch and Time.

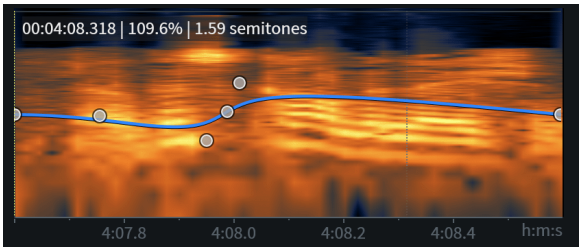
- **Pitch:** The vertical y-axis of the contour curve represents pitch in semitones.
  - The Pitch axis ranges from -24 (bottom) to +24 (top) semitones.
  - The center of the Pitch axis equates to 0 semitones.
- **Time:** The horizontal x-axis represents time.
  - The time format used here is determined by the Time Format Display selection in the transport section of the main editor window.
  - The range of the time ruler matches the length of the current selection.

To zoom the ruler in and out:

1. Hover over the ruler and use a mousewheel or trackpad to zoom in and out.
2. Click and drag left or right on the ruler when zoomed in to change the ruler position.
3. Double-click on a ruler display to reset the zoom level to default.

Contour Curve Readout

When the cursor is positioned over the spectrogram panel, a text readout will appear in the upper left hand corner of the panel. This readout displays information about the processing that will be applied by the current contour curve.



The readout displays the following information about the current cursor position, from left to right:

- **Time**: Current time position of the cursor within the spectrogram panel.
- **Pitch Shift (%)**: Percentage of pitch shift that will be applied at the cursor's current time position.
- **Pitch Shift (Semitones)**: Amount of pitch shift that will be applied at the cursor's current time position.

Contour Curve Editing

The following section describes the methods and controls available for editing the contour curve.

Add Nodes

Click in the spectrogram panel to add a new node to the contour curve.

> ⚠ **CONTOUR CURVE NODE LIMIT**
> The contour curve supports adding up to 25 nodes.

Semitone Adjustments

Click and drag a node up or down to adjust its semitone value.

Time Adjustments

Click and drag a node left or right to move the associated pitch adjustment point earlier or later in time.

- Nodes cannot be moved outside of the time bounds of the current selection.
- The contour curve shape will be maintained when the selection changes.
- The contour curve shape will be maintained after rendering.

Remove Nodes

Individual nodes can be deleted from the curve using the following methods:

▶ Click and drag a node past the top or bottom border of the contour curve display to quickly remove it from the curve.

▶ Control-click (Mac) or ctrl-click (Windows) on a node to remove it from the curve.

Reset Individual Nodes

Double-click on a node to reset it to the default value of 0 semitones. Double-clicking a node only resets the semitone value to default, it will not change the time position of the node.

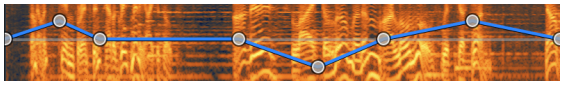
Reset Curve

Removes all custom nodes from the curve, resetting it to default. Two nodes are present in the default curve, one at the start and one at the end of the current selection. The default nodes are set to 0 semitones (no pitch adjustment).

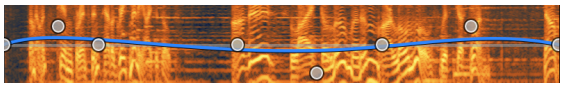
Smoothing

Adjusts the amount of smoothing applied between nodes on the contour curve. Smoothing is a global control and is applied to all nodes on the curve.

- **Lower smoothing values:** Applies little to no smoothing between nodes on the curve, resulting in strict transitions between nodes.



- **Higher smoothing values:** Applies more smoothing between nodes on the curve, resulting in a gradual, rounded slope between nodes.



Controls

The following section describes the Preserve Time control and the parameters available for refining the results of processing when Preserve Time (Radius Mode) is enabled.

Preserve Time

Determines the pitch processing algorithm that is used when rendering the current selection. When Preserve Time is enabled, the iZotope Radius™ processing algorithm is used. When Preserve Time is disabled, a resampling processing algorithm is used. See the descriptions below for more information about these processing algorithms and how they affect the rendered output.

Preserve Time Disabled (Resampling Mode)

When Preserve Time is disabled, a resampling algorithm is used to render the pitch contour curve. In this mode, time and pitch are adjusted synchronously, as if tape playback speed was adjusted. **The length of the rendered selection will change to account for changes in pitch.**

- Negative semitone adjustments *increase* the length of the rendered selection.
- Positive semitone adjustments *decrease* the length of the rendered selection.

> **USE CASE FOR RESAMPLING MODE**
>
> This mode can be useful for synchronizing two recordings of the same performance that have been captured at slightly different clock speeds.

Preserve Time Enabled (Radius Mode)

When Preserve Time is enabled, the Radius algorithm is used to render the pitch contour curve. **The mode will adjust pitch without altering the length of the current selection.**

> **PRESERVE TIME: ADDITIONAL CONTROLS**
>
> When Preserve Time mode is enabled, additional controls become available for refining the results of processing: Pitch Coherence & Transient Sensitivity.

Pitch Coherence

Adjusts the amount of timbre preservation applied during processing. This control is useful for maintaining the natural timbral qualities of pitched audio content (such as human speech, sung vocals or saxophone) by correcting the smearing or phasing issues that can be introduced during processing.

Increasing pitch coherence can help to combat phase effects that can occur when applying processing to a solo voice or a small group of related instruments. A possible tradeoff of increasing pitch coherence is that it may introduce roughness or unwanted modulation when applying processing to polyphonic or multi-instrument content.

NOTE

Pitch Coherence is only available for adjustment when Preserve Time is enabled.

Transient Sensitivity

Adjusts how the Radius algorithm detects and subsequently preserves transient content when processing. Higher values can improve transient clarity at the cost of applying heavier processing to non-transient material. Bowed instruments, such as cello or violin, may present stuttering artifacts after processing with higher transient sensitivity values. Lowering the transient sensitivity value can help to reduce the unwanted artifacts in sustained content.

TRANSIENT SENSITIVITY CONTROL AVAILABILITY

Transient Sensitivity is only available for adjustment when Preserve Time mode is ON.

Compare Settings

The Variable Pitch module does not allow for real-time preview playback. To audition different settings before rendering, click the Compare button in the module footer area to send the settings to the Compare Settings window.



Learn more about **Compare Settings** in the Common Module Controls.

Alternative Modules

- For pitch envelope editing tailored to speech, try using the Dialogue Contour module in RX 11 Advanced.
- Make static time-stretch and pitch-shift adjustments with the Time & Pitch module.
- Make adjustments to the time stretch ratio envelope of a selection with the Variable Time module.

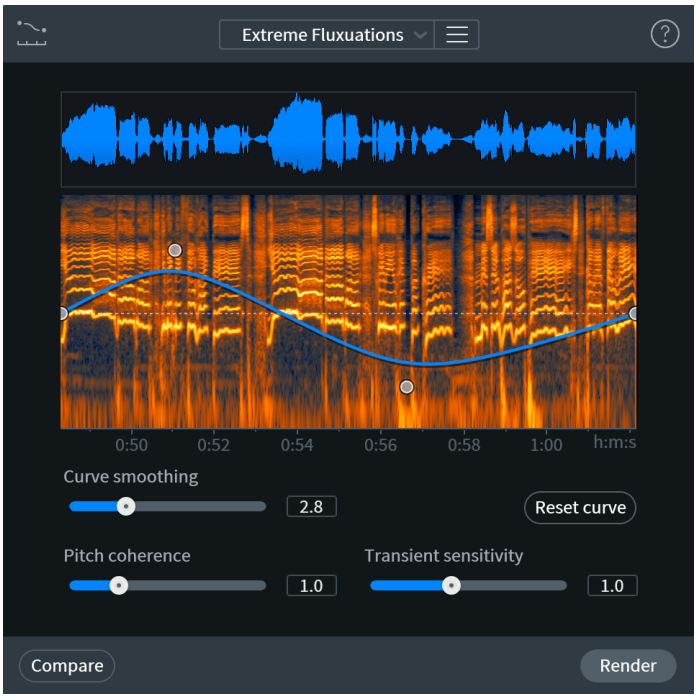
61. Variable Time

Overview

Variable Time can be used to correct timing issues or creatively adjust speed over the course of a selection. This module leverages the iZotope Radius algorithm to allow for high-quality time compression and expansion without also adjusting pitch.

Displays

Variable Time features a waveform panel and a spectrogram panel that each display information about the current selection in the active file tab. These panels will dynamically update when the selection is changed. If no selection is made in the active file tab, no information will be displayed in the spectrogram or waveform panels.



WINDOW RESIZING

Click and drag on the bottom right-hand corner of the module window to customize the window size.

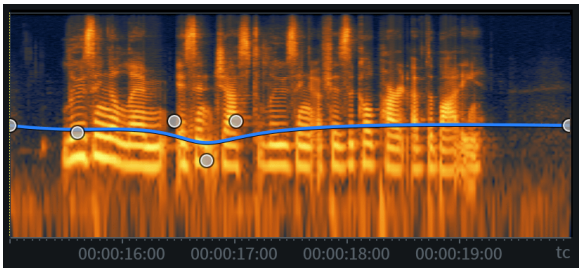
Waveform Display

The single waveform drawn in this panel represents the sum of all enabled channels in the current selection. The waveform drawing is normalized to allow for consistent vertical resolution when working with selections of varying amplitude.



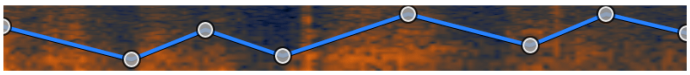
Spectrogram Display

The spectrogram drawn in this panel represents the sum of all enabled channels in the current selection.



Contour Curve Display

The blue line overlaid on the spectrogram panel represents the time stretch ratio contour curve. Nodes can be added to this curve and adjusted to make changes to time stretch ratio over the course of the active selection.



Contour Curve Axes

The contour curve allows for adjustments along two axes: Time Stretch Ratio (y-axis) and Time (x-axis).

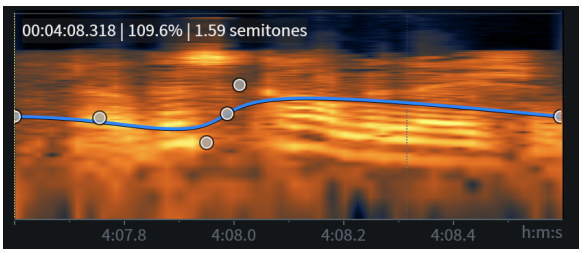
- **Time Stretch Ratio:** The vertical y-axis of the contour curve represents time stretch ratio values in increments associated with both Speed (multiplier) and Length (percentage) adjustments. The effect of adjustments along this axis can be thought of as affecting the processed signal in two ways, adjusting length or adjusting speed. The minimum and maximum values of this axis equate to the following speed and length adjustments:
  - **Minimum value:**
    - **Speed**: 0.125x - 0.125 times *slower* than the original
    - **Length**: 800.0% - 8 times *longer* than the original
  - **Maximum value:**
    - **Speed**: 8.000x - 8 times *faster* than the original
    - **Length**: 12.5% - 1/8th the length of the original; 8 times *shorter.*
  - The center of the y-axis is equivalent to no adjustment in length or speed.
- **Time (Current Selection Length):** The horizontal x-axis represents time across the current selection.
  - The time format used here is determined by the Time Format Display selection in the transport section of the main editor window.
  - The range of the time ruler matches the length of the current selection.

To zoom the ruler in and out:

1. Hover over the ruler and use a mousewheel or trackpad to zoom in and out.
2. Click and drag left or right on the ruler when zoomed in to change the ruler position.
3. Double-click on a ruler display to reset the zoom level to default.

Contour Curve Readout

When the cursor is positioned over the spectrogram panel, a text readout will appear in the upper left hand corner of the panel. This readout displays information about the processing that will be applied by the current contour curve when it is rendered.



The readout displays the following information about the current cursor position, from left to right:

- **Time**: Current time position of the cursor within the spectrogram panel.
- **Length Adjustment (%)**: Percentage of length adjustment that will be applied at the cursor's current time position.
- **Speed Adjustment (x - multiplier)**: Amount of speed adjustment that will be applied at the cursor's current time position.

Contour Curve Editing

The following section describes the methods and controls available for editing the contour curve.

Add Nodes

Click in the spectrogram panel to add a new node to the contour curve.

> ⚠ **CONTOUR CURVE NODE LIMIT**
> The contour curve supports adding up to 25 nodes.

Time Stretch Ratio Adjustments

Click and drag a node up or down to adjust its time stretch ratio value.

Time Adjustments

Click and drag a node left or right to move the associated time stretch ratio adjustment earlier or later in time.

- Nodes cannot be moved outside of the time bounds of the current selection.
- The contour curve shape will be maintained when the selection changes.

- The contour curve shape will be maintained after rendering.

Remove Nodes

Individual nodes can be deleted from the curve using the following methods:

- ▶ Click and drag a node past the top or bottom border of the contour curve display to quickly remove it from the curve.
- ▶ Control-click (Mac) or ctrl-click (Windows) on a node to remove it from the curve.

Reset Individual Nodes

Double-click on a node to reset it to the default value of 100%/1.000x (no change to time stretch ratio). Double-clicking on a node will only reset its time stretch ratio value to default, it will not change the time position of the node within the selection.

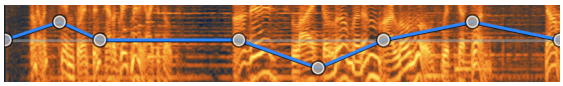
Reset Curve

Removes all custom nodes from the curve, resetting it to default. Two nodes are present in the default curve, one at the start and one at the end of the current selection. The default nodes are set to 100%/1.000x (no time stretch ratio adjustment).

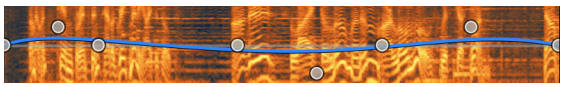
Smoothing

Adjusts the amount of smoothing applied between nodes on the contour curve. Smoothing is a global control and is applied to all nodes on the curve.

- **Lower smoothing values:** Applies little to no smoothing between nodes on the curve, resulting in strict transitions between nodes.



- **Higher smoothing values:** Applies more smoothing between nodes on the curve, resulting in a gradual, rounded slope between nodes.



Controls

The following section describes the controls available for refining the Radius algorithm processing results.

Pitch Coherence

Adjusts the amount of timbre preservation applied during processing. This control is useful for maintaining the natural timbral qualities of pitched audio content (such as human speech, sung vocals or saxophone) by correcting the smearing or phasing issues that can be introduced during processing.

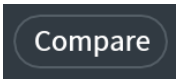
Increasing pitch coherence can help to combat phase effects that can occur when applying processing to a solo voice or a small group of related instruments. A possible tradeoff of increasing this control is that it may introduce roughness or unwanted modulation when processing polyphonic or multi-instrument content.

Transient Sensitivity

Adjusts how the Radius algorithm detects and subsequently preserves transient content when processing. Higher values can improve transient clarity at the cost of applying heavier processing to non-transient material. Bowed instruments, such as cello or violin, may present stuttering artifacts after processing with higher transient sensitivity values. Lowering the transient sensitivity value can help to reduce the unwanted artifacts in sustained content.

Compare Settings

The Variable Time module does not allow for real-time preview playback. To audition different settings before rendering, click the Compare button in the module footer area to send the settings to the Compare Settings window.



Learn more about Compare Settings in the Common Module Controls chapter.

Alternative Modules

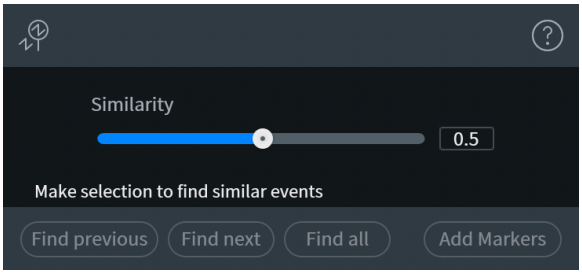
- Make static time-stretch and pitch-shift adjustments with the Time & Pitch module.
- Adjust the pitch envelope over the course of a selection with the Variable Pitch module.
- For pitch envelope editing tailored to speech, try using the Dialogue Contour module in RX 11 Advanced.

62. Find Similar

Overview

There are times when you may have many similar audio events to manually repair. In cases like this, manually selecting and processing each event can be time consuming. RX includes a Find Similar Event tool which takes your selection and finds all related instances of that audio event. You can choose from Find Next, Find Previous, and Find All.

Controls



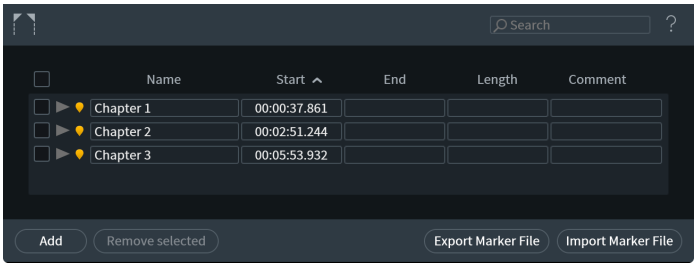
- **Similarity:** Lower values will find more events. The higher the value, the more similar an event must be to the original event selection for it to be detected.
- **Find Previous:** Searches before the current selection for a similar event.
- **Find Next:** Searches after the current selection for a similar event.
- **Find All:** Searches the entire file for events similar to the current selection.

63. Markers & Regions

Overview

Markers and Regions allows you to define and save particular points or selections in time for your audio file. All markers and regions that are created will be saved with your audio file when exporting.

Controls



- **Checkboxes**: The checkboxes to the left of your Marker list designate your selection. The checkbox at the top of the list toggles between selecting all and none.
- **Play Button**: Begins playback from the start point of your marker or region.
- **Find Button**: Moves the playhead to the exact position of your marker, but will not begin playback. For a region, Find will select the audio in the region.
- **Add**: Creates a new marker point at the exact current position of the playhead.
- **Remove Selected**: Deletes the desired marker or region.
- **Select all**: Selects all markers and regions.
- **Select None**: Deselects all markers and regions.
- **Import Marker File**: Opens exported marker information from a file.
- **Export Marker File**: Saves markers and regions in a tab-delimited text file that can be opened and used in another file or session of RX.

