

Shotcut User Guide

Table of Contents

(June, 2026)

- Glossary / Terminology
- Installing Shotcut
- Launching Shotcut
- Quick Start Guide
- Project Management
- Adding Media
- Trimming Clips
- Three Point Editing
- Undo and Redo
- Trimming Filters
- A Guide for Blind Users
- How to Reduce Memory Usage
- Editing HDR Source Clips
- Editing Variable Frame Rate Source Clips
- New Project
- Keyboard Shortcuts
- File > New > Project and File > Close
- File > Save and Save As
- File > Backup and Save
- File > Open MLT XML as Clip
- File > Export > Frame
- File > Export > EDL
- View > Notes
- View > Filters
- View > History
- View > Jobs
- View > Layout
- View > Recent
- View > Resources
- View > Subtitles
- View > Application Log
- Other Items in the View Menu

Help > Actions and Shortcuts

Adjusting Panels in the User Interface

Settings > Video Mode

Settings > Processing Mode

Settings > Audio Channels

Settings > Preview Scaling

Settings > Proxy Editing

Settings > Player > Scrub Audio

Settings > Player > Use JACK Audio

Settings > Player > Realtime (frame dropping)

Settings > Player > Progressive

Settings > Player > Deinterlacer

Settings > Player > Interpolation

Settings > Player > Synchronization

Settings > Timeline > Rectangle Selection

Settings > Timeline > Automatically Add Tracks

Settings > Timeline > Scrolling

Settings > Display Method

Settings > App Data Directory

Settings > Job Priority Setting

Settings > Clear Recent on Exit

About the Playlist

Playlist > Slideshow Generator

Source vs. Project Player

Player > Loop

About the Timeline

Cut, Copy, and Paste Editing

Timeline > Output

Timeline > Ripple All Tracks

Timeline > Nudge

Timeline > Markers

Timeline > Record Audio / Voice-over

Timeline > Align To Reference Track

Timeline > Apply Copied Filters

Timeline > Adjust Clip Gain/Volume

Timeline > Selection and Groups

Timeline > Freeze Frame

Timeline > Replace

Creating and Removing Transitions

Changing the Style of Transition

Creating a Custom Video Wipe

Making Multi-track Transitions

Subtitles > Speech To Text

Subtitles > Text to Speech

Export Basics

Hardware Encoding

Creating Export Presets

Advanced Export

Embedding metadata into Exports

Export at the Command Line

Audio/Video Properties

Properties > Extract Sub-clip

Image Properties

Color Generator

Text Generator

Drawing/Animation Generator

Image/Video From HTML Generator

Lissajous Generator

Noise Generator

Plasma Generator

Color Bars Generator

Audio Tone Generator

Count Generator

Blip Flash Generator

Screen Snapshot Generator

Screen Recording Generator

Open Other > Network

Open Other > SDI/HDMI

Open Other > Audio/Video Device

Open Other > Video4Linux

Open Other > PulseAudio

Open Other > ALSA Audio

About Keyframes

Keyframe Types and Easing

Simple Keyframes

Audio Loudness Scope

Audio Peak Meter Scope

Audio Spectrum Scope

Audio Surround Scope

Audio Vector Scope

Audio Waveform Scope

Video Scopes

Video Histogram Scope

Video RGB Parade Scope

Video RGB Waveform Scope

Video Vector Scope

Video Waveform Scope

Video Zoom Scope

Add, Search, Favorite, Remove, and Reorder Filters

Copy and Paste Filters

Filter Sets

360: Cap Top & Bottom Video Filter

360: Equirectangular Mask Video Filter

360: Equirectangular to Rectilinear Video Filter

360: Equirectangular to Stereographic Video Filter

360: Equirectangular Wrap Video Filter

360: Hemispherical to Equirectangular Video Filter

360: Rectilinear to Equirectangular Video Filter

360: Stabilize Video Filter

360: Transform Video Filter

Alpha Channel: Adjust Video Filter

Alpha Channel: View Video Filter

Alpha Strobe Video Filter

Audio Dance Visualization

Audio Level Visualization Video Filter

Audio Light Visualization

Audio Spectrum Visualization

Audio Waveform Visualization

Blend Mode Video Filter

Blur: Box Video Filter

Blur: Exponential Video Filter

Blur: Gaussian Video Filter

Blur: Low Pass Video Filter

Blur: Pad Video Filter

Brightness Video Filter

Choppy Video Filter

Chroma Hold Video Filter

Chroma Key: Advanced Video Filter

Chroma Key: Simple Video Filter

Color Grading Video Filter

Contrast Video Filter

Corner Pin Video Filter

Crop: Circle Video Filter

Crop: Rectangle Video Filter

Crop: Source Video Filter

Diffusion Video Filter

Distort Video Filter

Dither Video Filter

Drop Shadow Video Filter
Elastic Scale Video Filter
Fade In Video Filter
Fade Out Video Filter
Fisheye Video Filter
Flip Video Filter
Glitch Video Filter
Glow Video Filter
GPS Graphic Video Filter
GPS Text Video Filter
Gradient Map Video Filter
Grid Video Filter
Halftone Video Filter
HSL Primaries Video Filter
HSL Range Video Filter
Invert Colors Video Filter
Key Spill: Advanced Video Filter
Key Spill: Simple Video Filter
Lens Correction Video Filter
Levels Video Filter
Lut (3D) Video Filter
Mask: Apply Video Filter
Mask: Chroma Key Video Filter
Mask: Draw (Glaxnimate) Video Filter
Mask: From File Video Filter
Mask: Simple Shape Video Filter
Mirror Video Filter
Mosaic Video Filter
Motion Tracker Video Filter
Nervous Video Filter
No Sync Video Filter
Noise: Fast Video Filter

Noise: Keyframes Video Filter
Old Film: Dust Video Filter
Old Film: Grain Video Filter
Old Film: Projector Video Filter
Old Film: Scratches Video Filter
Old Film: Technocolor Video Filter
Opacity Video Filter
Outline Video Filter
Posterize Video Filter
Reduce Noise: HQDN3D Video Filter
Reduce Noise: Wavelet Video Filter
Reduce Noise: Quantization Video Filter
Reduce Noise: Smart Blur Video Filter
Reflect Video Filter
Reframe Output Video Filter
RGB Shift Video Filter
Saturation Video Filter
Scan Lines Video Filter
Sepia Tone Video Filter
Sharpen Video Filter
Size, Position & Rotate Video Filter
Sketch Video Filter
Spot Remover Video Filter
Stabilize Video Filter
Subtitle Burn-In Video Filter
Text: Rich Video Filter
Text: Simple Video Filter
Text: Typewriter Video Filter
Threshold Video Filter
Timer Video Filter
Track Auto Fade Video Filter
Trails Video Filter