More Information

- You can add a marker or region at the current location of the playhead by pressing the M key or selecting Add Marker or Region from the Window menu.
- Markers and regions can also be created from the Markers and Regions window.
- If you have a large number of markers and regions, you may use the search box in the upper right to quickly locate the ones you're looking for.

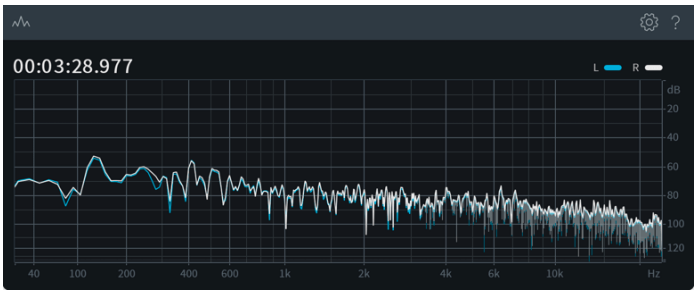
64. Spectrum Analyzer

Overview

A spectrum analyzer uses a fast Fourier transform (FFT) to extract frequency information from a waveform. Depending on the size of the FFT, the signal energy of thousands of frequency bands can be visually represented on a graph.

The RX Spectrum Analyzer will show the momentary spectrum of audio around the current playhead position, the average spectrum of a selected time and frequency range, or the real-time spectrum of the audio at the output of RX's playback.

Controls



Channel View Selection Buttons

When viewing the spectrum analyzer of a stereo file, you can toggle the Left or Right channel spectrum display on or off by clicking on the L or R buttons in the upper right hand corner of the spectrum.

Peak Finding

The RX Spectrum Analyzer has a peak-finding feature that automatically detects peaks in the spectrum data. If you hover your mouse cursor near a peak in the spectrum, a readout will appear displaying the exact frequency of the peak, its amplitude, and the closest musical note. This peak-finding readout can provide much higher accuracy than simply inspecting the graph by zooming in on the display and/or increasing the FFT size in the settings window.

The circle displays the exact amplitude and frequency of the spectral peak. It is usually slightly above the spectrum, because each spectral peak consists of several FFT bins, and their power is added together. This effect is known as spectral smearing (or frequency smearing) and is controlled by the choice of a weighting window.

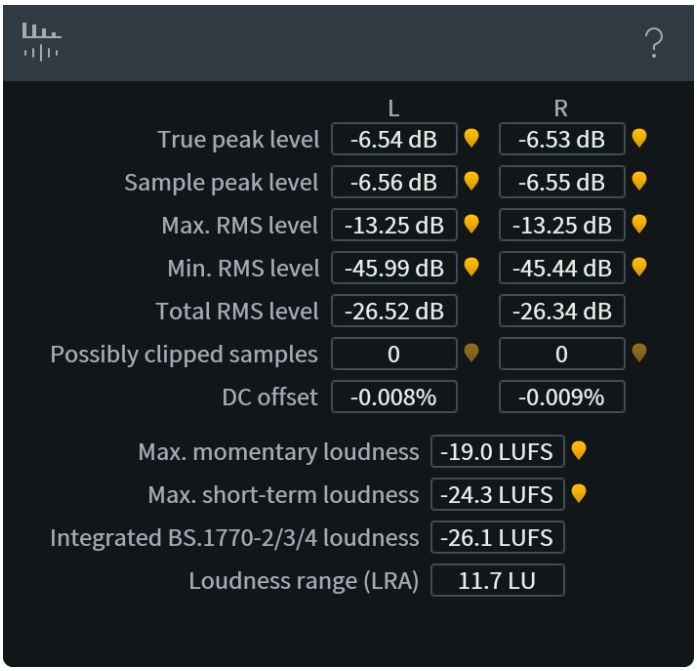
65. Waveform Statistics

Overview

The Waveform Statistics window shows information about the waveform within your current selection. This information can help you find and fix problems like clipping more easily. It is also useful for comparing two similar selections or files. The Waveform Stats window will update to display information about your active track and can remain open as you switch between track tabs.

- If no selection is made in the active file tab, the Waveform Statistics will populate with values based on the entire file.
- Values that appear in the Waveform Statistics window are based on the full bandwidth of the selection. This means that the Waveform Statistics are gathered based on time selections, and will display values for all frequencies in a time selection, even if the current selection only includes a specific frequency range.
- Selecting the cursor icon next to a value readout will move the playhead to the position in the file where the level was detected.

Controls



- **TRUE PEAK LEVEL:** The highest peak level detected, including signal levels between digital samples (called ISPs or Intersample Peaks).
- **SAMPLE PEAK LEVEL:** The maximum level of digital samples in your selection.
- **MAX. RMS LEVEL:** The highest RMS level detected in the selection.
- **MIN. RMS LEVEL:** The lowest RMS level detected in the selection.
- **TOTAL RMS LEVEL:** The RMS level of the entire selection.
- **POSSIBLY CLIPPED SAMPLES:** The number of samples where the true peak signal level exceeds 0 dBTP.

- **DC OFFSET:** The amount of DC offset, in percent of the full scale. Hover this value with your cursor to see this level in decibels.
- **MAX. MOMENTARY LOUDNESS:** Momentary loudness is computed on a K-weighted audio signal using 400-ms windows, as defined by the BS.1770 specification.
- **MAX. SHORT-TERM LOUDNESS:** Short-term loudness is computed using 3000-ms windows, as defined by the BS.1770 specification.
- **INTEGRATED LOUDNESS:** Displays the integrated loudness level (as defined by BS.1770-2). Settings in the Loudness Control module may switch this field to BS.1770-1 standard.
- **LOUDNESS RANGE (LRA):** Displays the loudness range (as defined by BS.1770). This value reflects the dynamics of audio levels in the selection.

More Information

True peak

True peak reflects the expected peak level of the analog waveform after digital-to-analog conversion. In practice, it is calculated by oversampling of the digital waveform according to BS.1770-3 standard. Different software may measure true peak levels slightly differently, because the standard leaves some freedom in choice of oversampling filters.

RMS

RX measures RMS using 50-ms windows (more specifically, a Hann window with a 100-ms period) and references levels using either an AES-17 standard (full-scale sine wave = 0 dB RMS) or a "scientific"standard (full-scale square wave = 0 dB RMS). These RMS measurement standards differ by 3 dB and can be chosen in RX Preferences.

66. RX Plug-ins

Plug-in Formats

The RX 11 installer will install only the following plug-in formats:

- **AU**
- **AAX** (Realtime & Audiosuite)
- **VST3**
- **AU ARA (Spectral Editor ARA & Music Rebalance ARA - Compatible with Logic Pro only)**

All plug-in formats are 64-bit only. VST2 is no longer supported.

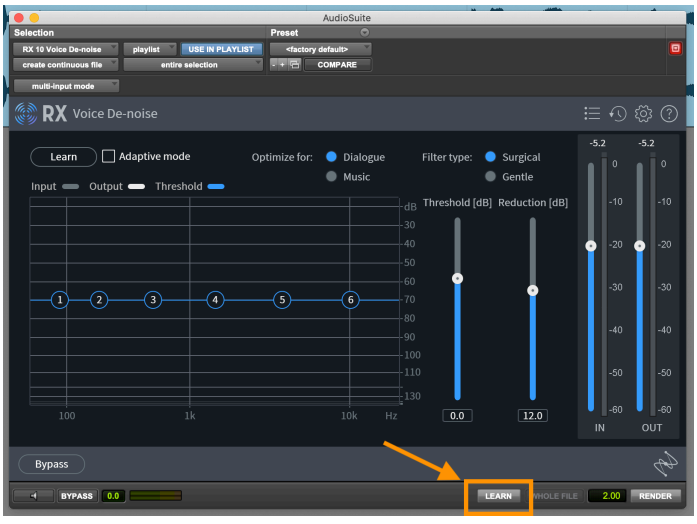
RX 11 Ambience Match, Dialogue Isolate, and De-rustle plug-ins are available as AAX Audiosuite plug-ins in Pro Tools only.

Learning in Audiosuite

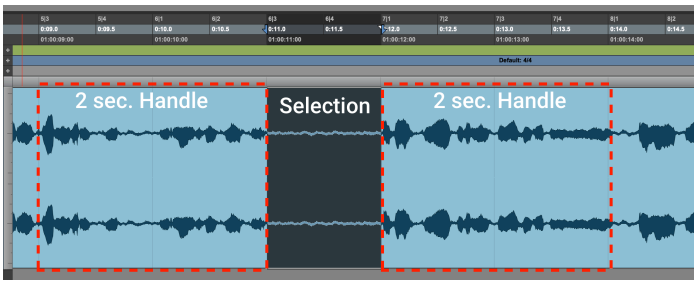
When using RX AudioSuite plugins in Pro Tools, there are two main workflows for learning noise profiles. Each has its own benefits and drawbacks. These two workflows are outlined as follows:

Method 1 - AS Learn (Audiosuite Learn)

Make a selection and click the Learn button that appears in the bottom right of the AudioSuite plugin wrapper. With this method, learning happens offline, saving you time.



It is important to note that learning with this method also includes “handles” in the learn pass. Handles are a Pro Tools feature which extend rendered selections by a set amount. This creates a safety net of additional processed material underneath the clip edges to allow for flexibility during editing. By default, Pro Tools handles are set to 2.00 seconds. With this handle setting, for a given selection, the learn pass would *also include audio 2 seconds before and 2 seconds after the selection (pictured below).*



This may account for unexpectedly intense learned noise profiles, because the plugin is receiving noise + dialogue or other material on either side of the noise selection. To avoid this, you can set handles to 0 via the text box next to the Render button, or learn in the middle of a long section of noise.

If you regularly use handles in your workflow and want to avoid manually turning handles on and off, then Method 2 below may provide better results for you.

Method 2 - UI Learn

1. Make a selection and activate the Learn button that is inside the RX plugin UI, usually on the upper left side.
2. While the learn state is active & listening (button turns blue), click Preview (button with speaker icon) in the lower left corner of the AudioSuite plugin wrapper. This passes the selected audio into the learn pass in realtime.
3. Stop transport when you are finished learning, and the learn pass will turn itself off automatically.



A benefit to this method is that only audio you've selected will playback into the plugin. There is no hidden audio provided to this learning pass via Pro Tools' handles feature. The drawback is that you must listen to your learn pass in real time. This "UI Learn" method may provide better results if you typically work with handles on (which is common in post-production).

Presets

From the Preset Manager, you can select from default presets and presets you have saved.

To browse presets, press the Presets button and click the name of any preset. If you like what you hear, press the Presets button again to hide the window.

- **ADD:** Clicking this button adds the current settings as a new preset. You can type a name and optionally add comments for the preset.

  > ⓘ Note that characters that are not allowed in Windows file names such as * or / cannot be used as preset names and will be ignored when entered. This is to ensure consistency with the names of the underlying .xml files for easy backup and transferring.

- **REMOVE:** To permanently delete a preset, select the preset from the list and click the Remove button.
- **UPDATE:** When you click the Update button, your current settings before you opened the preset window are assigned to the selected preset (highlighted). This is useful for selecting a preset, tweaking it, and saving your changes to the existing preset.
- **IMPORT:** Imports a preset into the preset folder.
- **FOLDER:** Opens a dialog that shows your current preset folder. You can also select a new preset folder from this dialog.
- **RENAMING PRESETS:** You can double click on the name of a preset to enter the edit mode and then type a new name for that preset.
- **CANCEL:** Press Escape to close the preset system dialog and revert to the settings when you opened the preset manager.

History

Pressing the History button opens the History window. This view gives you a list of all of the actions you've performed inside the plug-in, allowing you to step back to previous settings and undo changes.

- **CLEAR:** Resets the History list.

I/O Meters

- Select plug-ins feature an input and output gain control and input and output metering.
- In stereo instances of the plug-ins, the gain controls are linked stereo gains.

Options

General Options

Authorization & Updates

The Authoriztion and Updates section of the General Options tab includes the following options:

- **Authorization:** Learn more in the Authorization chapter.

Host Performance

- **Enable Multicore (Spectral De-noise only)** This option is available in the Spectral De-noise plug-in. Enabling this option lets RX process Spectral De-noise Quality Modes C and D more efficiently by spreading its processing across multiple computer cores.

I/O Options

The RX I/O meter displays a lower bar representing the average level (RMS) and a higher bar representing peak level. There is also a moving line above the bar representing the most recent peak level or peak hold.

- **PEAK HOLD TIME:** If peak hold is on, you can choose different peak hold times. The choices are 250 ms, 500 ms, 1000 ms, 5000 ms and Infinite. If set to infinite, the peak value will be held until you click on the current peak display.
- **INTEGRATION TIME:** Specifies the integration time for RMS calculation. In most RMS meters, the integration time is set to around 300 ms, which makes the RMS meter similar in ballistics to VU meters.
- **READOUT:** Allows you to control what level value is displayed on top of the meters: peak or actual (real time). If set to Max Peak, the display will show the highest peak level. If set to Current, the display will reflect the meter's current value.
- **ENABLE I/O METERS:** Turns the level meters on or off.
- **SHOW PEAK HOLD:** Turns the peak hold display for the level meters on or off.

Latency

Some of the processing modes in the RX plug-ins are very CPU intensive and result in a delay of the signal. That is, RX needs some time to process the audio before it can send it back to the host application. That time represents a delay when listening or mixing down.

Most modern DAWs and NLEs provide delay compensation—a means for RX plug-ins to tell the application it has delayed the signal, and the host application should "undo" the delay on the track (usually by adding compensating delays to other tracks in real-time processing, or by adjusting the rendered file after processing in offline processing). When Enable Delay Compensation is enabled in the Latency menu, we will report our latency to the host application.

If your application doesn't support it, or skips/stutters with this option on, you can always manually correct the delay offset in the host application (manually edit out the short delay of silence). To help you perform manual correction, the delay RX introduces is shown below as Total System Delay in both samples and milliseconds.

67. RX Monitor

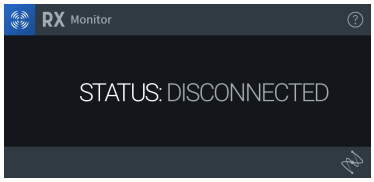
Plug-in only

Overview

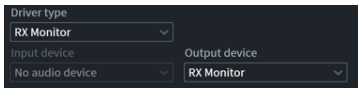
When using the RX Connect plug-in, some DAW/NLEs monopolize the system's audio drivers, preventing RX from playing audio through the same output device. The RX Monitor plug-in allows you to listen to the output of the RX Audio Editor through the audio driver output of your host application. This is particularly useful when using RX Connect with a host application that would need to be closed for the RX Audio Editor to access the output driver it is using.

Workflow

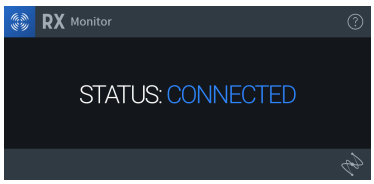
1. Insert the RX Monitor plug-in on an Aux or Instrument track in your DAW/NLE. When you first insert the RX Monitor plug-in, you will see a **Status: Disconnected** message. You need to configure RX Monitor as the Audio Driver in the RX Audio Editor to change the status to **Connected**.



2. Change the Audio Driver in the RX Audio Editor:
   a. Open **Preferences**.
   b. Select the **Audio** tab.
   c. In the **Driver Type** dropdown menu, select **RX Monitor**.



3. Navigate back to your DAW and confirm the RX Monitor plug-in is showing a "Connected" status.



4. Playback audio in the RX Audio Editor, the output will be routed to the track you added RX Monitor to in your DAW/NLE.

> **NOTE**
>
> If you are running virus protection software or a firewall on your DAW, you may need to grant permission for RX Monitor to run.

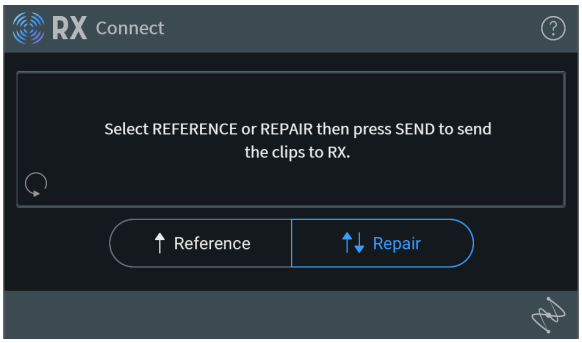
68. RX Connect

Plug-in Only

Overview

The RX Connect plug-in sends a clip, or multiple clips, to the RX 11 standalone application for editing and repair. This gives you access to all of RX 11's modules in one place, and provides the benefits of RX's offline processing and visual interface. RX Connect is available from the AudioSuite menu in Pro Tools, or as an AU or VST3 plug-in from your DAW's effects menu.

Controls



There are two modes for using RX Connect:

- **Send For Reference:** This is meant for analysis only. The clips are imported into RX 11 but cannot be sent back to your host.

  NOTE ABOUT REFERENCE MODE

  This mode of RX Connect will not open RX 11 Audio Editor automatically. Opening the Audio Editor after using send for reference will reveal the file in the RX Audio Editor.

- **Send For Repair:** Selected clips are sent to the RX 11 Audio Editor for repair, and you can send them back to your host from the RX Audio Editor.

  MORE INFORMATION

  For more information on using RX Connect in different DAWs or NLEs, please refer to the following sections, or check out the RX Connect FAQ on our Support website.

DAW/NLE specific instructions

The following sections outline host specific instructions for using RX Connect to Send audio to the RX Audio Editor and back to your host application:

Adobe Audition CC RX Connect Workflow

1. Inside of Audition, select the Waveform view.
2. Highlight the area of audio that requires editing.
3. In the Effects menu, load the **RX 11 Connect** plug-in from **VST3** > **Restoration** > **iZotope** (If you do not see the **RX 11 Connect** plug-in, open the Audio Plug-in Manager and Scan for Plug-ins, then make sure **RX 11 Connect** is enabled).
4. When the plug-in window opens, click **Apply**.
5. RX 11 will automatically load. Perform your desired audio edit, then click **SEND BACK** to send the audio back to Adobe Audition. The Waiting for Connect message will appear.
6. Re-load the **RX 11 Connect plug-in** from the **Effects** menu. It will now display a message **Press Apply to commit changes**.
7. Click **Apply** to apply the audio edit from RX to your audio file in Adobe Audition.

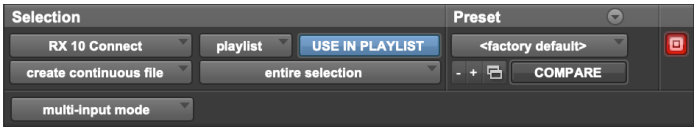
Avid Pro Tools RX Connect Workflow

1. Choose the audio to be sent to the RX Audio Editor by selecting the audio clip(s) in the timeline that you want to edit, and opening RX Connect from the AudioSuite **Noise Reduction** menu.
2. If you just need to load a noise profile or analyze some audio, choose **Reference** to send the audio one-way, but for the complete round-trip workflow click **Repair** and then hit **Send**. You'll see this opens the audio in the RX Audio Editor.
3. With HDX systems, Pro Tools will have control of your audio drivers, so you aren't able to hear the output of the RX Audio Editor. However, the RX Monitor tool is built to solve just this problem. In Pro Tools, create a dedicated aux track for monitoring RX, and insert RX 11 Monitor from the **Noise Reduction** or **Sound Field** menus,
4. Then, go to the 'Preferences' menu in the RX Audio Editor by clicking on the wrench icon in the top-right of the window. In the **Audio** tab, set your **Driver type** to be **RX Monitor**. Now we can hear the output of the RX Audio Editor through your Pro Tools output chain.
5. After you've made the desired edits in RX, click **Send back** at the top of the window. Click **Render** in the RX Connect window, and the repaired audio will be placed back into your session.



Some engineers might choose to create duplicate playlists before making any repairs to their audio, but you can 'undo' these RX Connect changes just like any AudioSuite process. If you make extensive repairs inside of the RX Audio Editor, you can also save an .rxdoc of the file, which will preserve all your adjustments so you can modify them later if you need to.

Audiosuite modes



When using Audiosuite plug-ins, there are various user definable input and output options, which affect how you may use RX Connect. These options are:

Input

- **Clip-by-clip:** Recognizes individual clips in the timeline, as well as fades.
- **Entire selection:** Treats the entire selected area as one clip.

Input channel modes

- **Mono mode:** Treats mono, dual mono and stereo clips, as well as multi-channel clips, all as discrete mono clips (e.g. a stereo clip will send as two separate mono files).

> **MONO MODE NOTE**
>
> Please note, this can result in large groups of audio clips being sent to RX, potentially exceeding the maximum file limit of 32.

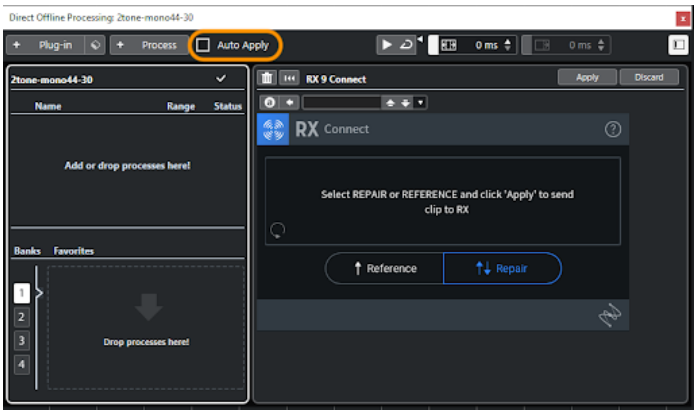
- **Multi-input mode** Treats dual mono and stereo audio clips as one entity.

Output

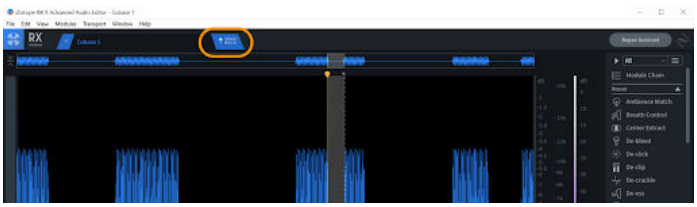
- **Overwrite files:** Destructive processing of the audio clip(s) in the session, overwriting the original file with the new file sent from RX.
- **Create individual files:** Nondestructive processing of the audio file(s) in the session, replacing them with the audio processed in RX. This mode preserves individual clips and fades/handles.
- **Create continuous files:** Nondestructive processing of the original audio file. Creates a new audio file with the audio sent back from RX, consolidated into one continuous clip.

Steinberg Cubase and Nuendo RX Connect Workflow

1. Select the audio clip you wish to apply changes to.
2. Navigate to the **Audio** menu in Cubase/Nuendo and select **Direct Offline Processing**.
3. Uncheck **Auto-Apply** in the Direct Offline Processing Window.



1. Click the **+ Plug-in** button and select RX 11 Connect.
2. Press **Apply** to send the file to RX.
3. The RX Audio Editor application will automatically open with the file you sent loaded in a tab named “Cubase 1” or “Nuendo 1”.
4. Make the desired changes to your file in the RX Audio Editor.
5. Click the **Send back** button in the RX file tab display to send the updated file back to Cubase/Nuendo.



1. Click the **Apply** button in the **Direct Offline Processing** window to apply the changes to the file in the session.

> **APPLY BUTTON MISSING?**
>
> If **Auto Apply** is enabled in the **Direct Offline Processing window**, the **Apply** button will not be available. Disable **Auto Apply** and restart the process with a new instance of RX Connect.

2. Navigate to the **Audio** menu and select **Make Direct Offline Processing Permanent.**

Using RX as an External Audio Editor

Some hosts don't support the use of RX Connect for round-trip editing, please refer to the instructions below for host specific workflows.

Adobe Premiere Pro CC with RX as an external audio editor

1. Inside of Premiere, right-click on an audio clip in your timeline and select **Reveal in Finder** (Mac) or **Reveal in Explorer** (Windows).
2. Open the resulting file in the **RX 11 Audio Editor**.
3. Perform desired processing in the RX application.
4. When you have made the desired changes to your file, go to the RX **File** menu and select **Overwrite Original File**.

> **NOTE**
>
> If you have Adobe Audition installed as well, you can right-click on an audio clip in your timeline, and select **Edit Clip In Adobe Audition**. Then follow these steps for using RX Connect with Adobe Audition.

Apple Logic Pro X with RX as an external audio editor

RX is a powerful audio editor that Apple Logic Pro X users can use to get better sounding audio. To use RX with Logic, you must first set it up as an external audio editor.

How to set up RX as an external audio editor

1. Open Logic **Preferences** > **Advanced**.
2. Under Additional options, enable the **Audio** check box.
3. In **Preferences**, click the **Audio** tab and select the **Audio File Editor** tab.
4. Under **Audio File Editor**, click on the **External Sample Editor** to select **iZotope RX 11 Audio Editor** from your applications folder.

Workflow

1. Select the clip you wish to edit in your timeline.
2. Select **Edit** > **Open in iZotope RX 11 Audio Editor**, `Shift` + `W`.
3. The file will open in RX 11. Once you've completed your edits, open the RX 11 File menu and select **Overwrite Original File**.
4. Close the tab, navigate back to Logic Pro X and wait for the waveform to update.

DaVinci Resolve with RX as an External Editor

To use RX as an external audio editor with Resolve, you must first set it up as an external process in the Resolve Preferences menu.

Configure RX as an external process in DaVinci Resolve

1. Open the **DaVinci Resolve** menu and select **Preferences**
2. Go to the **Audio Plugins** tab
3. Click **Add** in the **Setup External Processes** section
4. Configure the New External Process:
   a. **Name**: Double-click on the default name of the new external process to edit it. Edit the text and hit return to save the changes. We recommend changing this name to 'RX' so it is easily identifiable in context menus
   b. **Path**: Double-click on the **Path** field to set the path to the RX Audio Editor application. In the system dialog that appears, navigate to your Applications/Program Files directory, select the **iZotope RX Audio Editor** application, and click "Open"
   c. **Type**: Use the default setting of "Command Line"
5. Click **Save** in the bottom of the Preferences window and click **OK** in the 'Preferences Updated' dialog that appears

Working with RX as an external editor with DaVinci Resolve

Once you have configured the RX application as an external process in Resolve, use the following steps to send clips to RX for editing:

1. Open the **Fairlight** tab (musical notes icon at the bottom of Resolve)
2. Right-click on the audio clip you want to edit in RX
3. Go to **External Audio Processes** in the context menu, select the **RX** process you configured in Resolve Preferences
4. The clip will open in the RX application, make your edits, and save your changes in RX (using **File** menu > **Save**, or `Command`+`S` /`ctrl`+`S`) to update the clip in your Resolve project

> **NOTE**
>
> When you send a clip from Resolve to RX, saving changes in RX will immediately update the clip in your Resolve project. You may want to create a copy of the clip in your Resolve project before opening and editing it in RX if you want to maintain access to the original.

69. RX Spectral Editor

Overview

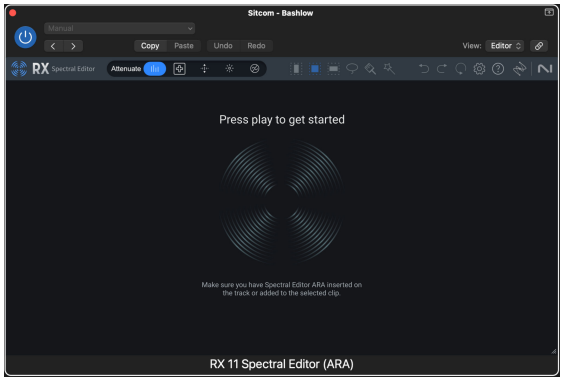
The Spectral Editor plug-in combines ARA 2.0 technology with the award-winning DSP and selection-based editing workflows of the RX Audio Editor. It allows you to identify and reduce undesirable sounds in your tracks without leaving your DAW.

Getting Started with Spectral Editor (ARA)

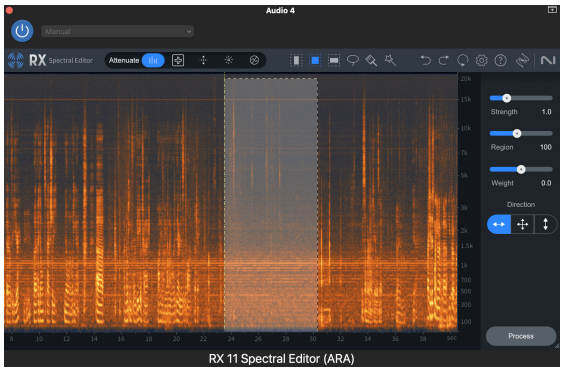
Logic Pro

Spectral Editor is currently only available in Logic Pro as an ARA 2.0 plug-in (Rosetta only). You can add RX Spectral Editor (ARA) to your project using the following steps:

1. Make sure you are using the latest version of Logic Pro. The RX Spectral Editor ARA plug-in is supported in **Logic Pro 10.7+ or higher.** We recommend that you update Logic Pro to the latest available version to ensure the best experience using RX Spectral Editor (ARA).
2. Add **RX Spectral Editor (ARA)** as the **first insert** on an **Audio** track.



3. Select a clip in Logic and start playback. The selected audio will appear in the spectrogram view of the Spectral Editor (ARA) plug-in after starting playback.
4. Make a selection in the spectrogram view of RX Spectral Editor.



5. Process your selection using one of the following modes described below.
6. To work on another clip in your session, select it and start playback. Once playback has started, you will see your new clip load in the spectrogram. Trimmed clips appear the same way in the spectrogram as they do in Logic.

Pro Tools

You can add RX Spectral Editor (ARA) to your Pro Tools project using the following steps:

1. Make sure you are using the latest version of Pro Tools. The RX Spectral Editor ARA plug-in is supported in Pro Tools 2024.6 or higher.
2. Clip menu:
    a. Select a clip on the timeline
    b. Open the Clip menu at the top of Pro Tools
    c. Select RX Spectral Editor > Edit to open that clip in the Spectral Editor.
3. You can add Spectral Editor to an entire track in one of two ways:
    a. The Elastic Audio/ARA plug-in selector drop down on the track itself
    b. Track menu
        i. Select a track
        ii. Go to Track Menu at the top of Pro Tools
        iii. Select RX Spectral Editor > Edit.
        iv. Once added to a track, select a clip on the track to open it in Spectral Editor.

All clips that exist on a track at the time Spectral Editor is added to it will be visible to the editor. If you add additional clips to the track or process pre-existing clips with AudioSuite, you will need to manually enable Spectral Editor on those new clips with the Clip menu method above.

Closing & Reopening the Editor

You can close and reopen the Spectral Editor window by clicking on the insert that appears in the Elastic / ARA slot on a track. You can also access it via the RX Spectral Editor tab at the bottom of the Pro Tools Edit window.

Committing edits to a new file on disk

When you're finished making edits with Spectral Editor, if you prefer, you can commit your edits directly to a new audio file on disk. This is similar to an AudioSuite workflow. To accomplish this, right click on a clip you want to commit and select RX Spectral Editor > Render. From this same menu, you can also clear or bypass Spectral Editor ARA processing.

If you edit a clip with Spectral Editor ARA and then process that clip with an AudioSuite plugin, your edits will be committed to the new file.

Studio One

The VST3 version of Spectral Editor ARA works directly on clips as an "Event FX" in Studio One. Event FX are applied to individual audio events rather than as inserts on entire tracks.

> **MACOS NOTE**
>
> Although an AU version of Spectral Editor ARA will show up as available in Studio One, it is not supported. The UI will direct you to use the VST3 version.

You can add RX Spectral Editor (ARA) to an audio event in your Studio One session using one of the following methods:

1. Make sure you are using the latest version of Studio One. The RX Spectral Editor ARA plug-in is supported in **Studio One 6.0 or higher**. Studio One 6.5 introduced additional improvements that make it easier to access ARA plugins.

2. Event FX Context menu [**Studio One 6.5 or higher**]:
   a. Right click on an audio event
   b. In the context menu, select Event FX > RX 11 Spectral Editor
3. Effects Browser:
   a. **[Studio One 6.5 or higher]**, Spectral Editor will appear in the "Event Editors" category at the top of the Effects browser.
      ▶ Click & drag RX 11 Spectral Editor onto an audio event
   b. In **Studio One 6 versions before 6.5**, Spectral Editor will appear under the regular VST3 location.
      i. While dragging Spectral Editor onto your event, if the status text doesn't say "Add Event FX", then hold OPT (Mac) or ALT (Win) and that message should appear.
      ii. The plugin will display a warning message if it has been incorrectly added as a normal channel effect.
4. Inspector:
   a. Open the Inspector (the 'i' icon in the top left)
   b. Select an audio event
   c. Underneath the Mixer section of the inspector, there is a section with info about the currently selected event.
   d. Click the Event FX Enable button, which will open a toggle dropdown
   e. In the Inserts section that appears, use the + sign to add Spectral Editor directly to the clip as an Event FX.

Reopening the Editor

To re-open the Spectral Editor window on a clip after you've closed the UI, either

- Select the clip and double click the Spectral Editor Event FX insert at the bottom of the Inspector
- **[Studio One 6.5 or higher]** Right click on the clip and select Event FX > RX 11 Spectral Editor from the context menu again.

Committing edits to a new file on disk

When you're finished making edits on a clip with Spectral Editor, if you prefer, you can commit your edits directly to a new audio file on disk. You do this with the Render command. With the Inspector open, ('i' icon in the top left), clicking on a clip will display info about that clip at the bottom. Click the orange Render button under the Event FX section at the bottom. You can also click Restore if you wish to bring the Editor back after Render.

A note about duplicated/copied audio

In Studio One, all copies of the same source audio will inherit the same edits when Spectral Editor is applied to each copy. For example, if there are two copies of Clip A in your session and you add the Spectral Editor plugin to each, then any processing applied to copy 1 will be synchronized to copy 2, and vice versa.

If you wish to decouple edits for multiple copies of the same source audio, right click on your clip and select Audio > Audio Parts > New Clip Version. You will see a small number appear in the bottom left of your clip. Clips with the same version number will share edits, as long as Spectral Editor ARA is applied on all the clips. In the session Pool, all clip versions will be grouped under the parent clip and can be swapped in/out from there.

Consult Studio One's help documentation for more information about Clip Versions and working with Event FX.

PLAYBACK IS REQUIRED TO REFLECT CHANGES IN SPECTROGRAM

ARA plug-ins are updated when playback is started. If you trim or otherwise change a clip that was already loaded in Spectral Editor, the spectrogram will not be updated until you've stopped and restarted playback with the clip selected.

Troubleshooting

If you find that RX Spectral Editor (ARA) is missing from the plug-in selection menu, try the following:

1. Ensure you are adding the RX Spectral Editor (ARA) plug-in to the **first** insert slot of an **Audio** track.
2. **Clear the AudioUnit Cache and rescan plug-ins**:
   - **Logic 10.7 or higher**: Go to the Logic Pro menu > Preferences… > Plug-in Manager… and select "Full Audio Unit Reset" at the bottom of the window to reset the cache and rescan plug-ins.
3. See the ARA Requirements and Limitations section below for more information about working with Spectral Editor (ARA) in Logic Pro.

Controls

Once your clip is loaded in the spectrogram view of the Spectral Editor, you'll have access to the following controls:

Processing Mode

Select Attenuate, Replace, Gain, De-click or De-hum and the relevant processing parameters will appear on the right side of the plugin.



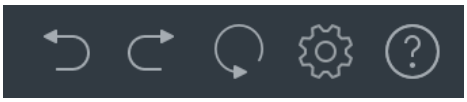
Selection Tools

Determines the selection tool type used when making audio selections in the spectrogram.



See the Interactive Tools chapter for more information about selection tools and modifiers.

General Controls



- **Undo:** Reverts to the previous processing step.

> **EDITING COPIES OF CLIPS**
>
> The Undo function in Spectral Editor will apply to all non-unique copies of a clip. If you wish to edit different copies of the same clip independently using Spectral editor, please bounce them first.

> **UNDO HISTORY WARNING**
>
> Undo history is cleared when a session is saved.

- **Redo:** Re-applies the next processing step in your history.
- **Reset:** Reverts all changes in the audio to the unprocessed state.
- **Settings:** Opens the plug-in General Settings menu.
- **Help**: Opens the installed HTML help documentation in your default web browser.

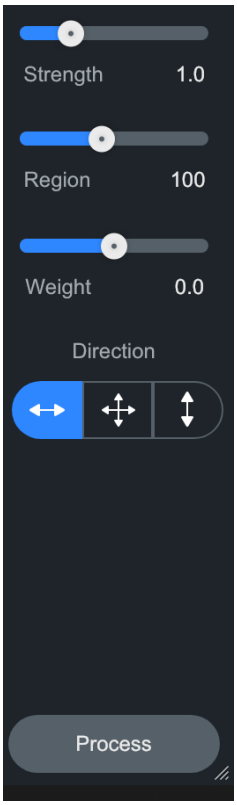
> **USE THE TRACKPAD TO NAVIGATE THE SPECTROGRAM**
>
> Scroll with the trackpad vertically to zoom in and out, and horizontally to scroll forward/backward in time.