Unpremultiply Alpha Video Filter
Vertigo Video Filter
Vibrance Video Filter
Vignette Video Filter
Wave Video Filter
White Balance Video Filter
Ambisonic Decoder Audio Filter
Ambisonic Encoder Audio Filter
Band Pass Audio Filter
Balance Audio Filter
Copy Channel Audio Filter
Compressor Audio Filter
Delay Audio Filter
Downmix Audio Filter
Equalizer: 15-band Audio Filter
Equalizer: 3-Band (Bass & Treble) Audio Filter
Equalizer: Parametric Audio Filter
Expander Audio Filter
Fade In Audio Filter
Fade Out Audio Filter
Gain / Volume Audio Filter
High Pass Audio Filter
Invert Audio Filter
Limiter Audio Filter
Low Pass Audio Filter
Mid-Side Matrix Audio Filter
Mute Audio Filter
Noise Gate Audio Filter
Normalize: One Pass Audio Filter
Normalize: Two Pass Audio Filter
Notch Audio Filter
Pan Audio Filter

Pitch Audio Filter

Reverb Audio Filter

Stereo Enhancer Audio Filter

Swap Channels Audio Filter

Track Auto Fade Audio Filter

Track Seam Audio Filter

Declick Audio Filter

Speed: Forward & Reverse Time Filter

Speed: Forward Only Time Filter

Time Remap Filter

Preset Locations: Filters

Alpha channel

An alpha channel is a part of the image data that describes each pixel's level of opacity or translucency. Contrast that with the other data, which describes light and color. Many people think of an alpha channel as a transparency mask. However, that can lead one to think of it as rather binary - opaque or transparent - whereas most often an alpha channel has at least 256 levels.

Aspect ratio

Aspect ratio is how the width of a rectangle compares with its height. In other words, wide or tall, and by how much? This is typically expressed as a fraction of width over height. For example, if a rectangle is twice as wide as it is high, we can say it has a 2:1 aspect ratio. This can get complicated in video where pixels themselves can have an aspect ratio that is distinct from the display aspect ratio. For example, SD NTSC video has a resolution of 720x480 but can be either 4:3 or 16:9 display aspect ratio neither of which are square pixels! Modern video standards such as ATSC and UHD avoid using non-square pixels such that the common HD resolution 1920x1080 reduces mathematically to 16:9, which is also the most common video aspect ratio in use today.

B frames

A Bi-directional predicted (B) frame is a frame of a video that uses the motion-compensated differences between itself and both the preceding and following frames to specify its content. By doing so they use fewer bits to store the information than both I-frames and P-frames.

Bitrate

Bitrate refers to the quantity of data (or number of bits) per unit of time. Most digital audio and video formats use compression to save space, and the compression ratio will vary. For instance, a clear blue sky contains relatively little information and will generally compress well, while a more complex scene will compress less well. Bitrate is also affected by other factors such as the compression algorithm used and the quality level chosen.

Clip

In video editing, a **clip** refers to a segment of video footage or audio that is a distinct unit within the overall project. Clips can vary in length and content, ranging

from just a few frames to several minutes long. They are the building blocks of a video project and can be manipulated, arranged, and edited together to create the final product. Clips can be sourced from various recordings, imported media files, or generated within Shotcut itself in “**Open Other...**” (ex.: *Color* clip, *Text* clip, *Animation (Glaxnimate)* clip, etc...). Additionally, clips may contain both video and audio components, or they can be exclusively video or audio depending on the nature of the project. Clips are primarily used in the **Timeline** for editing but can also be utilized in the **Source** panel and in the **Playlist** for previewing and organization.

Codec

A codec comprises two components, an encoder and a decoder, hence the name. Examples of video codecs are H.264, H.265, VP9 etc.. Codecs use various technologies to compress data. The compression can either be lossless, in which case decoding the data will produce exactly the same data that was encoded, or lossy, in which case decoding the data will lose some of the data that was encoded. The higher the compression the more data is lost. In general the use of lossless codecs result in much larger files than lossy ones. Some codecs are more efficient than others in the amount of data they need to produce videos of equivalent quality e.g. H.265 produces smaller files than H.264 at equivalent quality; however, the more complex methods needed to do this usually mean that they take longer to encode and decode the video.

See also [Prologue | Codec Wiki](#)

Colorspace

Color space is about how color and light is represented especially numerically. In video and computer images, the two most popular systems of organization are RGB and YUV (or Y'CbCr). This is a complicated subject area; you can [read more on Wikipedia](#).

Deinterlacer


A deinterlacer is an algorithm to convert interlaced video to progressive scan. See below for definitions on these two terms.

Field order

Interlace video consists of two fields per frame. This term describes which field appears before the other in storage and/or display.

GOP

A GOP, or Group Of Pictures, specifies the order in which intra-frames (I-frames) and inter-frames (B- and P-frames) are arranged. The GOP is a collection of successive frames within a coded video stream. Each coded video stream consists of successive GOPs, from which the visible frames are generated. Encountering a new GOP in a compressed video stream means that the decoder doesn't need any previous frames in order to decode the following ones, and allows fast seeking through the video. The GOP structure is often referred by two numbers, for example, M=3, N=12. The first number tells the distance between two anchor frames (I or P). The second one tells the distance between two full images (I-frames): it is the GOP size. For the example M=3, N=12, the GOP structure is IBBPBBPBBPBBI.



GUI

GUI is short for Graphical User Interface. As opposed to a Command Line Interface (CLI), which enables users to interface with application only by typing commands, a GUI consists of various widgets (graphical elements such as buttons, scrollbars, color-pallettes etc.) that enable the user to control the application and receive feedback to enable them to make decisions on how next to proceed.

I frame

An Intra-coded (I) frame, also called a keyframe (not to be confused with keyframes used for animating filter parameters), is a frame of a video that is coded independently of all other frames. Consequently they use the more bits than B-frames and P-frames to store the information. Each GOP begins (in decoding order) with this type of frame.

Interlace

Interlace is a simple form of video compression that uses two half vertical resolution frames to represent a full frame. Basically, you can double the refresh rate for the same data rate. Each half vertical resolution image is called a field. Typically the fields are interleaved in storage and then displayed one after the other on play back by skipping every other line.

Interpolation

Interpolation is the computation of values based on neighboring values. With respect to **Settings**, it is easiest to think of this as the quality level when changing the

size of an image. Interpolation is also a term used for animating parameters in **Keyframes**.

Keyframe

A keyframe defines a specific value or set of values at a specific point in time. The term is used when talking about animating parameter values in **Keyframes**. It is also used in temporal video compression (so-called delta or P- or B-frames).

Metadata

Metadata is data about another data. In the context of multimedia, the media data (audio/video) is the core data, and all other data in the file is metadata. There can be metadata about the media attributes such as resolution or number of audio channels. And there can be metadata about the context of the media file such as its creator, creation data, title, etc.

MLT

MLT is another open source software project that is the engine of Shotcut. Shotcut is primarily the user interface running on top of this engine. This engine provides some effects of its own, but it also uses other libraries such as FFmpeg, Qt, WebKit, frei0r, lads, etc.

P frame

A Predicted (P) frame is a frame of a video that uses the motion-compensated differences between itself and the preceding frame to specify its content. By doing so they use fewer bits to store the information than I-frames, but more than B-frames.

Progressive

In video, this refers to a scan mode where (“scan” refers to old tube-based TV technology where a cathode ray draws video by drawing lines) each frame of video is a whole picture from a single point in time. This is the opposite of interlace.

Resolution

Display Resolution is the actual pixel size of video and images. Example sizes: 3840x2160 (4k), 1920x1080 (HD), 1280x720 (HD), 720 × 480 (SD), 1080x1380 (Vertical)

Ripple

Ripple means that an operation can affect the clips on the timeline that are later or after the clip being changed. For example, a ripple delete not only removes the clip but also the space it occupies. This requires changing the start time of all of the following clips.

Sample rate

Sample rate refers to how many times per second a sample is taken of the audio or video. For instance, in the real analog world, a car driving down the street moves continuously, but a video of that consists of individual snapshots taken at regular intervals. If you could slow down time and watch the car, the real world car would still move smoothly; however, the video would show the car jumping from one position to another. Common sample rates for video include 24, 25, 30, and 60 frames per second; common sample rates for audio include 44,100 and 48,000 samples per second.

Scrubbing


Scrubbing is seeking by clicking some object and dragging it. Typically this is the play head in the player or timeline. But it can also refer to simply rewind and fast forward playback through media.

Skimming

Skimming is seeking based on the horizontal position of the mouse over the video image or timeline. You press and hold the **Shift** + **Alt** keys in Shotcut to enable skimming.

Snapping

Snapping is a setting on the timeline. Aids in sliding two clips together on the same track with no gap in between them.



Scan mode

Scan mode indicates whether video is progressive or interlaced - see related definitions above.

Threads

Threads are a software programming mechanism to let multiple things occur at the same time. Most CPUs now consist of multiple execution units typically called “cores.” Often, these cores support a CPU-based “thread,” which you can think of as a light core (not completely parallel). While it is important that your operating system let multiple things run at the same time to use these CPU cores and threads, it also important that Shotcut run things (i.e. parallel processing) at the same time because media decompression, processing, and compression is very computationally heavy. You can learn more about how Shotcut uses multiple cores and threads in the [FAQ](#).

Timecode

Timecode is a way to represent time numerically. Shotcut uses a standard called SMPTE from the Society of Motion Picture and Television Engineers. It is a display of the running time of video that is frame-accurate yet easier for humans to understand than pure frame count. It shows hours, minutes, seconds, and frames in the format HH:MM:SS:FF. On many video modes that use a non-integer frame rate (e.g. 29.970030 or 30000/1001 fps), the semicolon (;) is the delimiter between seconds and frames to indicate that it is using drop-frame timecode. Drop-frame is a technique to make the timecode follow the real time over long durations. For 30000/1001 fps, drop-frame subtracts two frames every minute except every tenth minute.

UI

UI is short for User Interface. This is the mechanism whereby users interact with the application (Shotcut). The goal of this interaction is to allow effective operation and control of the application from the human end, while the application concurrently feeds back information that aids the user’s decision-making process.

VUI

VUI is short for Visual User Interface, In the context of Shotcut this is an interface that appears in the preview area of the screen when certain filters, like the Text: Rich filter, or the Size, Position & Rotate filter are active. The VUI enables users to manipulate the parameters by dragging handles and/or typing directly in the preset window rather than changing the parameters manually in the filters panel.

XML

XML is a text format that is designed to be both human and machine readable and writable. It is standardize, structured, and extensible - the X in eXtensible Markup Language. There are many dialects of XML, and when one video editor says it reads “XML” it does not mean it can read the XML that another video editor can export. They need to be the same *kind* of XML. Shotcut reads and writes MLT XML, but at this time it can only fully understand the MLT XML that it writes.



Shotcut supports Windows, macOS, and Linux. Here's how you can install it on each platform:

Windows

1. Go to the [Shotcut download page](#).
2. Choose the appropriate installer for your system (e.g., Windows 10/11 on Intel/AMD or ARM CPU)
3. Download the installer and run it.
4. Follow the on-screen instructions to complete the installation.

Alternatively, you can purchase it for a small price in the Microsoft Store. It is single lifetime purchase, and doing so helps fund the development. In return, you get automated updates without needing to re-download and install.

macOS

1. Go to the [Shotcut download page](#).
2. Select the macOS version compatible with your system (e.g., macOS 12+ or older versions).
3. Download the app bundle and open it.
4. Drag the Shotcut app to your Applications folder.

Linux

1. Go to the [Shotcut download page](#).
2. Download the Linux portable tar or ApplImage file.
3. For the portable tar, extract the archive and run the `shotcut` shell script.
4. For the ApplImage, make it executable and run it directly.

Alternatively, you can install the Flatpak (via Flathub), the Snap, or a package from your Linux distribution. Often, the software manager in your desktop environment makes some of these options available to you.

Please be aware that the portable, ApplImage, and Snap are all the same build based on the glibc version in Ubuntu 20.04 and thus stuck to using an older version of Qt. The Flatpak or distribution package works better on the **Wayland** graphics subsystem. Also, the Flatpak or distribution package supports **ARM** processors whereas the portable, ApplImage, and Snap do not.

Windows

The installer puts the app icon in the Start menu only (not on the desktop). Use the Start menu or search for it. When using the portable version, you **must** extract the zip file. You should not open the zip file as a virtual folder in Explorer and run it from there. In Explorer, navigate to where you extracted the Shotcut portable zip file, locate `shotcut.exe` and double-click it. If you want to create a shortcut on your desktop, right click `shotcut.exe` and choose **Send to > Desktop** from the context menu.

macOS

It depends on where you drag the app out of the downloaded DMG file, but usually it is **Applications** in **Finder**. If you did add it to Applications, then you can also find it in Launchpad and search it in Spotlight. After it starts, right-click the app icon in the Dock, and choose **Keep in Dock** if you want to launch it from there in the future.

If you want to run it from the command line run `open -a shotcut` or manually run the executable at `/Applications/Shotcut.app/Contents/MacOS/shotcut` per the app nimble standard. To run a version not installed to `Applications` at the command line, use `open /path/to/Shotcut.app`. To run another instance of Shotcut at the same time, use `open -na shotcut` OR `open -n /path/to/Shotcut.app`.

Linux

It depends very much on which app package you used, the desktop environment you use, and sometimes the distribution.

If you are using the **portable tar** version, first extract the tar file. Then, in your file manager, double-click the icon in Shotcut. That does not always work. In some desktop environments, it will only work on the Desktop. In others, it does not work at all. In GNOME and extracted to the Desktop, you might need to right-click the app icon and choose **Allow Launching**. If the file manager does not launch Shotcut, open the `Shotcut.app` folder, and try double-clicking the `shotcut` shell script. **Do not try to run `bin/shotcut` directly.**

If you are using the **AppImage**, simply **make it executable**. Then, double-click the file to run it.