Processing Modes

Attenuate



Attenuate removes sounds by comparing the content inside of your selection to the content outside of your selection. It modifies dissimilar audio in your selection to be more similar to the surrounding audio. Attenuate does not re-synthesize any audio.

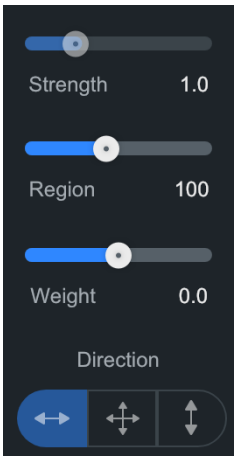


- **Strength:** Adjusts the amount of attenuation.
- **Region:** Defines how much of the surrounding content will be used for interpolation.
- **Weight:** Gives more weight to the surrounding audio before or after the selection.
- **Direction:** Determines where the material used in the repair process is located in relation to the current selection. **Note**: Replace mode always uses Horizontal interpolation when processing a selection.
- **Horizontal**: Signal to the left and right of the current selection will be used for interpolation.
- **2D**: Signal above, below, to the left and to the right of the current selection will be used for interpolation.
- **Vertical**: Signal above and below the current selection will be used for interpolation.
- **Process**: Applies processing to the current selection.
- **Resize**: Allows you to click and drag to change the size of the plug-in window.

Replace



Replace resynthesizes badly damaged sections (such as gaps) in tonal audio. It completely replaces the selected content with audio interpolated from the surrounding data.



- **Strength:** This control is not active in Replace.
- **Region:** Defines how much of the surrounding content will be used for interpolation.
- **Weight:** Gives more weight to the surrounding audio before or after the selection.
- **Direction:** This control is not active in Replace. The Replace mode always uses Horizontal interpolation when processing a selection.

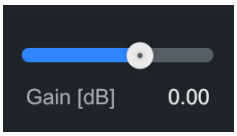
PROCESSING SELECTIONS LONGER THAN 10 SECONDS PROCESSING SELECTIONS LONGER THAN 10 SECONDS

Attenuate: Horizontal, Attenuate: 2D, and Replace mode do not support processing selections longer than 10 seconds. If your selection is longer than 10 seconds, the processing will automatically default to Attenuate: Vertical.

Gain



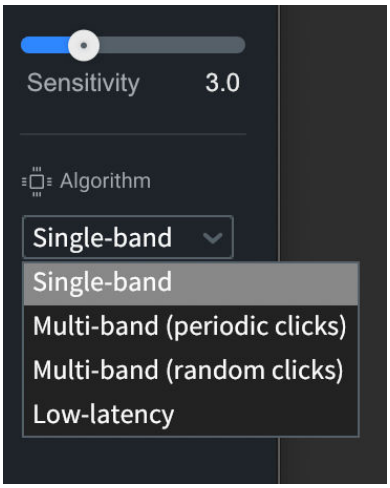
Gain boosts or cuts the level of the signal by the designated decibel amount.



De-click



De-click analyzes audio for amplitude irregularities and smoothes them out. Use De-click to remove a variety of short impulse noises, such as clicks caused by digital errors, mouth noises, and interference from cell phones.

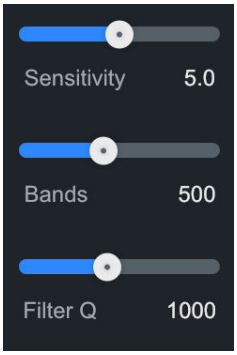


- **SENSITIVITY**: Determines how many clicks are detected in the signal.
- **ALGORITHM**: Provides options for adjusting the configuration and processing quality of click interpolation, depending on the types of clicks and pops present in the audio.
  - SINGLE BAND: Processes quickly and works well on very narrow “digital” clicks
  - MULTI-BAND (PERIODIC CLICKS): Uses multi-band processing for removing regularly repeating clicks with a wider spectrum, or regular clicks that have concentrated low or high energy (like thumps or optical soundtrack perforation noise)
  - MULTI-BAND (RANDOM CLICKS): Uses multi-band processing for wider vinyl clicks and thumps, with a protective algorithm for preserving periodic audio elements characteristic to certain instruments such as brass or vocals
  - LOW LATENCY: Works well on mouth clicks and other clicks that cannot be handled by other algorithms. Low latency mode is suitable for real-time work.

De-hum



De-hum is designed to remove persistent tonal noise, like AC noise, electrical buzz or interference introduced by EMF (Electro-magnetic fields). De-hum uses the Dynamic mode with Adaptive learning enabled.

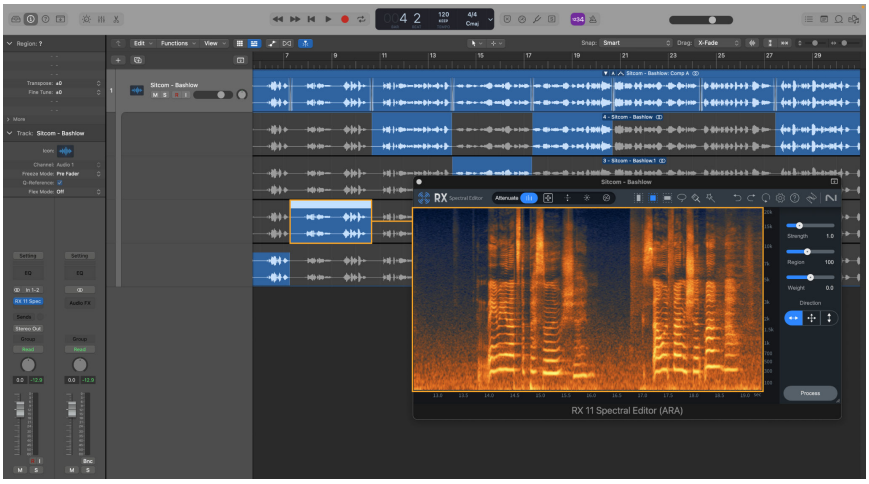


- **SENSITIVITY**: Adjusts the amount of hum that will be removed.
- **BANDS**: Controls the number of dynamic notch filters, which update their gains based on the input signal to avoid ringing. Adjust to higher values for complex hums, at the expense of transparency.

- **FILTER Q**: Controls the bandwidth of the notch filters.

Working with Comp Takes in Logic Pro

Spectral Editor supports comp takes in Logic. 7.5+ (Rosetta only). To process comp takes, the individual takes that make up the comp must be selected and processed.



The Logic comp region at the top will not appear in the spectrogram if selected. Choose the individual takes and toggle playback to work on comps in Spectral Editor.

Keyboard Shortcuts

The following keyboard shortcuts are available in Spectral Editor. CMD is for Mac, CTRL is for Windows:

- Undo in Spectral Editor: CMD/CTRL+Z
- Redo in Spectral Editor: Shift+CMD/CTRL+Z
- Switch to Attenuate mode: CMD/CTRL+1
- Switch to Replace mode: CMD/CTRL+2
- Switch to Gain mode: CMD/CTRL+3
- Switch to De-click mode: CMD/CTRL+4
- Switch to De-hum mode: CMD/CTRL+5
- Process: CMD/CTRL+Enter
- Choose Time selection tool: 1
- Choose Time-frequency selection tool: 2
- Choose Frequency selection tool: 3
- Choose Lasso selection tool: 4
- Choose Brush selection tool: 5
- Choose Magic wand selection tool: 6

ARA Requirements and Limitations

- Spectral Editor (ARA) must be the first plug-in on a track.

- Spectral Editor (ARA) cannot be added to a software instrument track or bus.
- Spectral Editor (ARA) is not supported by Selection Based Processing.
- Host presets are not supported by Spectral Editor (ARA).
- Automation is not supported by Spectral Editor (ARA).
- Playback must be started to gather audio data. After editing, playback must be stopped and started again to update data.
- Spectral Editor (ARA) cannot process Apple Loops, compressed audio, flex audio, or reversed regions. These regions can be bounced in place in order to be processed by Spectral Editor (ARA).
- Spectral Editor (ARA) will output silence for stereo files on mono tracks.
- Spectral Editor (ARA) will output silence for regions whose sample rate does not match the sample rate of the session.

70. Keyboard Shortcuts

The RX Audio Editor includes options for defining custom keyboard shortcut commands.

The column named "RX Shortcut Command Name" in the table below lists the underlying name associated with each default shortcut in the RX Audio Editor. These are useful to know when assigning custom shortcuts in the Keyboard tab.

IMPORTING KEYBOARD SHORTCUTS FROM RX 10

If you would like to carry over custom shortcuts you set in RX 10, you will need to export the keybindings from RX 8 and then import that keybindings file to RX 11 using the Import option in the Keyboard tab.

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
New...	command+N	ctrl+N	File.New
New from Clipboard	shift+command+N	ctrl+shift+N	File.NewFromClipboard
Open...	command+O	ctrl+O	File.Open
Save	command+S	ctrl+S	File.Save
Save As...	shift+command+S	ctrl+shift+S	File.SaveAs
Save RX Document			File.SaveRXDocument
Save RX Document As...			File.SaveRXDocumentAs
Overwrite Original File	option+command+S	ctrl+alt+S	File.SaveOverwriteOriginal
Export...	command+E	ctrl+E	File.Export
Export Selection...	shift+command+E	ctrl+shift+E	File.ExportSelection
Export Regions to Files...	option+command+E	ctrl+alt+E	File.ExportRegions
Close	command+W	ctrl+W	File.Close
Close All	shift+V	ctrl+shift+W	File.CloseAll
Zoom out full all rulers	command+0	ctrl+0	Zoom.AllOutFull
Zoom in on amplitude ruler	shift+up arrow	shift+up arrow	Zoom.AmpIn
Zoom out on amplitude ruler	shift+down arrow	shift+down arrow	Zoom.AmpOut

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Zoom in on frequency ruler	shift+command+up arrow	shift+ctrl+up arrow	Zoom.FreqIn
Zoom out on frequency ruler	shift+command+down arrow	shift+ctrl+down arrow	Zoom.FreqOut
Zoom in on time ruler	up arrow	up arrow	Zoom.TimeIn
Zoom in on time ruler	command+=	ctrl+=	Zoom.TimeIn
Zoom on left side of time ruler	command+[	ctrl+[	Zoom.TimeLeftEdge
Zoom out on time ruler	command+-	ctrl+-	Zoom.TimeOut
Zoom out on time ruler	down arrow	down arrow	Zoom.TimeOut
Zoom out full on time ruler	shift+command+-	ctrl+shift+-	Zoom.TimeOutFull
zoom on right side of time ruler	command+]	ctrl+]	Zoom.TimeRightEdge
Zoom to time selection	command+\	ctrl+[	Zoom.TimeSelection
Undo	command+Z or option+command+Z	ctrl+Z or ctrl+alt+Z	Edit.Undo
Redo	command+Y or shift+command+Z	ctrl+Y or ctrl+shift+Z	Edit.Redo
Cut	command+X	ctrl+X	Edit.Copy
Copy	command+C	ctrl+C	Edit.Cut
Paste	command+V	ctrl+V	Edit.Paste
Paste Special> Insert	option+command+V	ctrl+alt+V	Edit.PasteInsert
Paste Special> Replace	option+shift+command+V	ctrl+alt+shift+V	Edit.PasteReplace
Paste Special> Mix	shift+V	shift+V	Edit.PasteMix
Paste Special> Invert and Mix	option+V	alt+V	Edit.PasteMixInvert

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Paste Special> To Selection Only	option+shift+V	alt+shift+V	Edit.PasteToSelection
Paste Special> Clip Gain Only	shift+command+V	ctrl+shift+V	Edit.PasteClipGainOnly
Deselect	command+D	ctrl+D	Edit.Deselect
Reselect	shift+command+D	ctrl+shift+D	Edit.Reselect
Select All	command+A	ctrl+A	Edit.SelectAll
Invert Selection	shift+command+I	ctrl+shift+I	Edit.SelectInverse
Invert Selection Frequencies	command+I	ctrl+I	Edit.SelectInverseFreq
Select Harmonics…	shift+command+H	ctrl+shift+H	Edit.SelectHarmonicsByNumbers
Begin Selection At Playhead	[	[	Edit.SetSelectionStart
End Selection At Playhead	]	]	Edit.SetSelectionEnd
Delete Selection	Del	Del	Edit.SilenceDelete
Trim to Selection	command+T	ctrl+T	Edit.TrimToSelection
Snap	shift+command+;	ctrl+shift+;	View.ToggleSnapping
Find Similar Event Window	command+F	ctrl+F	Edit.FindSimilarEvent
Find Next Similar Event	shift+command+F	ctrl+shift+F	Edit.FindNextSimilarEvent
Find Previous Similar Event	option+command+F	ctrl+alt+F	Edit.FindPrevSimilarEvent
Add Marker or Region	M	M	Edit.AddMarkerOrRegion
Edit Cursor Mode > Select Time	T	T	Edit.EditorCursorMode.SelectTime

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Edit Cursor Mode > Select Time/ Freq	R	R	Edit.EditorCursorMode.SelectTimeFreq
Edit Cursor Mode > Select Freq	F	F	Edit.EditorCursorMode.SelectFreq
Edit Cursor Mode > Lasso	L	L	Edit.EditorCursorMode.SelectLasso
Edit Cursor Mode > Selection Brush	B	B	Edit.EditorCursorMode.SelectBrush
Edit Cursor Mode > Selection Wand	W	W	Edit.EditorCursorMode.SelectWand
Edit Cursor Mode > Zoom Time	Z	Z	Edit.EditorCursorMode.ZoomTime
Edit Cursor Mode > Zoom Time/ Freq	shift+Z	shift+Z	Edit.EditorCursorMode.ZoomTimeFreq
Edit Cursor Mode > Zoom Freq	option+Z	alt+Z	Edit.EditorCursorMode.ZoomFreq
Edit Cursor Mode > Grab Time	G	G	Edit.EditorCursorMode.GrabTime
Edit Cursor Mode > Grab Time/Freq	shift+G	shift+G	Edit.EditorCursorMode.GrabTimeFreq
Edit Cursor Mode > Grab Freq	option+G	alt+G	Edit.EditorCursorMode.GrabFreq
Open Batch Processing window	command+B	ctrl+B	File.BatchProcessing
Send Connect Clips back to host	Command+Return	ctrl+return	File.SendConnectClipsBackToHost
Discard Connect Clips	command+delete	ctrl+backspace	File.DiscardConnectClips

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Remove Clip Gain from selection	shift+delete	shift+backspace	Editor.RemoveClipGain
Remove All Clip Gain	shift+command+delete	ctrl+shift+backspace	Editor.RemoveAllClipGain
Toggle Follow Playhead	command+P	ctrl+P	Transport.TogglePlayheadFollow
Toggle Follow Playhead Mode > Page / Continuous	shift+command+P	ctrl+shift+P	Transport.CyclePlayHeadFollowMode
Show Clip Gain	command+G	ctrl+G	View.ToggleGainCurveOverlay
Show Channels Separately	shift+command+C	ctrl+shift+C	View.ToggleCompositeAudioDisplay
Show Spectrogram Settings	shift+command+,	ctrl+shift+,	View.ToggleSpectrogramSettingsVisible
Decrease Spectrogram FFT Size	Shift+	Shift+	Spectrogram.FFTSizeDecrement
Increase Spectrogram FFT Size	Shift+.	Shift+.	Spectrogram.FFTSizeIncrement
Show Preferences Window	command+	ctrl+	Edit.Preferences
Show File Info window	shift+option+command+I	shift+alt+ctrl+I	File.Info
Enter Full Screen	^+command+F	ctrl+^+F	View.ToggleFullScreen
Exit Full Screen	esc	Esc	View.ExitFullScreen
Toggle Instant Process	I	I	Edit.EditorCursorMode.ToggleInstant
Toggle Preview Bypass	shift+B	shift+B	TogglePreviewBypass
Toggle Window Opacity	shift+command+O	ctrl+shift+O	View.ToggleFloatingWindowOpacity

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Toggle Input Monitoring	option+I	alt+I	ToggleInputMonitoring
Start or Stop Playback	Spacebar	Spacebar	Transport.PlayOrStop
Start or Stop Preview Playback	shift+Spacebar	shift+Spacebar	Transport.PreviewOrStop
Rewind Transport	return	Home	Transport.Rewind
Seek to End of file		End	Transport.SeekToEnd
Toggle Looping	command+L	ctrl+L	Transport.ToggleLooping
Toggle Playhead Follow	command+P	ctrl+P	Transport.TogglePlayHeadFollow
Toggle Playhead Return	command+R	ctrl+R	Transport.TogglePlayHeadReturn
Select Both Channels	shift+command+B	ctrl+shift+B	Editor.ChannelSelectBoth
Select Left Channel	shift+command+L	ctrl+shift+L	Editor.ChannelSelectLeft
Select Right Channel	shift+command+R	ctrl+shift+R	Editor.ChannelSelectRight
Extend selection left	shift+left Arrow	shift+left arrow	Editor.ExtendSelectionLeft
Extend selection left by page	shift+up arrow	shift+page up	Editor.ExtendSelectionPageLeft
Extend selection to the right boundary of current view	shift+down Arrow	shift+page down	Editor.ExtendSelectionPageRight
Extend selection by increment to the right	shift+right arrow	shift+right arrow	Editor.ExtendSelectionRight
Move playhead to next marker or selection boundary	option+right arrow	alt+right arrow	Editor.GoToNextMarkerOrSelectionBoundary

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Move playhead to previous marker or selection boundary	option+left arrow	alt+left arrow	Editor.GoToPreviousMarkerOrSelectionBoundary
Nudge playhead to the left	left arrow	left arrow	Editor.NudgeLeft
Nudge playhead to the right	right arrow	right arrow	Editor.NudgeRight
Page Left	page up	page up	Editor.PageLeft
Page Right	page down	page down	Editor.PageRight
Select to End	shift+end	shift+end	Editor.SelectToEnd
Select to Start	shift+home	shift+home	Editor.SelectToStart
Process Reverse	shift+R	shift+R	Process.Reverse
Process Silence	shift+S	shift+S	Process.Silence
Process Gain	option+command+6	ctrl+alt+6	Apply.Gain
Process Leveler	option+command+0	ctrl+alt+0	Apply.Leveler
Process Loudness	option+command+4	ctrl+alt+4	Apply.Loudness
Process Ambience Match	option+command+2	ctrl+alt+2	Apply.MatchAmbience
Process Mixing	option+command+8	ctrl+alt+8	Apply.ChannelMix
Process De-plosive	command+5	ctrl+5	Apply.DePlosive
Process De-click	command+2	ctrl+2	Apply.Declick
Process De-clip	command+1	ctrl+1	Apply.Declip
Process Deconstruct	command+7	ctrl+7	Apply.Deconstruct
Process De-reverb	command+8	ctrl+8	Apply.Dereverb
Process EQ	option+command+7	ctrl+alt+7	Apply.EQ

Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Process EQ Match	option+command+1	ctrl+alt+1	Apply.EQMatch
Process Plug-in	option+command+5	ctrl+alt+5	Apply.Plug-in
Process De-Hum	command+3	ctrl+3	Apply.RemoveHum
Process Resampler	option+command+9	ctrl+alt+9	Apply.Resampler
Process Spectral Repair	command+6	ctrl+6	Apply.SpectralRepair
Process Pitch Contour	option+command+3	ctrl+alt+3	Apply.TimeStretchPitchShift
Process Voice De-noise	command+4	ctrl+4	Apply.VoiceDenoise
Open Gain module	shift+option+6	shift+option+6	View.Module.ToggleGain
Open Leveler module	shift+option+0	shift+option+0	View.Module.ToggleLeveler
Open Loudness module	shift+option+4	shift+option+4	View.Module.ToggleLoudness
Open Ambience Match module	shift+option+2	shift+option+2	View.Module.ToggleMatchAmbience
Open Mixing module	shift+option+8	shift+option+8	View.Module.ToggleChannelMix
Open De-plosive module	shift+5	shift+5	View.Module.DePlosive
Open De-click module	shift+2	shift+2	View.Module.ToggleDeclick
Open De-clip module	shift+1	shift+1	View.Module.ToggleDeclip
Open Deconstruct module	shift+7	shift+7	View.Module.ToggleDeconstruct
Open De-reverb module	shift+8	shift+8	View.Module.ToggleDereverb
Open EQ module	shift+option+7	shift+option+7	View.Module.ToggleEQ

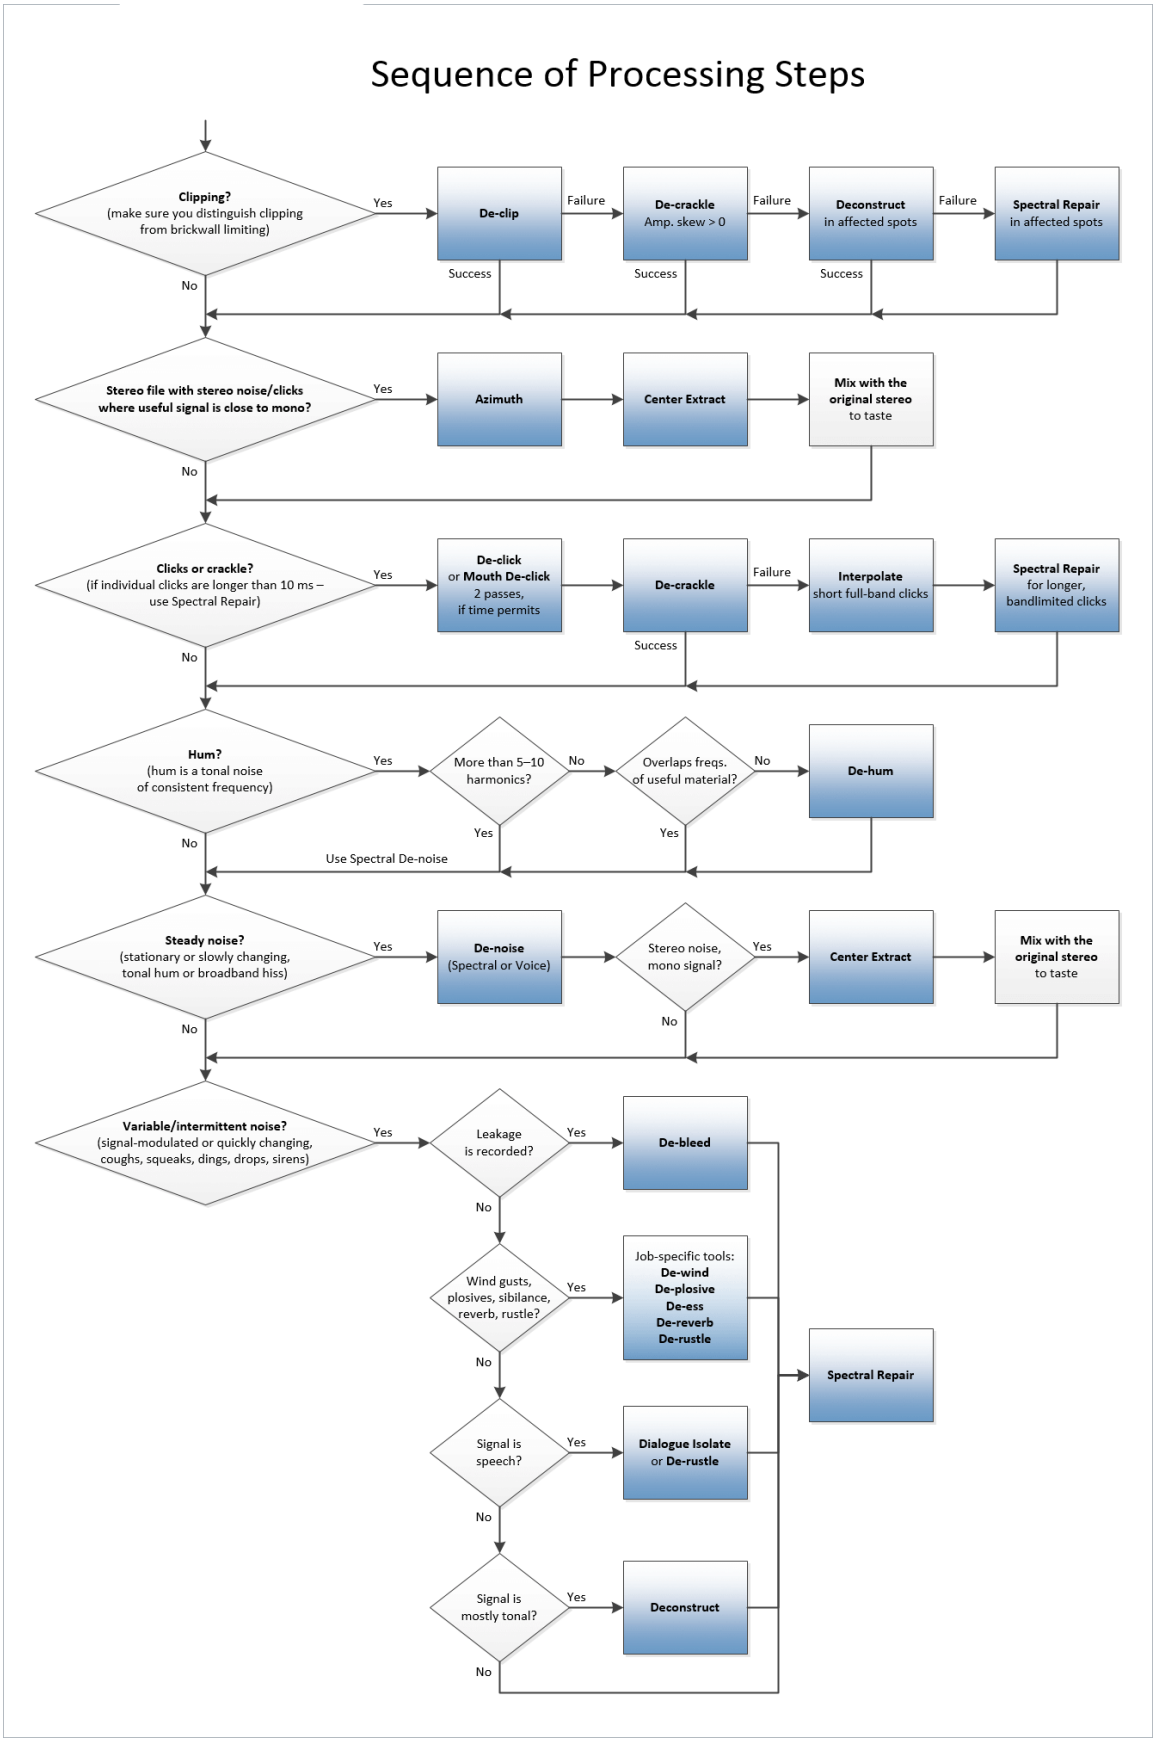
Shortcut	Default Mac Shortcut	Default Windows Shortcut	RX Shortcut Command Name
Open EQ Match module	shift+option+1	shift+option+1	View.Module.ToggleEQMatch
Open Plug-in window	shift+option+5	shift+option+5	View.Module.TogglePlugIn
Open De-hum module	shift+3	shift+3	View.Module.ToggleRemoveHum
Open Resample module	shift+option+9	shift+option+9	View.Module.ToggleSRC
Open Spectral Repair module	shift+6	shift+6	View.Module.ToggleSpectralRepair
Open Time & Pitch module	shift+option+3	shift+option+3	View.Module.ToggleTimeStretchPitchShift
Open Voice De-noise module	shift+4	shift+4	View.Module.ToggleVoiceDenoise
Open Markers window	option+m	option+m	View.ToggleMarkerPanelVisible
Open Module Chain window	c	c	View.ToggleModuleChainVisible
Open Spectrum Analyzer window	option+r	option+r	View.ToggleSpectrumAnalyzerVisible
Open Waveform Stats window	option+d	option+d	View.ToggleWaveformStatsVisible
Learn Ambience Match	shift+option+command+9	shift+alt+ctrl+9	Apply.DereverbTrain
Suggest De-hum	shift+option+command+4	shift+alt+ctrl+4	Apply.RemoveHumTrain
Learn Voice De-noise	shift+option+command+5	shift+alt+ctrl+5	Apply.VoiceDenoiseTrain
Learn EQ Match	shift+option+command+2	shift+alt+ctrl+2	Apply.EQMatchTrain

71. 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.

RX Processing Step Flowchart

To get started, it's important to identify the problems with your file and select the tool(s) that will give you the best results. We have created a process flowchart to help identify the right RX module to use and what modules to consider using if you aren't getting the results you are looking for.


Sequence of Processing Steps
Clipping?
(make sure you distinguish clipping from brickwall limiting)
Yes
No
De-clip
Failure
De-crackle
Amp. skew > 0
Failure
Deconstruct
in affected spots
Failure
Spectral Repair
in affected spots
Success
Success
Success
Stereo file with stereo noise/clicks where useful signal is close to mono?
Yes
No
Azimuth
Center Extract
Mix with the original stereo to taste
Clicks or crackle?
(if individual clicks are longer than 10 ms – use Spectral Repair)
Yes
No
De-click or Mouth De-click 2 passes, if time permits
De-crackle
Failure
Interpolate
short full-band clicks
Spectral Repair
for longer, bandlimited clicks
Success
Hum?
(hum is a tonal noise of consistent frequency)
Yes
No
More than 5–10 harmonics?
No
Overlaps freqs. of useful material?
No
De-hum
Yes
Yes
Use Spectral De-noise
Steady noise?
(stationary or slowly changing, tonal hum or broadband hiss)
Yes
No
De-noise
(Spectral or Voice)
Stereo noise, mono signal?
Yes
No
Center Extract
Mix with the original stereo to taste
Variable/intermittent noise?
(signal-modulated or quickly changing, coughs, squeaks, dings, drops, sirens)
Yes
Leakage is recorded?
Yes
De-bleed
No
Wind gusts, plosives, sibilance, reverb, rustle?
Yes
Job-specific tools:
De-wind
De-plosive
De-ess
De-reverb
De-rustle
No
Signal is speech?
Yes
Dialogue Isolate or De-rustle
No
Signal is mostly tonal?
Yes
Deconstruct
No
Spectral Repair

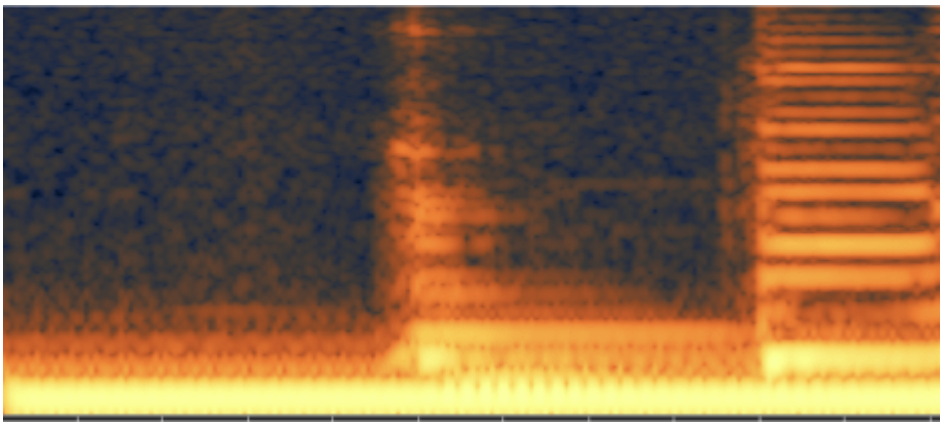
Using the Spectrogram Display

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.

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.

Hum

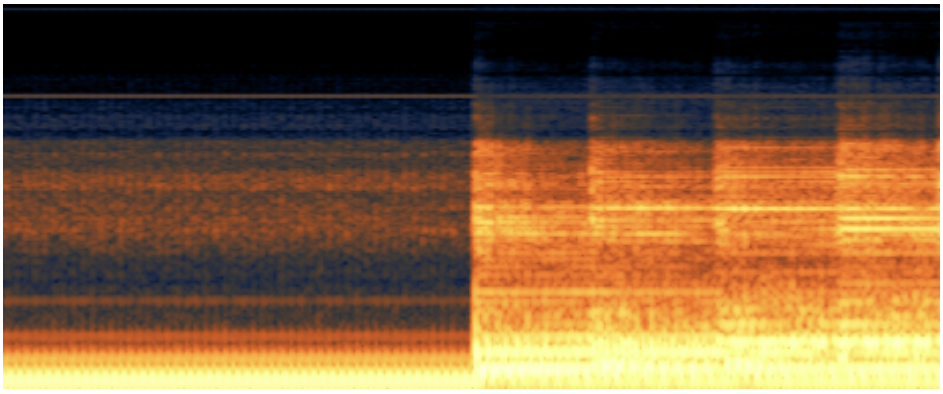
Hum is usually the result of electrical noise somewhere in the recorded signal chain. It's normally heard as a low-frequency tone based at either 50 Hz or 60 Hz depending on where the recording was made If you zoom in to the low frequencies, you'll be able to see hum as a series of horizontal lines,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



De-hum works best when frequencies of the hum do not overlap with any useful transient signals.

Buzz

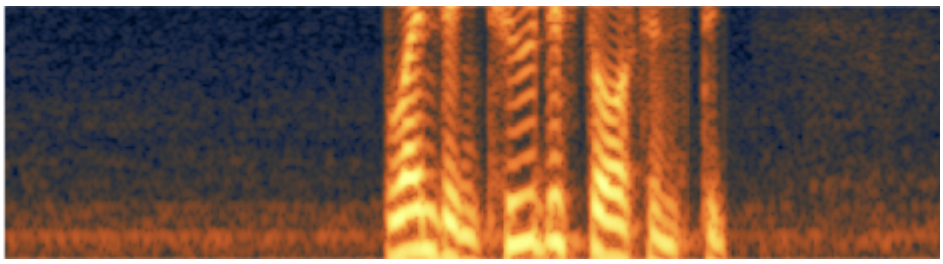
In some cases, electrical noise will extend up to higher frequencies and manifest itself as a background buzz. See the example below:



Hum-removal tools usually focus on low-frequency hum, so when the harmonics extend to frequencies above 400 Hz, the Spectral De-noise tool is often more effective at removing the problem.

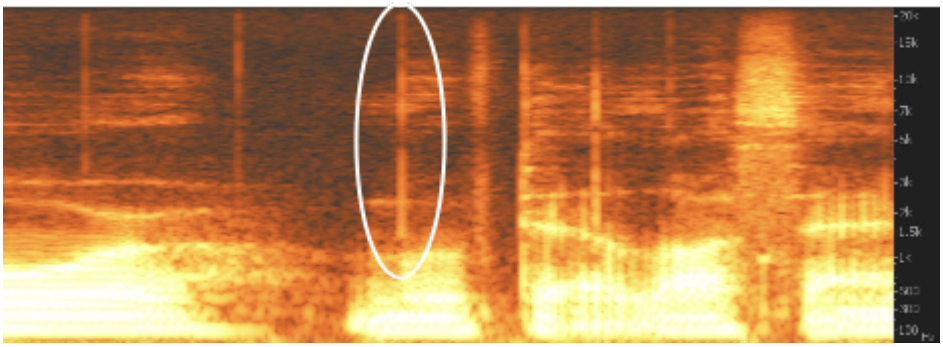
Hiss and other Broadband Noise

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 the example below:



Clicks, Pops, & Short Impulse Noises

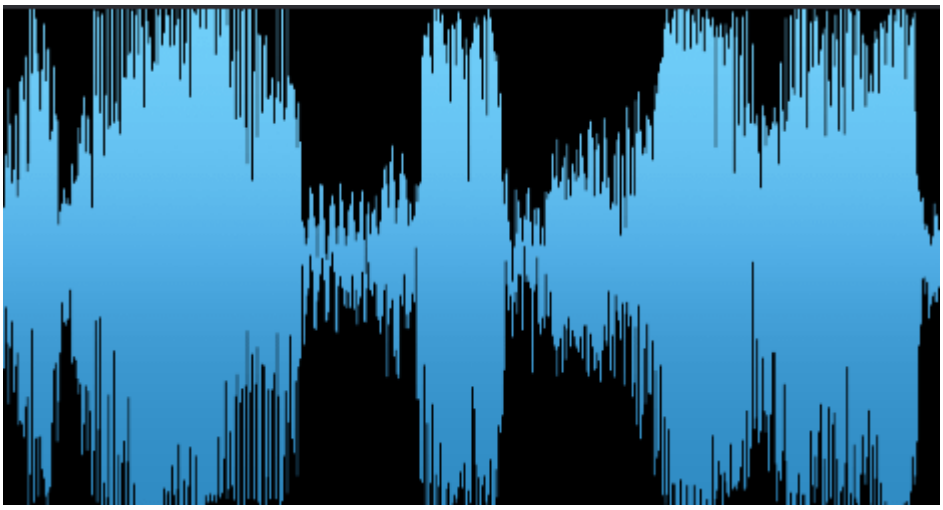
Clicks and pops are common on recordings made from vinyl, shellac and other grooved media, but can also be introduced by digital errors, including recording into a DAW with improper buffer settings, or making a bad audio edit that missed a zero crossing. Even mouth noises such as tongue clicks and lip 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:



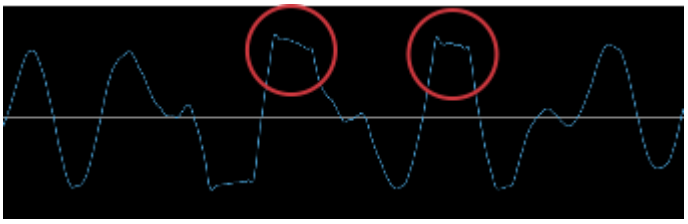
The De-click tool can recognize, isolate, and then reduce and remove clicks like these.