If you are using the **Snap**, it depends on your distribution. On Ubuntu, after installing, Shotcut will be in your menu or list of programs and in your `$PATH` to simply run `shotcut` at the command line.

If you are using the **Flatpak**, after installing, it should be in your menu or list of programs. At the command line run `flatpak run org.shotcut.Shotcut`. The upper case `s` is important.

Quick Start Guide on how to use Shotcut

New Project

Option	Description
Projects Folder	This is where Shotcut will look for your file for editing by default, and where it will save videos when you export them.
Project Name	This creates a sub-folder within the named Projects folder, and creates the project save file when you hit start.
Video Mode	Important to set this before beginning your project. Leaving this in Automatic may not produce the results you desire. Example: If you want to Export a HD 1080p 30 FPS video, set it here first.

Project File Saves

Function	Description	Shortcut Win/Linux	Shortcut macOS
Save	Saves to a project MLT file. If already saved previously, updates opened MLT project file.	Ctrl + S	command + S
Save As	Allows you to choose a new file name (and/or location) for your current project. Open project will take on the new file name and/or file location.	Ctrl + Shift + S	shift + command + S
Backup & Save	Performs Save with the ability to make a backup copy of <code>named.mlt</code> to <code>named yyyy-mm-ddThh-mm-ss.mlt</code> while keeping <code>named.mlt</code> active. Must perform function twice to get date/time saved file.	Ctrl + Alt + S	option + command + S

Editing

There are three modes of editing available within Shotcut.

Modes	Description
Timeline	The most popular. Has multiple Video/Audio tracks (or one). Needed for transitions and voiceovers. No need for the Playlist.
Playlist	Arrange clips in order, apply filters if needed, and then export. No need for the Timeline.


Modes	Description
Source	Just a single file. Trim if needed and/or apply filters as needed, then export. No need to use the Playlist or Timeline.

Exporting

Options	Description
Presets	Choose a preset that best suits your export needs. If you are a beginner Stock → Default is your best choice.
From	You can export from Timeline, Playlist, Source, or Marker. If you have no clips in the playlist or the timeline, or no markers set, these options will not appear.
Hardware Encoder	Graphics card encoder may aid with exporting (Read FAQ). Click Configure, then Detect to set. If nothing is detected then your graphics card is not compatible.
Export File	Prompts you to name your video export file. Once named, click Save and export will start. Do not close Shotcut until the export finishes.
Reset	Resets all options back to Default.
Advanced	This is where you will find all of the advanced options such as Parallel Processing, Codec, Rate control, CRF, Bitrate, etc. Warning: <i>The Advanced moded does not prevent creating an invalid combination of options!</i>

If you're a beginner at video editing follow these steps:

To export with the Default settings (frame rate will match your Video Mode)

1. Click on the Export icon at the top, right.
2. 
3. Name your video file.
4. Wait for Shotcut to finish exporting.

After the export is done, you can double click on the job to play your finished video within Shotcut.

The longer the video project, amount of filters, and computer specifications will all be factors in how fast your video export takes to complete.

Troubleshooting

At times exporting may fail. Video filters may not perform the way you expect them to. Or some other issue that you may not understand. Video editing is a very complicated task for anyone and their computer. Do not be afraid to ask for help.

- First start by reading the [FAQ](#) (Frequently Answered Questions). Several issues are addressed in there.

- Search the forum for the issue you're having. If you can't find the answer, feel free to make an account and post your issue. The forum is made up of many volunteers who also use Shotcut.
- Look through the [Tutorials](#) section of the forum.
- Watch [Tutorial Videos](#) specifically chosen by the developers.

Introduction

Shotcut does not provide very much in the way of project management at this time. Rather, you are required to manage many aspects of this. There are a few things you should know:

- Shotcut saves to a project file in a text format called **MLT XML**. (Technical detail on that format can be found [here](#).)
- The project file does not contain any of your media. Rather, it only *links* to your media files by file system path and name.
- There are 3 different kinds of projects:

1. Clip

If you put anything in **Playlist** or **Timeline** and you open a media file (or generator), then you can save it as a clip-only project along with trimming information as well as filters and their keyframes.

2. Playlist

If you add some things to the **Playlist** but not the **Timeline**, then you have saved a playlist-only project.


3. Timeline

This is what most people use, but keep in mind it is not necessary to use the timeline. Any of the above can be exported as a final audio/video file!

Project Folder

Shotcut supports the concept of a project folder, which is optional. A project folder is when you create a file system folder for your project and you put project file in there along with whatever companion files get created as well as media files if you choose.

When Shotcut first starts (or after you choose **File > New** or **File > Close**), this view appears in the area for the **Source** player:



You use this to create a project folder. When you use it, Shotcut creates a folder on the file system with the name in **Project name** within **Projects folder**. Then, it saves an empty project with the name in **Project name** followed by `.mlt`. Thereafter, companion files that Shotcut creates are saved automatically into the project folder instead of requiring you to name them:

- **Stabilize** video filter
- **Text: HTML** video filter
- **Properties > Reverse**
- **Text to Speech**
- **Record Audio**
- **Screen Snapshot**
- **Screen Recording** (depending on operating system, see [Screen Recording Generator](#))
- **Image/Video from HTML**
- **New > Animation**
- **Mask: Draw** video filter
- **Settings > Proxy > Storage > Use Project Folder**

We encourage you to store all or most of the media files that you use in the project in the project folder. However, Shotcut does not provide a way to automatically copy or move the files into there automatically at this time. Use your operating system's file manager for that. If you move the files after they were added to the project, the next time you open the project it will show a missing files dialog where you can re-link them to the project folder.

You can also store multiple `.mlt` (MLT XML) files in your project folder. You could be nestinge/embedding child MLT XML files within the parent project file, or you could use this to have multiple timelines in a project (each timeline is its own MLT XML file). If you have Settings Backup turned on, the backup files are saved in this project file too.

Use the **Files** panel (menu item **View > Files**)  folder path drop-down and choose **Current Project** to view its contents. Or use from the menu **File > Show Project in Folder** to open your operating system's file manager to the project folder.

Of course, you can skip using this because it is entirely optional, and the view goes away as soon as you open something. Also, you can create a project folder manually by not using this, but Shotcut then does not provide the automatic naming and placement of the companion files.

Relative vs. Absolute File Names

Shotcut saves MLT XML with file paths and names rather than embed any media. So, it can save these paths in full (i.e. absolute) or partially (i.e. relative). On Windows, a full path begins with a drive letter such as `c:`. On Linux and macOS, a

full path begins with a slash (/). Shotcut saves with a relative path if the file is in the same folder as the `.mlt` project file or a sub-folder of this folder. It saves with an absolute path otherwise. This means your project file can be a mix of the two; it is determined per file.