Clipping

Clipping is an all-too-common problem. It can occur when a loud signal distorts the input to an audio interface, analog-to-digital 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. As you'll see in the image below, the clipping appears as "squared-off" sections of the waveform.



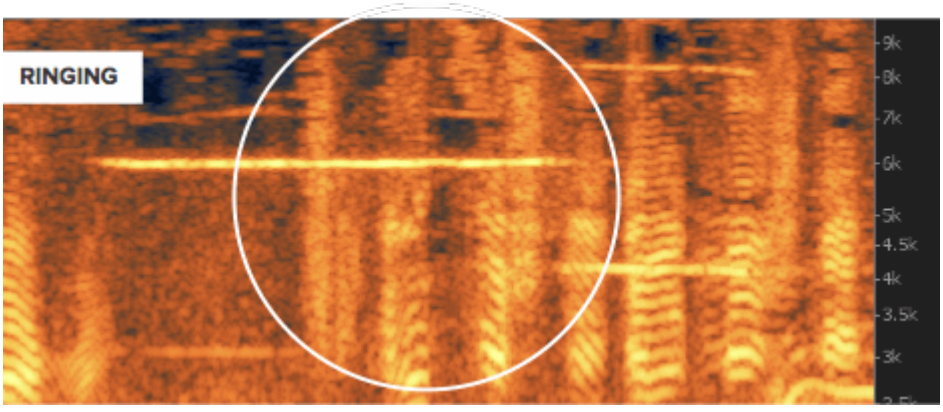
You can zoom in on a waveform and see in detail where the waveform has been truncated because of clipping.

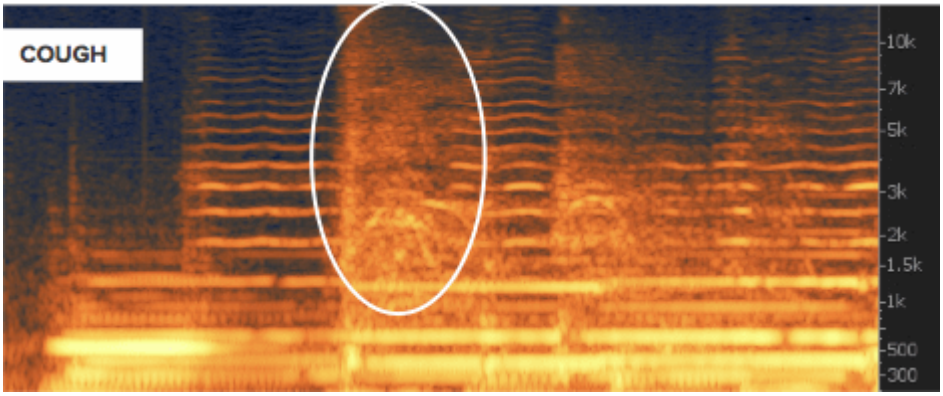


The De-clip tool can intelligently redraw the waveform to where it might have naturally been if the signal hadn't clipped. Sometimes, brickwall limited audio will also appear "squared off" when zoomed out, but this doesn't necessarily mean it will sound as heavily distorted as clipped waveforms that have been truncated. You can zoom in to see if the tops of individual waveforms 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:

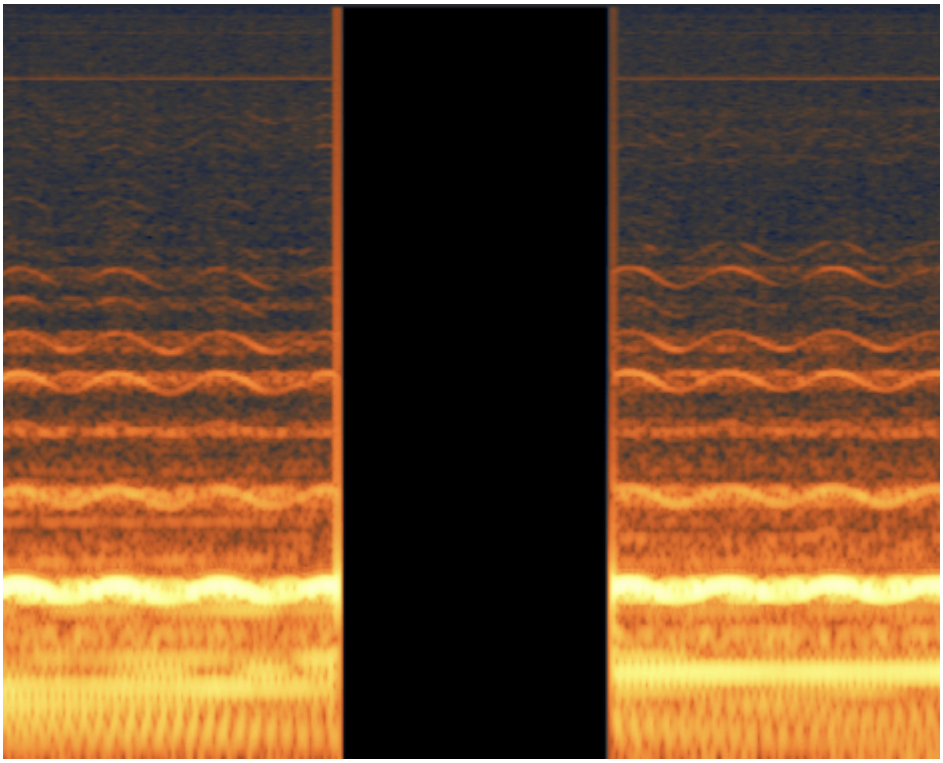




The Spectral Repair tool can help isolate these intermittent sounds, analyze the audio around them and attenuate or replace them.

Gaps and Drop Outs

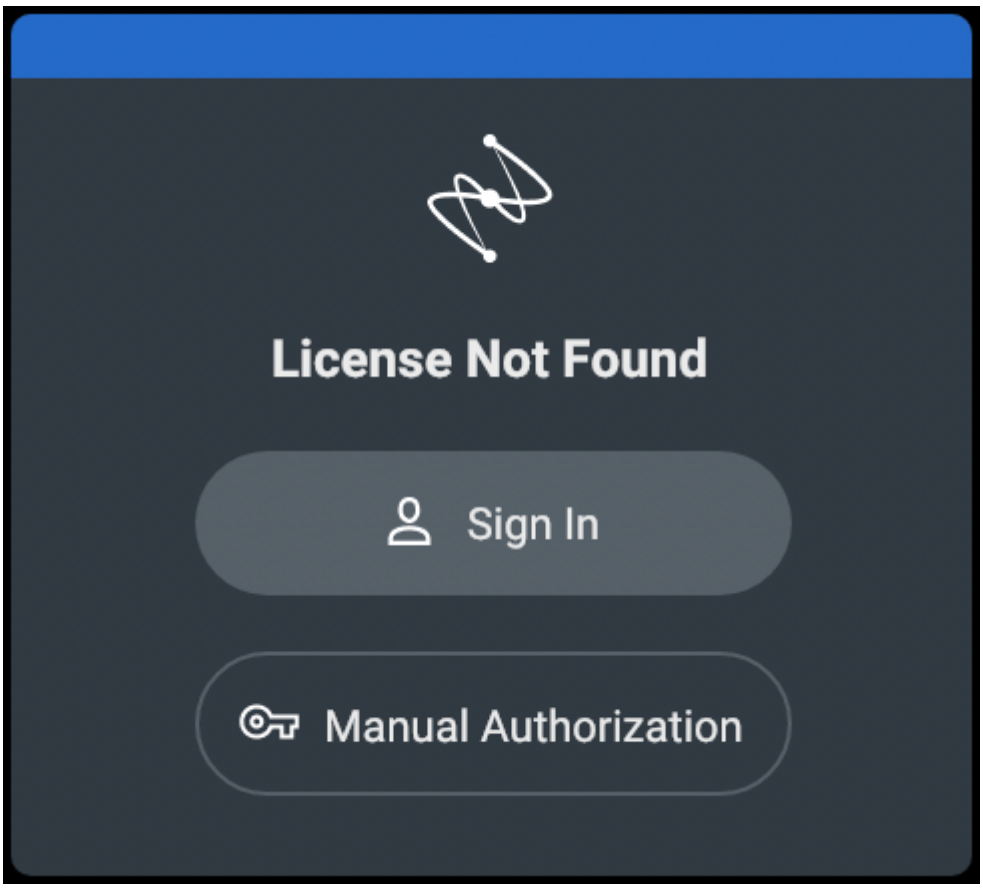
Sometimes a recording may have short sections of missing or corrupted audio. These are usually very obvious to both the eye and the ear! See the example below:



Deleting the gap and then applying Spectral Repair to replace any missing audio can help fix these problems.

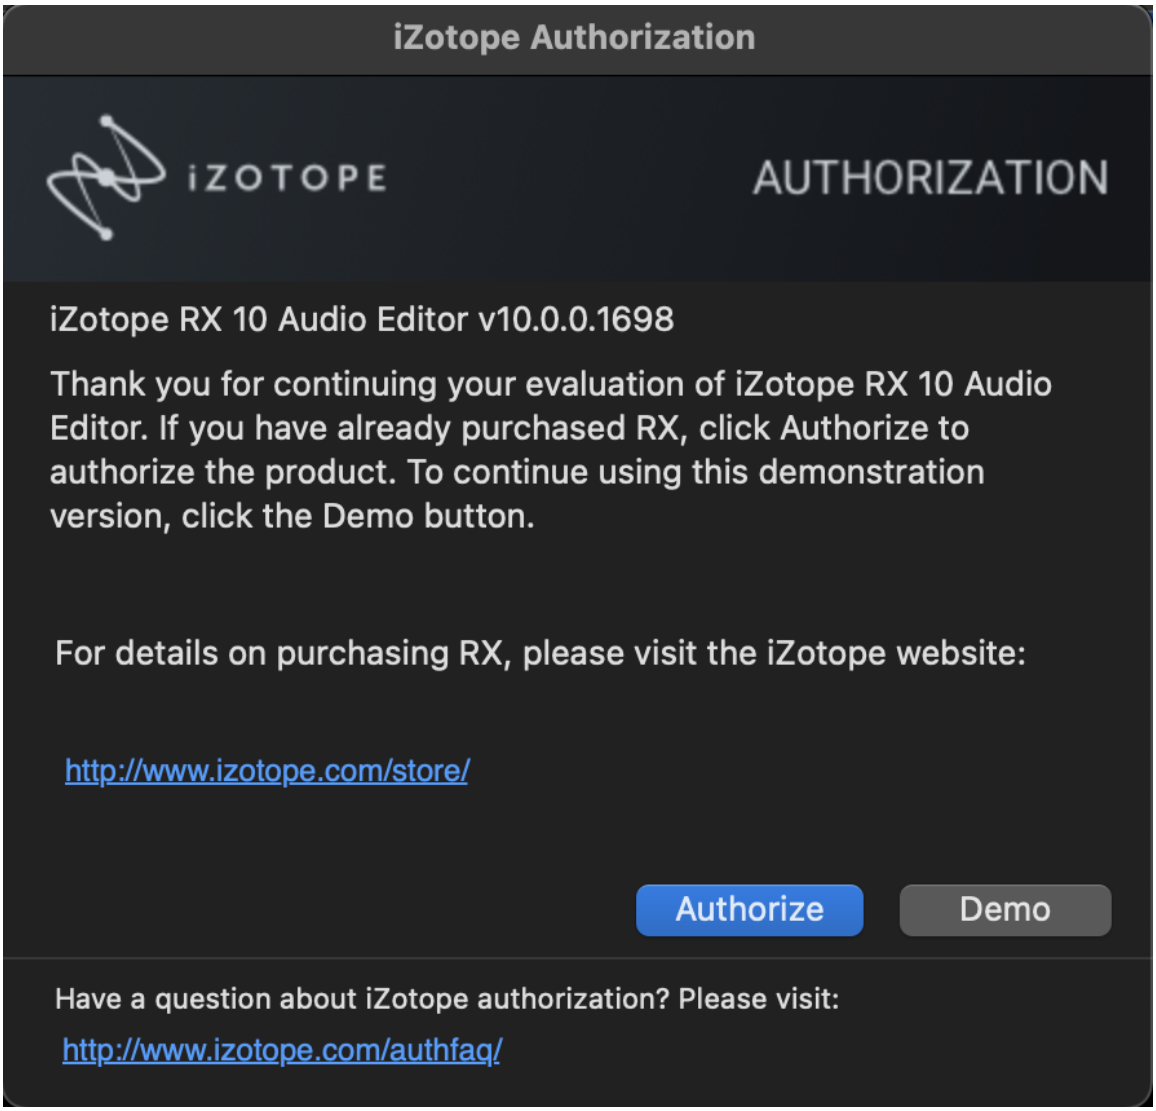
72. Authorization

The first time you open the RX 11 Audio Editor or an RX plug-in, the Authorization window will appear.



Selecting Sign In will open the Product Portal if this application is installed on your computer. If Product Portal is not installed you will be brought to the web page to download Product Portal.

Selecting Manual Authorization will prompt the manual Authorization window:



The Authorization window allows you to:

- **TRIAL:** Select the Continue in Trial option to continue a Trial period evaluation prior to purchasing.
- **DEMO:** Continue evaluating the product with Demo limitations (after the 10 day Trial period ends).
- **AUTHORIZE:** Authorize the product with a serial number.

Trial Mode

Trial mode allows you to evaluate RX 11 over a 10 day trial period. The trial period begins when you first open the RX 11 Audio Editor or an RX plug-in in a DAW/NLE. The Authorization window will display the number of days remaining in your trial period. Click the Continue button to exit the Authorization window.

Demo Mode

After your 10 day trial period expires, you have the option to operate RX 11 in Demo mode. To continue evaluating RX 11 in demo mode, click the Demo button.

DEMO MODE LIMITATIONS

All controls and user interaction will be disabled in the RX 11 audio editor and plug-ins while in Demo mode. However, playback and automation will remain intact so your existing sessions are not negatively affected.

Authorization Methods

To disable Trial or Demo mode, you must authorize the product with a valid serial number. We offer three authorization methods for RX 11:

- **Online Authorization:** Authorize RX 11 on a computer online.
- **Offline Authorization:** Authorize RX 11 on a computer offline.
- **iLok Authorization:** Authorize RX 11 using iLok.

More Authorization Help

For information about Authorization, please visit our Support Portal

73. License Information

Abseil

Copyright 2020 The Abseil Authors.

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Anti-Grain Geometry

Version 2.4

Copyright (c) 2002-2005 Maxim Shemanarev (McSeem).

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

ARA

ARA Audio Random Access is a Celemony product.



ARM_NEON_2_x86_SSE

Copyright (C) 2012-2020 Intel Corporation. All rights reserved.

IMPORTANT: READ BEFORE DOWNLOADING, COPYING, INSTALLING OR USING.

By downloading, copying, installing or using the software you agree to this license. If you do not agree to this license, do not download, install, copy or use the software.

License Agreement

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- The name of the copyright holders may not be used to endorse or promote products derived from this software without specific prior written permission.

This software is provided by the copyright holders and contributors “as is” and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the Intel Corporation or contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

base64

v0.4.0

Copyright (c) 2005-2007, Nick Galbreath

Copyright (c) 2013-2019, Alfred Klomp

Copyright (c) 2015-2017, Wojciech Mula

Copyright (c) 2016-2017, Matthieu Darbois

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Better Enums

Version 0.11.1

Copyright (c) 2012-2016, Anton Bachin. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Boost

Boost Software License - Version 1.0 - August 17th, 2003

Permission is hereby granted, free of charge, to any person or organization obtaining a copy of the software and accompanying documentation covered by this license (the “Software”) to use, reproduce, display, distribute, execute, and transmit the Software, and to prepare derivative works of the Software, and to permit third-parties to whom the Software is furnished to do so, all subject to the following:

The copyright notices in the Software and this entire statement, including the above license grant, this restriction and the following disclaimer, must be included in all copies of the Software, in whole or in part, and all derivative works of the Software, unless such copies or derivative works are solely in the form of machine-executable object code generated by a source language processor.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR ANYONE DISTRIBUTING THE SOFTWARE BE LIABLE FOR ANY DAMAGES OR OTHER LIABILITY, WHETHER IN CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Bravura

Copyright © 2015, Steinberg Media Technologies GmbH http://www.steinberg.net/, with Reserved Font Name "Bravura".

This Font Software is licensed under the SIL Open Font License, Version 1.1. This license is copied below, and is also available with a FAQ at: http://scripts.sil.org/OFL

SIL OPEN FONT LICENSE Version 1.1 - 26 February 2007

PREAMBLE

The goals of the Open Font License (OFL) are to stimulate worldwide development of collaborative font projects, to support the font creation efforts of academic and linguistic communities, and to provide a free and open framework in which fonts may be shared and improved in partnership with others.

The OFL allows the licensed fonts to be used, studied, modified and redistributed freely as long as they are not sold by themselves. The fonts, including any derivative works, can be bundled, embedded, redistributed and/or sold with any software provided that any reserved names are not used by derivative works. The fonts and derivatives, however, cannot be released under any other type of license. The requirement for fonts to remain under this license does not apply to any document created using the fonts or their derivatives.