There are different schools of thought around which to use. If you use absolute paths, you can keep media where you have it already organized and freely move the project document without any impact. If you use relative paths, you can put everything into a single folder (optionally with sub-folders), and freely move around the project folder. Take your pick; Shotcut does not force you into either one.

Example : Relative File Structure

- D:/Shotcut/ZProject (File Folder)
- D:/Shotcut/ZProject/Zproject.mlt (Current project MLT file)
- D:/Shotcut/ZProject/BugV.png (Source Location)
- D:/Shotcut/ZProject/Square/BugS.png (Source Location)

When saved, the MLT file has the Source file locations as such:

- BugV.png
- Square/BugS.png

You can copy the file folder D:/Shotcut/Zproject which is helpful for sharing with people or to another computer. Also useful for backing up saved projects for later use.

Example: Absolute File Structure

- D:/Shotcut/**XProject** (File Folder)
- D:/Shotcut/**XProject**/Xproject.mlt (Current project MLT file)
- D:/Shotcut/ZProject/BugV.png (Source location)
- D:/Shotcut/ZProject/Square/BugS.png (Source location)

When saved, the MLT file has the Source file locations as such:


- D:/Shotcut/ZProject/Bug Vertical.png
- D:/Shotcut/ZProject/Square/Bug Square.png

The folder of D:/Shotcut would need to be copied for sharing with people or to another computer. If you tried to share XProject.mlt along with your source files individually to another person they would be presented with this dialog box.

Autosave

Autosave is always working and cannot be turned on or off. It does not save to your existing project. Rather, it saves to a hidden file in your app data directory. Then, it is checked when you reopen a project. It is only meant for crash recovery. Every 60

seconds it checks if the project is modified (the window title bar shows an asterisk except on macOS which uses the red dot) and saves to the backup file. When you reopen a project after a crash, if the backup is found, Shotcut prompts you to use it:




Once you save successfully - including save at exit when prompted, - the backup file is removed. This also works for projects that were never saved - an Untitled project. Except in that case, simply restarting Shotcut will attempt to locate a backup and prompt. Otherwise, for a named project, you need to open the project for it to check.


Missing Files

Shotcut presents a missing dialog box when you open the project if a file has been moved or deleted or file folder has changed name. This allows you to still use your existing project, but you must tell Shotcut where to find your source file.


Double click each item that's missing as assign the source files a new location.



One files have been found, then click OK.





You will now have two project files. The original, and a Repaired project MLT file.



There are many ways to import media files in Shotcut.

Open File

- From the main toolbar or from the **File** menu, click on **Open File**. Or use the keyboard shortcut **Ctrl+O** (**command+O** on macOS).



- Navigate to the folder containing the media file (or files) you want to import.
- Select your file(s) and click on the **Open** button.

If only **one** file is opened, it will automatically be added to the **Source panel**. From there it can be dragged to the **Playlist**, or to the **Timeline**.


If **multiple** files are opened, they will all be added to the Playlist. The first file in the list will open in the Source panel.

NOTE: **Open File** cannot be used to import a *MLT XML file as clip*.

From the Playlist

The Playlist panel can be used to import files in a project, including *MLT XML file as clip*.

- At the bottom of the Playlist panel, click on the **Add files to playlist** button.



- Navigate to the folder containing the file (or files) you want to import.
- Select your file(s) and click on the **Open** button.

The file(s) will be added to the Playlist. From there, any or all of them can be moved to the Timeline when needed.

Drag from a folder

You can add files (including *MLT XML file as clip*) to your project by dragging them from a folder

- Open the folder containing the file(s) you want to import.
- Select the file(s).
- Grab and drag the file(s) in the Playlist or in the Timeline..
- You can also grab and drag the file(s) in the Source panel.

If you drag **multiple** files at once in the Source panel, they will all be added to the Playlist, and the first file in the list will open in the Source panel.

In video editing, **trimming** is the process of adjusting the start (in-point) or end (out-point) of a video or audio clip to shorten or lengthen its duration. This is different from moving the clip on the timeline; it's about changing the content and duration of the clip itself.


While some software uses a dedicated trim tool, in Shotcut you perform trimming directly within the **Timeline** or **Source** player using controls and keyboard shortcuts.

A good workflow involves using the **Source** player to prepare clips *before* adding them to the timeline. This helps you select and trim the best parts of your footage first, which saves time and makes the editing process on the timeline more efficient and less disruptive.

Source Player

You can open a clip into the **Source** player a number of ways:

- **File > Open** (menu, main toolbar, shortcut)
- drag from a file manager & drop onto the player
- double-click in:
 - **Recent**
 - **Playlist**
 - **Files**



Drag  to change the in-point.

Drag  to change the out-point.

Alternatively, play and seek within the clip to position the playhead: 

Then, press to set the in point and/or (letter o, not the number 0) to set the out point.

As you are trimming the bottom, **right** corner of the player shows some time values. Remember, there are tool tips through the application. Hold the mouse a couple of seconds over one of the time values to reveal its purpose: In Point / Selected Duration. Note that, technically, duration is not the out point. If you actually move the playhead to the out-point, the time field in the bottom, **left** corner shows the out point. However, more often the duration of this sub-clip is more interesting than the out point.

Sub-Clip and Playlist

A **sub-clip** is a new, shorter clip created from the larger, original media file. Instead of keeping the entire clip, you create a sub-clip that contains only the specific, trimmed segment you want to use. This is a common practice in the **Source** player, where you set the in-point and out-point of a desired section and then add it to the **Playlist**. Sub-clips are useful for organization, as they allow you to work with manageable, focused pieces of a long recording without affecting the original footage. You can have multiple sub-clips from the same source in the playlist or timeline. You can also add **Filters** to the sub-clips while preparing it in the **Source** player.

Timeline

Clips on the timeline also have interactive controls for trimming. Move the mouse cursor to the left edge of a clip, and you see a green bar appear:



Drag it to trim the in point.

Move the mouse cursor to the right edge of a clip, and a red bar appears:



Drag it to trim the out point.

You can also use the and keys to on the timeline.

Regardless of the approach, in the timeline the behavior also depends on something called **Ripple**, which is a fundamental concept in video editing and especially in Shotcut. For the interactive mouse control, ripple is used when this timeline toolbar button is active:



When the button is *not* active, you can hold while trimming to do a ripple trim.

There are ripple variants of the keyboard shortcuts:

- Shift + I for **Ripple Trim Clip In**
- Shift + 0 for **Ripple Trim Clip Out**

Another thing to note about using the timeline keyboard shortcuts is that you do not need to first select the clip. It operates on the topmost clip under the playhead. But an important consideration is that shortcuts on the timeline cannot *increase* the duration. It can only decrease by removing more from the beginning or end.