DEFINITIONS

"Font Software" refers to the set of files released by the Copyright Holder(s) under this license and clearly marked as such. This may include source files, build scripts and documentation.

"Reserved Font Name" refers to any names specified as such after the copyright statement(s).

"Original Version" refers to the collection of Font Software components as distributed by the Copyright Holder(s).

"Modified Version" refers to any derivative made by adding to, deleting, or substituting – in part or in whole – any of the components of the Original Version, by changing formats or by porting the Font Software to a new environment.

"Author" refers to any designer, engineer, programmer, technical writer or other person who contributed to the Font Software.

PERMISSION & CONDITIONS

Permission is hereby granted, free of charge, to any person obtaining a copy of the Font Software, to use, study, copy, merge, embed, modify, redistribute, and sell modified and unmodified copies of the Font Software, subject to the following conditions:

1. Neither the Font Software nor any of its individual components, in Original or Modified Versions, may be sold by itself.
2. Original or Modified Versions of the Font Software may be bundled, redistributed and/or sold with any software, provided that each copy contains the above copyright notice and this license. These can be included either as stand-alone text files, human-readable headers or in the appropriate machine-readable metadata fields within text or binary files as long as those fields can be easily viewed by the user.

3. No Modified Version of the Font Software may use the Reserved Font Name(s) unless explicit written permission is granted by the corresponding Copyright Holder. This restriction only applies to the primary font name as presented to the users.
4. The name(s) of the Copyright Holder(s) or the Author(s) of the Font Software shall not be used to promote, endorse or advertise any Modified Version, except to acknowledge the contribution(s) of the Copyright Holder(s) and the Author(s) or with their explicit written permission.
5. The Font Software, modified or unmodified, in part or in whole, must be distributed entirely under this license, and must not be distributed under any other license. The requirement for fonts to remain under this license does not apply to any document created using the Font Software.

TERMINATION

This license becomes null and void if any of the above conditions are not met.

DISCLAIMER

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

C++ Rest SDK

Version 2.10.15

Main Library:

Copyright (c) 2014, Peter Thorson. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the WebSocket++ Project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL PETER THORSON BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Bundled Libraries:

****** Base 64 Library (base64/base64.hpp) ******

base64.hpp is a repackaging of the base64.cpp and base64.h files into a single header suitable for use as a header only library. This conversion was done by Peter Thorson (webmaster@zaphoyd.com) in 2012. All modifications to the code are redistributed under the same license as the original, which is listed below.

base64.cpp and base64.h

Copyright (C) 2004-2008 René Nyffenegger

This source code is provided 'as-is', without any express or implied warranty. In no event will the author be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this source code must not be misrepresented; you must not claim that you wrote the original source code. If you use this source code in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original source code.
3. This notice may not be removed or altered from any source distribution.

René Nyffenegger rene.nyffenegger@adp-gmbh.ch

****** SHA1 Library (sha1/sha1.hpp) ******

sha1.hpp is a repackaging of the sha1.cpp and sha1.h files from the shallsha1 library (http://code.google.com/p/smallsha1/) into a single header suitable for use as a header only library. This conversion was done by Peter Thorson (webmaster@zaphoyd.com) in 2013. All modifications to the code are redistributed under the same license as the original, which is listed below.

Copyright (c) 2011, Micael Hildenborg

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Micael Hildenborg nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY Micael Hildenborg "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL Micael Hildenborg BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

****** MD5 Library (common/md5.hpp) ******

md5.hpp is a reformulation of the md5.h and md5.c code from http://www.opensource.apple.com/source/cups/cups-59/cups/md5.c to allow it to function as a component of a header only library. This conversion was done by Peter Thorson (webmaster@zaphoyd.com) in 2012 for the WebSocket++ project. The changes are released under the same license as the original (listed below)

Copyright (C) 1999, 2002 Aladdin Enterprises. All rights reserved.

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

L. Peter Deutsch

ghost@aladdin.com

****** UTF8 Validation logic (utf8_validation.hpp) ******

utf8_validation.hpp is adapted from code originally written by Bjoern Hoehrmann bjoern@hoehrmann.de . See http://bjoern.hoehrmann.de/utf-8/decoder/dfa/ for details.

The original license:

Copyright (c) 2008-2009 Bjoern Hoehrmann bjoern@hoehrmann.de

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Caffe

COPYRIGHT

All contributions by the University of California:

Copyright (c) 2014-2017 The Regents of the University of California (Regents)

All rights reserved.

All other contributions:

Copyright (c) 2014-2017, the respective contributors

All rights reserved.

Caffe uses a shared copyright model: each contributor holds copyright over their contributions to Caffe. The project versioning records all such contribution and copyright details. If a contributor wants to further mark their specific copyright on a particular contribution, they should indicate their copyright solely in the commit message of the change when it is committed.

LICENSE

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

CONTRIBUTION AGREEMENT

By contributing to the BVLC/caffe repository through pull-request, comment, or otherwise, the contributor releases their content to the license and copyright terms herein.

cerberus

Cerberus is a lightweight and extensible data validation library for Python.

ISC License

Copyright (c) 2012-2016 Nicola Iarocci.

Permission to use, copy, modify, and/or distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

composable_kernel

https://github.com/ROCmSoftwarePlatform/composable_kernel

Copyright (c) 2018- , Advanced Micro Devices, Inc. (Chao Liu, Jing Zhang)

Copyright (c) 2019- , Advanced Micro Devices, Inc. (Letao Qin, Qianfeng Zhang, Liang Huang, Shaojie Wang)

Copyright (c) 2022- , Advanced Micro Devices, Inc. (Anthony Chang, Chunyu Lai, Illia Silin, Adam Osewski, Poyen Chen, Jehandad Khan)

Copyright (c) 2019-2021, Advanced Micro Devices, Inc. (Hanwen Chang)

Copyright (c) 2019-2020, Advanced Micro Devices, Inc. (Tejash Shah)

Copyright (c) 2020 , Advanced Micro Devices, Inc. (Xiaoyan Zhou)

Copyright (c) 2021-2022, Advanced Micro Devices, Inc. (Jianfeng Yan)

SPDX-License-Identifier: MIT

Copyright (c) 2018-2023, Advanced Micro Devices, Inc. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

coremltools

BSD-3-Clause License

https://github.com/apple/coremltools

Copyright (c) 2020, Apple Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holder(s) nor the names of any contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

© 2021 GitHub, Inc.

CpuInfo

Copyright (c) 2019 Google LLC

Copyright (c) 2017-2018 Facebook Inc.

Copyright (c) 2012-2017 Georgia Institute of Technology

Copyright (c) 2010-2012 Marat Dukhan

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Crashpad

Version 0.8.0

Copyright 2014 The Crashpad Authors. All rights reserved.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

curl

https://github.com/curl

COPYRIGHT AND PERMISSION NOTICE

Copyright (C) Daniel Stenberg, daniel@haxx.se, and many contributors, see the THANKS file.

All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

dcleblanc/SafeInt

MIT License

Copyright (c) 2018 Microsoft

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Distributed Machine Learning Common Codebase

Copyright (c) 2015 by Contributors

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

dlfcn-win32

dlfcn-win32 is licensed under the MIT License.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

DLPack

https://github.com/dmlc/dlpack

Copyright 2017 by Contributors

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

DLPack: Open In Memory Tensor Structure

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Eigen

Distributed under the Mozilla Public License v2.0 (MPLv2.0).

Full text of the license is available here: https://www.mozilla.org/en-US/MPL/2.0/

To receive a copy of the source code for the Eigen library distributed with this product under the under the terms of the MPLv2.0 please contact dev-support@native-instruments.com.

emsdk

https://github.com/emscripten-core/emsdk/

Copyright (c) 2018 Emscripten authors (see AUTHORS in Emscripten)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

This is the MIT/Expat Licence. For more information see:

1. http://www.opensource.org/licenses/mit-license.php
2. http://en.wikipedia.org/wiki/MIT_License

FarmHash

Copyright (c) 2014 Google, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

FLAC

libFLAC and libFLAC++

Version 1.3.2

Copyright (c) 2000-2009 Josh Coalson

Copyright (c) 2011-2016 Xiph.Org Foundation

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Xiph.org Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

FlashAttention

FlashAttention, https://github.com/Dao-AILab/flash-attention

BSD 3-Clause License

Copyright (c) 2022, the respective contributors, as shown by the AUTHORS file. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

FlatBuffers

Copyright 2014 Google Inc. All rights reserved.

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

FreeBSD: getopt.c file

Copyright (c) 1987, 1993, 1994

The Regents of the University of California.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

FreeType

Version 2.4.6

Portions of this software are copyright © 2011 The FreeType Project (www.freetype.org). All rights reserved.

FXDev

The MIT License (MIT)

Copyright (c) 2017 Facebook Inc.

Copyright (c) 2016-2017 Marat Dukhan

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

G3log: Asynchronous logger with Dynamic Sinks

This is free and unencumbered software released into the public domain.

Anyone is free to copy, modify, publish, use, compile, sell, or distribute this software, either in source code form or as a compiled binary, for any purpose, commercial or non-commercial, and by any means.

In jurisdictions that recognize copyright laws, the author or authors of this software dedicate any and all copyright interest in the software to the public domain. We make this dedication for the benefit of the public at large and to the detriment of our heirs and successors. We intend this dedication to be an overt act of relinquishment in perpetuity of all present and future rights to this software under copyright law.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

For more information, please refer to http://unlicense.org/

gemmlowp

Copyright 2021 gemmlowp authors

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

GLEW

The OpenGL Extension Wrangler Library

Copyright (C) 2002-2008, Milan Ikits <milan ikits[]ieee org>

Copyright (C) 2002-2008, Marcelo E. Magallon <mmagallo[]debian org>

Copyright (C) 2002, Lev Povalahev

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of the author may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Mesa 3-D graphics library Version: 7.0

Copyright (C) 1999-2007 Brian Paul All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL BRIAN PAUL BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Copyright (c) 2007 The Khronos Group Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and/or associated documentation files (the "Materials"), to deal in the Materials without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Materials, and to permit persons to whom the Materials are furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Materials.

THE MATERIALS ARE PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE MATERIALS OR THE USE OR OTHER DEALINGS IN THE MATERIALS.

google/glog

Copyright (c) 2008, Google Inc.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

A function gettimeofday in utilities.cc is based on

http://www.google.com/codesearch/p?hl=en#dR3YEbitojA/COPYING&q=GetSystemTimeAsFileTime%20license:bsd

The license of this code is:

Copyright (c) 2003-2008, Jouni Malinen j@w1.fi and contributors

All Rights Reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name(s) of the above-listed copyright holder(s) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

google/nsync

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

google/re2

Copyright (c) 2009 The RE2 Authors. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

google/sentencepiece

https://github.com/google/sentencepiece

(included when statically linked with onnxruntime-extensions)

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

google/XNNPACK

BSD License

For XNNPACK software

Copyright © Facebook, Inc. and its affiliates. All rights reserved.

Copyright 2019 Google LLC

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name Facebook nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

gsl

Copyright (c) 2015 Microsoft Corporation. All rights reserved.

This code is licensed under the MIT License (MIT).

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

HalideIR

Copyright (c) 2016 HalideIR contributors

Copyright (c) 2012-2014 MIT CSAIL, Google Inc., and other contributors

HalideIR is derived from the Halide project.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

HowardHinnant/date

The source code in this project is released using the MIT License. There is no global license for the project because each file is licensed individually with different author names and/or dates.

If you contribute to this project, please add your name to the license of each file you modify. If you have already contributed to this project and forgot to add your name to the license, please feel free to submit a new P/R to add your name to the license in each file you modified.

For convenience, here is a copy of the MIT license found in each file except without author names or dates:

The MIT License (MIT)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

huggingface/transformers

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

IcoMoon

IcoMoon-Free licensed under: CC BY 4.0

Intel Math Kernel Library (Intel MKL)

Intel Simplified Software License (Version April 2018)

Copyright (c) 2018 Intel Corporation.

Use and Redistribution.

You may use and redistribute the software (the "Software"), without modification, provided the following conditions are met:

- Redistributions must reproduce the above copyright notice and the following terms of use in the Software and in the documentation and/or other materials provided with the distribution.
- Neither the name of Intel nor the names of its suppliers may be used to endorse or promote products derived from this Software without specific prior written permission.
- No reverse engineering, decompilation, or disassembly of this Software is permitted.

Limited patent license. Intel grants you a world-wide, royalty-free, non-exclusive license under patents it now or hereafter owns or controls to make, have made, use, import, offer to sell and sell ("Utilize") this Software, but solely to the extent that any such patent is necessary to Utilize the Software alone. The patent license shall not apply to any combinations which include this software. No hardware per se is licensed hereunder.

Third party and other Intel programs. "Third Party Programs" are the files listed in the "third-party-programs.txt" text file that is included with the Software and may include Intel programs under separate license terms. Third Party Programs, even if included with the distribution of the Materials, are governed by separate license terms and those license terms solely govern your use of those programs.

DISCLAIMER. THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT ARE DISCLAIMED. THIS SOFTWARE IS NOT INTENDED FOR USE IN SYStemS OR APPLICATIONS WHERE FAILURE OF THE SOFTWARE MAY CAUSE PERSONAL INJURY OR DEATH AND YOU AGREE THAT YOU ARE FULLY RESPONSIBLE FOR ANY CLAIMS, COSTS, DAMAGES, EXPENSES, AND ATTORNEYS' FEES ARISING OUT OF ANY SUCH USE, EVEN IF ANY CLAIM ALLEGES THAT INTEL WAS NEGLIGENT REGARDING THE DESIGN OR MANUFACTURE OF THE MATERIALS.

LIMITATION OF LIABILITY. IN NO EVENT WILL INTEL BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. YOU AGREE TO INDEMNIFY AND HOLD INTEL HARMLESS AGAINST ANY CLAIMS AND EXPENSES RESULTING FROM YOUR USE OR UNAUTHORIZED USE OF THE SOFTWARE.

No support. Intel may make changes to the Software, at any time without notice, and is not obligated to support, update or provide training for the Software.

Termination. Intel may terminate your right to use the Software in the event of your breach of this Agreement and you fail to cure the breach within a reasonable period of time.

Feedback. Should you provide Intel with comments, modifications, corrections, enhancements or other input ("Feedback") related to the Software Intel will be free to use, disclose, reproduce, license or otherwise distribute or exploit the Feedback in its sole discretion without any obligations or restrictions of any kind, including without limitation, intellectual property rights or licensing obligations.

Compliance with laws. You agree to comply with all relevant laws and regulations governing your use, transfer, import or export (or prohibition thereof) of the Software.

Governing law. All disputes will be governed by the laws of the United States of America and the State of Delaware without reference to conflict of law principles and subject to the exclusive jurisdiction of the state or federal courts sitting in the State of Delaware, and each party agrees that it submits to the personal jurisdiction and venue of those courts and waives any objections. The United Nations Convention on Contracts for the International Sale of Goods (1980) is specifically excluded and will not apply to the Software.

*Other names and brands may be claimed as the property of others.

Intel neural-compressor

https://github.com/intel/neural-compressor

Copyright 2016-2019 Intel Corporation

Copyright 2018 YANDEX LLC

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

intel/dnnl

Copyright 2016-2018 Intel Corporation

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

sub-components:

xbyak

Copyright (c) 2007 MITSUNARI Shigeo. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of the copyright owner nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Intel® Integrated Performance Primitives (Intel® IPP)

Version 2019.0.5

LIMITATION OF LIABILITY. IN NO EVENT WILL INTEL BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. YOU AGREE TO INDEMNIFY AND HOLD INTEL HARMLESS AGAINST ANY CLAIMS AND EXPENSES RESULTING FROM YOUR USE OR UNAUTHORIZED USE OF THE SOFTWARE.

No support. Intel may make changes to the Software, at any time without notice, and is not obligated to support, update or provide training for the Software.

Termination. Intel may terminate your right to use the Software in the event of your breach of this Agreement and you fail to cure the breach within a reasonable period of time.

Feedback. Should you provide Intel with comments, modifications, corrections, enhancements or other input ("Feedback") related to the Software Intel will be free to use, disclose, reproduce, license or otherwise distribute or exploit the Feedback in its sole discretion without any obligations or restrictions of any kind, including without limitation, intellectual property rights or licensing obligations.

Compliance with laws. You agree to comply with all relevant laws and regulations governing your use, transfer, import or export (or prohibition thereof) of the Software.

Governing law. All disputes will be governed by the laws of the United States of America and the State of Delaware without reference to conflict of law principles and subject to the exclusive jurisdiction of the state or federal courts sitting in the State of Delaware, and each party agrees that it submits to the personal jurisdiction and venue of those courts and waives any objections. The United Nations Convention on Contracts for the International Sale of Goods (1980) is specifically excluded and will not apply to the Software.

Other names and brands may be claimed as the property of others.

JDAI-CV/DNNLibrary

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

JsonCpp

Version 1.2.1

Copyright (c) 2007-2010 Baptiste Lepilleur and The JsonCpp Authors

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

LAME

libmp3lame

Version 3.99.5

Copyright (c) 1999-2011 The L.A.M.E. Team

Distributed under the GNU Library General Public License v2 (LGPLv2).

Full text of the license is available here: https://www.gnu.org/licenses/old-licenses/lgpl-2.0.txt

To receive a copy of the source code for the LAME library distributed with this product under the under the terms of the LGPLv2 please contact dev-support@native-instruments.com.

libb64

https://github.com/libb64/libb64

Copyright-Only Dedication (based on United States law) or Public Domain Certification

The person or persons who have associated work with this document (the “Dedicator” or “Certifier”) hereby either (a) certifies that, to the best of his knowledge, the work of authorship identified is in the public domain of the country from which the work is published, or (b) hereby dedicates whatever copyright the dedicators holds in the work of authorship identified below (the “Work”) to the public domain. A certifier, moreover, dedicates any copyright interest he may have in the associated work, and for these purposes, is described as a “dedicator” below.

A certifier has taken reasonable steps to verify the copyright status of this work. Certifier recognizes that his good faith efforts may not shield him from liability if in fact the work certified is not in the public domain.

Dedicator makes this dedication for the benefit of the public at large and to the detriment of the Dedicator’s heirs and successors. Dedicator intends this dedication to be an overt act of relinquishment in perpetuity of all present and future rights under copyright law, whether vested or contingent, in the Work. Dedicator understands that such relinquishment of all rights includes the relinquishment of all rights to enforce (by lawsuit or otherwise) those copyrights in the Work.