Roll Edit

A **roll edit** is a specific type of trim that adjusts the cut point between two adjacent clips on the timeline simultaneously. The defining feature of a roll edit is that it shortens the end of the first clip and lengthens the beginning of the second clip by the exact same amount. This means the **total duration** of the two clips—and your entire video sequence—remains unchanged.

You use a roll edit to refine the timing and rhythm of a transition without affecting the rest of your timeline. It allows you to shift the “moment of the cut” to a more effective point, such as aligning a visual cut with a beat in music or a specific word in dialogue.

In Shotcut, hold Ctrl when trimming to perform a roll edit.

Scrub While Dragging

“Scrubbing” is essentially another word for seeking, but it is a bit more about playback than simply going to a point in time. Think of it as rapidly seeking short distances while updating the video preview. This can be useful when trimming or doing a roll edit. In Shotcut, this is on when the timeline toolbar button is active:




Otherwise, when it is off, you can still kind of see where you are inside the clip by viewing the thumbnail on the timeline.

Three-point editing is a technique used in video editing to manipulate footage by using three reference points: an in-point, an out-point, and a cut point. It's particularly useful when you need to extract a segment from a source clip and insert it in your **Timeline**. The segment can also be inserted into to the **Playlist**.

The Three Points

In Point


This is the starting time of the segment you want to insert into the **Timeline**. In the **Source** player, you set the in-point at the specific frame where you want the clip to begin.



The keyboard shortcut used to set the in-point is **I**.

Out Point

This is the ending time of the segment you want to insert. In the **Source** player, you set the out-point at the specific frame where you want the clip to end.




The keyboard shortcut used to set the out-point is **O** (the letter o, not the number 0).

Cut Point

This is the location in the **Timeline** where you want to insert or overwrite the new clip. To set the cut point, move the playhead at the frame where you want the new clip to be inserted. If there is more than one track in the **Timeline**, also make sure to change the current track to where the segment needs to go.

Edit

The last step is to execute the edit operation: **Paste** (insert) or **Overwrite**. Drag-n-drop is not recommended but still possible. With drag-n-drop the playhead is only a visual cue and not as precise.







The Undo and Redo functions in Shotcut are crucial for managing your editing workflow efficiently.

- **Undo:**  The Undo feature allows you to revert the last action you performed. Whether you added a clip, adjusted the filters, or made a cut, you can easily reverse the change by clicking the **Undo** button (or using the keyboard shortcut Ctrl + Z). This can be repeatedly used to step back through multiple actions.
- **Redo:**  The Redo function is the counterpart to Undo. If you have undone an action and decide you want to keep it, the **Redo** button (or the keyboard shortcut Ctrl+ Y on Windows or Ctrl + Shift + Z on Linux and macOS) will reapply the most recently undone action. Like Undo, you can use it repeatedly to step forward through your action history.
- **History Panel:** Shotcut also features a [History Panel](#), which provides a visual timeline of your editing actions. This panel allows you to see and navigate through your actions quickly, making it easier to manage complex edits. Access the History Panel by going to **View > History**.
- There is a default maximum of 50 items for the history, but it can be increased using the [configuration](#) key `undoLimit`.

Trimming a filter means you can choose when the filter effect starts and ends independent of the clip's starting and ending time. Obviously, if the filter is applied to a clip, it must still be within the clip's time frame. If the filter is applied to a track, however, it can be truly independent. Filter trimming takes place in the **Keyframes** panel:




The first row of the **Keyframes** panel shows the clip with its thumbnails and audio waveform. Drag the left, green edge to change when the filter starts. Drag the right, red edge to change when it ends. Here is how it looks when both ends of the filter have been trimmed:



The dimmed portion of the clip reflects when the filter is inactive.

With version 19.10.20, toolbar buttons and keyboard shortcuts for filter trimming were added.


These are the toolbar buttons:



These are the keyboard shortcuts:

- [to set the filter start
-] to set the filter end

Using these buttons and keyboard shortcuts will speed up the process of precise filter trimming considerably. Simply place the playhead at the desired point and press the filter start or end button/keyboard shortcut.





The timeline is the least accessible part of the Shotcut UI. You do not need to use the timeline in Shotcut. It is possible to make single-clip-only and playlist-only projects in Shotcut. I think we should start there. In a clip-only project, you can open a single clip, trim and filter it, and then export it. For this, you only need Ctrl+O to open a file into the Source player. Then, you can refer to all of the keyboard shortcuts under the player section. When you use left and right arrow keys here you should hear little audio blips as long as “Scrub Audio” is turned on in the Settings menu. If you still do not hear the audio, then this feature is incompatible with the format or encoding of some files. In the Properties panel there is a way to convert a file into what I call an edit-friendly format to fix that. Next, Press I to set the start frame, and press O to set the ending frame. Finally, you can export the sub-clip.

Now, you can also do the above repeatedly; but, instead of export, add each sub-clip to the playlist. When done, you can export the playlist, which will concatenate all of the sub-clips.

Initially, the playlist panel is not open, and that is OK. With a clip opened and trimmed, you can press shift+a to add it. That opens the playlist panel for you, but maybe you do not need to do anything else at this time except to export it of course. You can add more to the playlist by repeating everything you have done thus far. You can play the entire playlist by pressing the escape key and use the same playback shortcuts. Press escape key again to return the clip player, which has the label “Source”. Exporting the playlist is simple. As soon as you add something to the playlist, export automatically defaults to export from the playlist. This can be confirmed using the drop-down control labelled “From” in the export panel.

Let me give you a tip about the toggles for the different panels. When you use the View menu or keyboard shortcuts, it will toggle the visibility of the panel. However, the main toolbar buttons always make it visible - not a toggle - and raises the panel in case it is tabbed behind another panel. I just noticed that neither of these methods give focus to the panel. Panel focus is difficult in this UI and for most users. So, Shotcut demphasizes it. However, sometimes it is needed. I guess your voiceover software lets you search the UI and take focus. Each panel does have like a sub-window title, and maybe you can find it that way. This panel title is also reflected in the tab in case the panel is tabbed with another panel.

Here are ways to reduce the memory usage in general:

- Save and restart Shotcut
- Close other browser tabs.
- Close other applications.
- Reboot (does all of the above).

Here are ways to reduce the memory usage of your project:

- Reduce the resolution of the [Video Mode](#).
- Reduce the number of tracks. The number of tracks will increase the amount of caching to prevent files from being closed and reopened repeatedly from one frame to the next.

Here are some ways to reduce the memory usage while editing:

- Make sure Settings > Realtime is turned ON. When it is off, it uses more threads, and more threads needs more memory.
- Use [Preview Scaling](#) and [Proxy Editing](#).

Here are some ways to reduce the memory usage while exporting:

- See all of the things list above about reducing memory usage in general and of the project.
- Turn OFF **Export > Advanced > Video > Parallel processing**. That uses more threads, which needs more memory.
- Turn ON **Export > Use hardware encoder** (use **Configure... > Detect** if needed).
- It can be very helpful to reboot, start Shotcut and only Shotcut after reboot, open the project but do not play through it, and then export. You are then guaranteed to give the most of your available physical memory to the Shotcut export process.

Introduction


Some users report that their clip looks “washed out” or “pale” in Shotcut compared to media players. This is often because the clip is encoded as High Dynamic Range (HDR). Shotcut has very limited support for editing HDR, and most users do not want to bother with HDR output. So, HDR clips can be converted to SDR that Shotcut can edit.

Identifying HDR Clips

To identify HDR clips in Shotcut, open the clip and view the Properties panel. In the properties panel, look for the “Color transfer” field in the Video tab.

The following Color transfer types are HDR:

- SMPTE ST2084 (PQ)
- SMPTE ST428
- ARIB B67 (HLG)





Transcode HDR to SDR

When opening HDR clips in Shotcut, you may be prompted to convert the file. If you are not prompted, follow these steps:

1. Click **Convert...** in the **Properties** panel. A dialog window opens.
2. Open the “Advanced” section of the dialog.
3. Ensure that the **Convert to BT.709 colorspace** checkbox is checked.
4. Click **OK** to start the conversion.

After the conversion is complete, the clip will be replaced with the converted clip.



Convert Color With a LUT

Instead of converting a HLG HDR clip, you can use it as-is with or without a proxy and use a LUT to convert it to SDR. Our [Resources](#) page has a LUTs section with some freely downloadable LUTs from NBC/Universal that can convert from HLG HDR to BT.709 SDR. “SL” in the file name means scene light, and “DL” means display light. Use either of these with the **LUT (3D)** video filter.

Convert Using Another Tool

There are different ways to convert HDR to SDR since the result is only an approximation of the source. Shotcut's Convert gives a result some people do not like. A lot of people mention using the popular Free, open source, cross-platform Handbrake video transcoder.

If you are using macOS and shooting HDR video with an iPhone you can use QuickTime Player:

1. Choose **File > Export As**
2. Pick a resolution and a file dialog opens
3. At the bottom of the dialog is a drop-down where you can choose:
Greater Compatibility (H.264)

How to Edit HLG HDR Video With Shotcut

Version 25.05 of Shotcut added limited support for editing, processing, and exporting HDR video. This is focused on HLG HDR at this time because my cameras are iPhone and GoPro HERO13 that both can shoot in HLG. It is also possible to record in log color, but that is not the focus here.

Limitations

- No embedded HDR preview! You can only see the HDR video as HDR in Shotcut when using a Blackmagic Design SDI/HDMI device with **Settings > Player > External Monitor**. As an upside for this, since HLG is supposed to provide a satisfactory experience when viewed on a SDR screen, this gives you a way to simultaneously preview both SDR and HDR.
- If you want to do any filtering, transitions, or track blending/compositing you must enable **Settings > Processing Mode > Linear 10-bit GPU/CPU** and limit yourself to video filters with the tags `#gpu` or `#10bit` on 10-bit video sources. (GPU filters always support a 10-bit workflow but not all CPU video filters.)
- When not using GPU processing, you can do a simple cuts-only editing (trimming and sequencing). It is OK to include text on a solid color or still image with this, but its color will look a little different than the preview.
- There is a bug in version 25.05 converting color range of 10-bit video between full and limited. You should export the same range as the source. This has been

fixed for the next version.

A Simple, Inexpensive Setup for HLG HDR

This also includes how I do 5.1 surround sound.

- recording with GoPro MAX and HERO13 cameras – less than \$500 each,
- previewing on a 42" 4K HDR (800 nits) Samsung TV I bought in 2018 for about \$650,
- using a Blackmagic Design [DeckLink Mini Monitor 4K](#) that I bought this year for \$200, and
- playing through a Sony AV receiver I bought in 2018 for less than \$200.

Shotcut Settings

- Turn on **Settings > GPU Effects**.
- Choose **Settings > Audio Channels > 6 (5.1 surround)** if desired.
- Create and use a custom **Video Mode** with **Colorspace = ITU-R BT.2020**.
- Ignore conversion dialogs when importing video.
- Add and edit media as usual. For video filters on HLG video sources, use only `#gpu` and `#10bit` filters. You can use these tags when searching filters.
- To convert first order (4 channel) Ambisonic audio to 5.1 surround, use the **Ambisonic Decoder** audio filter.
- It improves performance to use proxies and preview scaling, but the proxy for HLG must have been created with version 25.05 or later.
- In **Export**, select a preset from the “ten_bit” category. I have tested AV1, HEVC, and ProRes.
- If the Shotcut version is 25.05 and the source HLG video is primarily full color range (for example, GoPro HERO13 is but not iPhone), set **Export > Video > Color range = Full (JPEG)**.
- In **Export > Other** add the line `color_trc=arib-std-b67`.




Introduction

Sometimes users report that Shotcut is not frame accurate. For example, they split a clip on a certain frame, but then later the split is not where they expected it. Or, they put a keyframe on a certain frame, but later the keyframe seems to be on a different frame. Also, sometimes people report that the audio and video out not in sync.

Symptoms like these can be caused by Variable Frame Rate (VFR) source clips. Some users report that VFR clips work in Shotcut if the VFR is subtle. Other clips, however, are completely unusable in Shotcut.

Identifying VFR Clips

Some programs will skip encoding frames as a way to reduce file size. Also, some devices will skip encoding frames if they do not have enough processing power to perform real time encoding. Typically, these clips can be identified by looking at the frame rate in the properties panel. If the clip frame rate does not match a standard frame rate, then it is probably VFR. If Shotcut is able to detect a VFR clip, it will add a “(variable)” tag after the frame rate. But Shotcut can not always detect VFR clips.




Converting VFR Clips

The best way to edit VFR clips is to convert them to a fixed frame rate. Follow these steps:

1. Set the desired **Settings > Video Mode** in Shotcut (do not use Automatic). In particular, choose a frame rate for your project that meets your needs.
2. Open the VFR clip
3. If Shotcut detects the VFR, it may prompt you to convert. If not, click the “Convert...” button in the properties panel to open the conversion prompt.
4. Click “OK” to start the conversion

After the conversion is complete, the clip will be replaced with the converted clip.

The **New Project** view is the primary screen shown when Shotcut is launched. It serves as an optional starting point for work, ensuring that your project is organized and that the technical foundation (resolution and frame rate) is correctly established before media is added.



To start a new project, you must complete the three primary fields in the center of the application window. If you are already working in a project, you can return to this screen via **File > New** or `Ctrl+N` (`Cmd+N` on macOS).

1. Projects Folder

This is the parent directory on your computer where Shotcut will save your project folders.

- **How to set:** Click the text or the folder icon to browse your system.
- **Best Practice:** Choose a dedicated folder for video editing so all your projects are kept in one location.

2. Project Name

Enter a descriptive name for your current work (e.g., “Summer_Vacation_2024”).