Dedicator recognizes that, once placed in the public domain, the Work may be freely reproduced, distributed, transmitted, used, modified, built upon, or otherwise exploited by anyone for any purpose, commercial or non-commercial, and in any way, including by methods that have not yet been invented or conceived.

libprotobuf-mutator

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

LibXML2

Version 2.7.8

Except where otherwise noted in the source code (e.g. the files hash.c, list.c and the trio files, which are covered by a similar licence but with different Copyright notices) all the files are:

Copyright (C) 1998-2003 Daniel Veillard. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE DANIEL VEILLARD BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Daniel Veillard shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from him.

madler/zlib

The deflate format used by zlib was defined by Phil Katz. The deflate and zlib specifications were written by L. Peter Deutsch. Thanks to all the people who reported problems and suggested various improvements in zlib; they are too numerous to cite here.

Copyright notice:

(C) 1995-2017 Jean-loup Gailly and Mark Adler

This software is provided ‘as-is’, without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler jloup@gzip.org madler@alumni.caltech.edu

If you use the zlib library in a product, we would appreciate *not* receiving lengthy legal documents to sign. The sources are provided for free but without warranty of any kind. The library has been entirely written by Jean-loup Gailly and Mark Adler; it does not include third-party code.

If you redistribute modified sources, we would appreciate that you include in the file ChangeLog history information documenting your changes. Please read the FAQ for more information on the distribution of modified source versions.

magic_enum

MIT License

Copyright (c) 2019 - 2023 Daniil Goncharov

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Microsoft Cognitive Toolkit (CNTK)

Copyright (c) Microsoft Corporation. All rights reserved.

MIT License

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

microsoft/mimalloc

microsoft/mimalloc, https://github.com/microsoft/mimalloc

MIT License

Copyright (c) 2018-2021 Microsoft Corporation, Daan Leijen

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

microsoft/wil

MIT License

Copyright (c) Microsoft Corporation. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE

mpi4py

https://github.com/mpi4py/mpi4py/

==========

LICENSE: MPI for Python

==========

:Author: Lisandro Dalcin

:Contact: dalcinl@gmail.com

Copyright (c) 2019, Lisandro Dalcin.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

MurmurHash3

MIT license

https://github.com/aappleby/smhasher

SMHasher is a test suite designed to test the distribution, collision, and performance properties of non-cryptographic hash functions.

This is the home for the MurmurHash family of hash functions along with the SMHasher test suite used to verify them.

SMHasher is released under the MIT license. All MurmurHash versions are public domain software, and the author disclaims all copyright to their code.

Netlib numeralgo na10 Aberth’s method

All the software contained in this library is protected by copyright. Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

THIS SOFTWARE IS BEING PROVIDED “AS IS”, WITHOUT ANY EXPRESS OR IMPLIED WARRANTY. IN NO EVENT, NEITHER THE AUTHORS, NOR THE PUBLISHER, NOR ANY MEMBER OF THE EDITORIAL BOARD OF THE JOURNAL “NUMERICAL ALGORITHMS”, NOR ITS EDITOR-IN-CHIEF, BE LIABLE FOR ANY ERROR IN THE SOFTWARE, ANY MISUSE OF IT OR ANY DAMAGE ARISING OUT OF ITS USE. THE ENTIRE RISK OF USING THE SOFTWARE LIES WITH THE PARTY DOING SO.

ANY USE OF THE SOFTWARE CONSTITUTES ACCEPTANCE OF THE TERMS OF THE ABOVE STATEMENT.

AUTHOR:

- DARIO ANDREA
  UNIVERSITY OF PISA, ITALY
  E-MAIL: bini@dm.unipi.it

REFERENCE:

- NUMERICAL COMPUTATION OF POLYNOMIAL ZEROS BY MEANS OF ABERTH’S METHOD NUMERICAL ALGORITHMS, 13 (1996), PP. 179-200

SOFTWARE REVISION DATE:

- JUNE, 1996

SOFTWARE LANGUAGE:

- FORTRAN

nlohmann/json

MIT License

Copyright (c) 2013-2019 Niels Lohmann

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

nvidia/cutlass

Copyright (c) 2017 - 2022 NVIDIA CORPORATION & AFFILIATES. All rights reserved.

SPDX-License-Identifier: BSD-3-Clause

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

OGG / Vorbis

libogg and libvorbis

Version 1.3.2

Copyright (c) 2014, Xiph.org Foundation

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/ or other materials provided with the distribution.
- Neither the name of the Xiph.org Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

onnx

Open Neural Network Exchange

Copyright (c) Facebook, Inc. and Microsoft Corporation. All rights reserved.

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

onnx/onnx-tensorrt

MIT License

Copyright (c) 2018, NVIDIA CORPORATION. All rights reserved.

Copyright (c) 2018 Open Neural Network Exchange

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

onnxruntime

MIT License

Copyright (c) Microsoft Corporation

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

THIRD PARTY SOFTWARE NOTICES AND INFORMATION

Do Not Translate or Localize

This software incorporates material from third parties. Microsoft makes certain open source code available at http://3rdpartysource.microsoft.com, or you may send a check or money order for US $5.00, including the product name, the open source component name, and version number, to:

Source Code Compliance Team

Microsoft Corporation

One Microsoft Way

Redmond, WA 98052

USA

Notwithstanding any other terms, you may reverse engineer this software to the extent required to debug changes to any libraries licensed under the GNU Lesser General Public License.

Open MPI

3-Clause BSD License

Most files in this release are marked with the copyrights of the organizations who have edited them. The copyrights below are in no particular order and generally reflect members of the Open MPI core team who have contributed code to this release. The copyrights for code used under license from other parties are included in the corresponding files.

Copyright (c) 2004-2010 The Trustees of Indiana University and Indiana University Research and Technology Corporation. All rights reserved.

Copyright (c) 2004-2017 The University of Tennessee and The University of Tennessee Research Foundation. All rights reserved.

Copyright (c) 2004-2010 High Performance Computing Center Stuttgart, University of Stuttgart. All rights reserved.

Copyright (c) 2004-2008 The Regents of the University of California. All rights reserved.

Copyright (c) 2006-2017 Los Alamos National Security, LLC. All rights reserved.

Copyright (c) 2006-2017 Cisco Systems, Inc. All rights reserved.

Copyright (c) 2006-2010 Voltaire, Inc. All rights reserved.

Copyright (c) 2006-2017 Sandia National Laboratories. All rights reserved.

Copyright (c) 2006-2010 Sun Microsystems, Inc. All rights reserved. Use is subject to license terms.

Copyright (c) 2006-2017 The University of Houston. All rights reserved.

Copyright (c) 2006-2009 Myricom, Inc. All rights reserved.

Copyright (c) 2007-2017 UT-Battelle, LLC. All rights reserved.

Copyright (c) 2007-2017 IBM Corporation. All rights reserved.

Copyright (c) 1998-2005 Forschungszentrum Juelich, Juelich Supercomputing Centre, Federal Republic of Germany

Copyright (c) 2005-2008 ZIH, TU Dresden, Federal Republic of Germany

Copyright (c) 2007 Evergrid, Inc. All rights reserved.

Copyright (c) 2008 Chelsio, Inc. All rights reserved.

Copyright (c) 2008-2009 Institut National de Recherche en Informatique. All rights reserved.

Copyright (c) 2007 Lawrence Livermore National Security, LLC. All rights reserved.

Copyright (c) 2007-2017 Mellanox Technologies. All rights reserved.

Copyright (c) 2006-2010 QLogic Corporation. All rights reserved.

Copyright (c) 2008-2017 Oak Ridge National Labs. All rights reserved.

Copyright (c) 2006-2012 Oracle and/or its affiliates. All rights reserved.

Copyright (c) 2009-2015 Bull SAS. All rights reserved.

Copyright (c) 2010 ARM ltd. All rights reserved.

Copyright (c) 2016 ARM, Inc. All rights reserved.

Copyright (c) 2010-2011 Alex Brick. All rights reserved.

Copyright (c) 2012 The University of Wisconsin-La Crosse. All rights reserved.

Copyright (c) 2013-2016 Intel, Inc. All rights reserved.

Copyright (c) 2011-2017 NVIDIA Corporation. All rights reserved.

Copyright (c) 2016 Broadcom Limited. All rights reserved.

Copyright (c) 2011-2017 Fujitsu Limited. All rights reserved.

Copyright (c) 2014-2015 Hewlett-Packard Development Company, LP. All rights reserved.

Copyright (c) 2013-2017 Research Organization for Information Science (RIST). All rights reserved.

Copyright (c) 2017-2018 Amazon.com, Inc. or its affiliates. All Rights reserved.

Copyright (c) 2018 DataDirect Networks. All rights reserved.

Copyright (c) 2018-2019 Triad National Security, LLC. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer listed in this license in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

The copyright holders provide no reassurances that the source code provided does not infringe any patent, copyright, or any other intellectual property rights of third parties. The copyright holders disclaim any liability to any recipient for claims brought against recipient by any third party for infringement of that parties intellectual property rights.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Opus

Copyright 2001-2011 Xiph.Org, Skype Limited, Octasic,

Jean-Marc Valin, Timothy B. Terriberry,

CSIRO, Gregory Maxwell, Mark Borgerding,

Erik de Castro Lopo

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Internet Society, IETF or IETF Trust, nor the names of specific contributors, may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Opus is subject to the royalty-free patent licenses which are specified at:

Xiph.Org Foundation:

https://datatracker.ietf.org/ipr/1524/

Microsoft Corporation:

https://datatracker.ietf.org/ipr/1914/

Broadcom Corporation:

https://datatracker.ietf.org/ipr/1526/

PIL

The Python Imaging Library (PIL) is

Copyright © 1997-2011 by Secret Labs AB

Copyright © 1995-2011 by Fredrik Lundh

Pillow is the friendly PIL fork. It is

Copyright © 2010-2023 by Alex Clark and contributors

Like PIL, Pillow is licensed under the open source HPND License:

By obtaining, using, and/or copying this software and/or its associated documentation, you agree that you have read, understood, and will comply with the following terms and conditions:

Permission to use, copy, modify, and distribute this software and its associated documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies, and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of Secret Labs AB or the author not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission.

SECRET LABS AB AND THE AUTHOR DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL SECRET LABS AB OR THE AUTHOR BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

portmidi

PortMidi Portable Real-Time MIDI Library

Copyright (c) 1999-2000 Ross Bencina and Phil Burk

Copyright (c) 2001-2009 Roger B. Dannenberg

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

posix pthread library

https://sourceforge.net/projects/pthreads4w

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

protocolbuffers/protobuf

Copyright 2008 Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Code generated by the Protocol Buffer compiler is owned by the owner of the input file used when generating it. This code is not standalone and requires a support library to be linked with it. This support library is itself covered by the above license.

pthreadpool

Copyright 2019 Google LLC

Copyright (c) 2017 Facebook Inc.

Copyright (c) 2015-2017 Georgia Institute of Technology

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Pytorch / Caffe2

From PyTorch:

Copyright (c) 2016- Facebook, Inc (Adam Paszke)

Copyright (c) 2014- Facebook, Inc (Soumith Chintala)

Copyright (c) 2011-2014 Idiap Research Institute (Ronan Collobert)

Copyright (c) 2012-2014 Deepmind Technologies (Koray Kavukcuoglu)

Copyright (c) 2011-2012 NEC Laboratories America (Koray Kavukcuoglu)

Copyright (c) 2011-2013 NYU (Clement Farabet)

Copyright (c) 2006-2010 NEC Laboratories America (Ronan Collobert, Leon Bottou, Iain Melvin, Jason Weston)

Copyright (c) 2006 Idiap Research Institute (Samy Bengio)

Copyright (c) 2001-2004 Idiap Research Institute (Ronan Collobert, Samy Bengio, Johnny Mariethoz)

From Caffe2:

Copyright (c) 2016-present, Facebook Inc. All rights reserved.

All contributions by Facebook:

Copyright (c) 2016 Facebook Inc.

All contributions by Google:

Copyright (c) 2015 Google Inc.

All rights reserved.

All contributions by Yangqing Jia:

Copyright (c) 2015 Yangqing Jia

All rights reserved.

All contributions from Caffe:

Copyright(c) 2013, 2014, 2015, the respective contributors

All rights reserved.

All other contributions:

Copyright(c) 2015, 2016 the respective contributors

All rights reserved.

Caffe2 uses a copyright model similar to Caffe: each contributor holds copyright over their contributions to Caffe2. The project versioning records all such contribution and copyright details. If a contributor wants to further mark their specific copyright on a particular contribution, they should indicate their copyright solely in the commit message of the change when it is committed. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the names of Facebook, Deepmind Technologies, NYU, NEC Laboratories America and IDIAP Research Institute nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

pytorch/cpuinfo

BSD 2-Clause "Simplified" License

https://github.com/pytorch/cpuinfo

Copyright © 2019 Google LLC

Copyright © 2017-2018 Facebook Inc.

Copyright © 2012-2017 Georgia Institute of Technology

Copyright © 2010-2012 Marat Dukhan

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

react-native

MIT License

https://github.com/facebook/react-native

Copyright © Facebook, Inc. and its affiliates.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

readerwriterqueue

Copyright (c) 2013-2015, Cameron Desrochers All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Resemblyzer

Copyright 2019 Resemblyzer Authors.

All rights reserved.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Roboto font family

Font data Copyright Google 2012

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at:

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

ruy

Copyright 2020 Google LLC. All Rights Reserved.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Skia

Copyright (c) 2011 Google Inc. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. Neither the name of Google Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

SQLite

SQLite is Public Domain

All of the code and documentation in SQLite has been dedicated to the public domain by the authors. All code authors, and representatives of the companies they work for, have signed affidavits dedicating their contributions to the public domain and originals of those signed affidavits are stored in a firesafe at the main offices of Hwaci. Anyone is free to copy, modify, publish, use, compile, sell, or distribute the original SQLite code, either in source code form or as a compiled binary, for any purpose, commercial or non-commercial, and by any means.

The previous paragraph applies to the deliverable code and documentation in SQLite - those parts of the SQLite library that you actually bundle and ship with a larger application. Some scripts used as part of the build process (for example the "configure" scripts generated by autoconf) might fall under other open-source licenses. Nothing from these build scripts ever reaches the final deliverable SQLite library, however, and so the licenses associated with those scripts should not be a factor in assessing your rights to copy and use the SQLite library.

All of the deliverable code in SQLite has been written from scratch. No code has been taken from other projects or from the open internet. Every line of code can be traced back to its original author, and all of those authors have public domain dedications on file. So the SQLite code base is clean and is uncontaminated with licensed code from other projects.

TagLib

Version 1.9.1

Copyright (c) 1995-2015 Scott Wheeler, Lukas Lalinsky, Ismael Orenstein, Allan Sandfeld Jensen, Teemu Tervo, Mathias Panzenböck, and Tsuda Kageyu.

Distributed under the Mozilla Public License v1.1 (MPLv1.1).

Full text of the license is available here: https://www.mozilla.org/en-US/MPL/1.1/

To receive a copy of the source code for the TagLib library distributed with this product under the terms of the MPLv1.1 please contact dev-support@native-instruments.com.

Tencent/rapidjson

https://github.com/Tencent/rapidjson

Tencent is pleased to support the open source community by making RapidJSON available.

Copyright (C) 2015 THL A29 Limited, a Tencent company, and Milo Yip. All rights reserved.

If you have downloaded a copy of the RapidJSON binary from Tencent, please note that the RapidJSON binary is licensed under the MIT License. If you have downloaded a copy of the RapidJSON source code from Tencent, please note that RapidJSON source code is licensed under the MIT License, except for the third-party components listed below which are subject to different license terms. Your integration of RapidJSON into your own projects may require compliance with the MIT License, as well as the other licenses applicable to the third-party components included within RapidJSON. To avoid the problematic JSON license in your own projects, it's sufficient to exclude the bin/jsonchecker/ directory, as it's the only code under the JSON license.

A copy of the MIT License is included in this file.

Other dependencies and licenses:

Open Source Software Licensed Under the BSD License:

The msinttypes r29

Copyright (c) 2006-2013 Alexander Chemeris

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS AND CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Open Source Software Licensed Under the JSON License:

json.org

Copyright (c) 2002 JSON.org

All Rights Reserved.

JSON_checker

Copyright (c) 2002 JSON.org

All Rights Reserved.

Terms of the JSON License:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

The Software shall be used for Good, not Evil.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Terms of the MIT License:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the “Software”), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Tensorflow

Copyright 2018 The TensorFlow Authors. All rights reserved.

Licensed under the Apache License, Version 2.0 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at

https://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an “AS IS” BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

TinyXML

Copyright (c) 2000-2002 Lee Thomason (www.grinninglizard.com)

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Tipue Search

Copyright (c) 2017 Tipue

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

vectorize

Copyright (c) 2012 Aaron Wishnick. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the <organization> nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL <COPYRIGHT HOLDER> BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

WebView2

Copyright (C) Microsoft Corporation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Microsoft Corporation, or the names of its contributors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

xsimd

Copyright (c) 2016, Johan Mabille, Sylvain Corlay, Wolf Vollprecht and Martin Renou

Copyright (c) 2016, QuantStack

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

- Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Yoga

Version 1.9.0

MIT License

Copyright (c) 2014-present, Facebook, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

ZeroMQ

Copyright (c) 2007-2015 iMatix Corporation

Copyright (c) 2009-2011 250bpm s.r.o.

Copyright (c) 2010-2011 Miru Limited

Copyright (c) 2011 VMware, Inc.

Copyright (c) 2012 Spotify AB

Copyright (c) 2013 Ericsson AB

Copyright (c) 2014 AppDynamics Inc.

Distributed under the GNU Lesser General Public License v3 (LGPLv3) with the following "Static linking exception":

The copyright holders give you permission to link this library with independent modules to produce an executable, regardless of the license terms of these independent modules, and to copy and distribute the resulting executable under terms of your choice, provided that you also meet, for each linked independent module, the terms and conditions of the license of that module. An independent module is a module which is not derived from or based on this library. If you modify this library, you must extend this exception to your version of the library.

Full text of the license is available here: https://www.gnu.org/licenses/lgpl-3.0.txt

To receive a copy of the source code for the ZeroMQ library distributed with this product under the under the terms of the LGPLv3 please contact dev-support@native-instruments.com.

zlib

Copyright (c) 1995-2004 Jean-loup Gailly and Mark Adler.

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly jloup@gzip.org

Mark Adler madler@alumni.caltech.edu