- **Behavior:** When you click **Start**, Shotcut creates a sub-folder with this name inside your **Projects folder** and creates an empty `.mlt` project file within it.
- **Storage:** You can optionally to store all the media (video, audio, and images) you plan to use inside this specific project folder to avoid “missing file” errors later.
- **Generated files:** Shotcut can generate files for certain tasks such as Stabilization or generating an image from HTML. When using this feature to make a project folder, Shotcut automatically saves them to this folder instead of asking you every time where to put it! This **only** works if you press the **Start** button, simply entering a name here does nothing in and of itself.

3. Video Mode

The **Video Mode** is the most critical setting. it defines the resolution (e.g., 1080p), the aspect ratio (e.g., 16:9), and the frame rate (e.g., 30 fps) of your project.

- **Automatic (Default):** Shotcut will leave the settings blank and automatically set the Video Mode based on the properties of the **first file** you drag onto the Timeline or Playlist. Not the first file at any later point in time as you make changes. Literally, the first time you add something, the video mode is set and locked in until you change it with **Settings > Video Mode**.
- **Manual Selection:** You can click the drop-down to choose from various presets (HD, 4K, Vertical/Social Media).
- **Important:** It is best to set this manually if you know your final delivery requirements. Changing the Video Mode mid-project can occasionally cause filters or transitions to misalign.
- This does not change **Settings > Video Mode** if you start a project *without* clicking the **Start** button!!

Starting the Project

Once these fields are filled, click the **Start** button.

- The Timeline, Playlist, and Player will be initialized according to your Video Mode.
- The window title bar will now display your project name followed by an asterisk (*) if there are unsaved changes.

Recent Projects List

On the right side of the **New Project** view, Shotcut maintains a list of recent **Projects**.

- **Quick Access:** Simply click a project name in this list to immediately reload that work.
- **Search:** Use the search bar at the top of the **Recent** panel to find older projects by name.
- **Removal:** To remove a project from the list, right-click it and choose **Remove**. Removal only removes the project from this list; it does not remove any files. Use your operating system's or preferred file manager, such as Explorer on Windows or Finder on macOS, to manage files.

See also [Project Management](#), [Settings > Video Mode](#), [View > Recent](#)



See [Keyboard Shortcuts page](#) on the main web site.

As of version 22.09 you can edit or add your own using [Help > Actions and Shortcuts](#).



Both of these commands do the same thing

1. If you have un-saved work you are first asked if you want to save your project.
2. If a file is open—whether a project, saved or unsaved, or a single piece of media—the current project is closed and all clips and timeline tracks are removed from the window. It will look similar to how it appears when you start the program by itself (not by launching it along with a project or media file).
3. **Close** does not close the Shotcut program; use **Exit** for that.
4. If you already did some things in Shotcut you need to choose **Close** or **New** in order to return to the **New Project** view.

See also [Project Management](#)




These actions save your project to a file format that is specific to Shotcut so you can resume editing at a later time. If you want to output a video (or audio only) file that plays in any other media player or to share use [Export](#).

Save



- Saves current project to a `.mlt` file you name.
- Repeat using **Save** continues to save current project to the same specified file.
- Shortcut `Ctrl` + `S` (macOS `Cmd` + `S`)

Save As



- Saves current project to a `.mlt` file you name (possibly different).
- Repeat using **Save As** will offer you a new file name to save as each time.
- The last **Save As** file you name will be the current project file open in Shotcut. Now when using **Save** it only saves to this current project file name.
- Shortcut `Ctrl` + `Shift` + `S` (macOS `Shift` + `Cmd` + `S`)


Backup and Save

This is only available if you have already saved a project (“Untitled” does not appear in the window title bar). It first gets the last date and time the project file was modified. Then, it duplicates the current project file with a new name that ends with the date and time of last modification. Finally, it saves the changes.


This menu item was added in version 23.09.

Making backup copies of a project can be very helpful in case we want to go back to a known good state of the project.

A project backup can be created, either manually with **File > Backup and Save** or automatically with **Settings > Backup**.



When a backup occurs, the project backup appears in **File > Other Versions**.




In Shotcut, the “**Open MLT XML as Clip**” feature allows you to import an MLT XML file directly into your project as a single clip. This means that instead of treating the XML file as a project file with multiple elements like video clips, audio tracks, and effects, Shotcut treats it as a single media item that can be placed on the timeline like any other clip.

Using this feature can be handy if you want to incorporate a pre-edited project into your current project without having to import each individual element separately. It simplifies the process by treating the entire XML file as a single entity, making it easier to manage and manipulate within your project.

Like any other type of clip, a MLT clip can be trimmed and have filters applied to it.

MLT clips work best when the profile of the clip matches the profile of the project it is being added to. To make this easier, the properties panel for an MLT clip will display the profile properties and a validation if the clip matches the project profile:




MLT clips with a profile that does not match the project profile can still be used. But it may result in lower quality or performance.

To import a “MLT XML” file as a clip you can either:

1. Use the “File” menu

- Open the **File** menu and click on “**Open MLT XML as clip...**”




- Navigate to the folder containing the MLT file you want to import.
- Select and open the file.

The MLT file will open in the the Source panel, and from there it can be moved to the Playlist, or to the Timeline.

2. Use the Playlist panel

- **Open the Playlist panel.**
- Click on the “**Add files to playlist**” button



- Navigate to the folder containing the MLT file you want to import.
 - Select and open the file.
- The MLT file will be added to the Playlist, and from there you can move it to the Timeline.

3. Drag it from a folder

- On your computer, open the folder containing the MLT file you want to import.
- Select the file and drag it in Shotcut, either to the Source panel, the Playlist, or the Timeline.


Export Frame is a feature that lets you save the frame currently displayed in the Shotcut viewer as a still image. The available formats are: *PNG*, *BMP*, *JPEG*, *PPM*, *TIFF*, or *WebP*.

Export Frame was added in Shotcut version **16.07**

After moving the playhead to the desired frame, you can export the image by either:


- Using the **Ctrl** + **Shift** + **E** keyboard shortcut.
- Going to **File > Export** and choosing **Frame...**

In the export dialog box, Shotcut automatically generates a file name corresponding to the position of the playhead, but you can use a custom name if you prefer. This is also where you can choose the file format.



This saves the frame as it is displayed in the viewer, including any changes made with filters or from options in the Settings menu (e.g., Proxy).

Upon saving, the new file is added to the **Recent** panel for easy access in Shotcut.



This feature makes it easy to achieve a freeze frame effect:

- Move the Playhead to the desired frame on the timeline.
- Export the frame.
- In the **Recent** panel, double-click on it to copy it in the **Source** viewer.
- In the Timeline toolbar, click on the **Paste** button to insert the still image at the position of the Playhead.

It's also convenient for creating a custom thumbnail for your web video:

- Export the frame you want to use as a thumbnail.
- In the Recent panel, double-click it to open it in the **Source** viewer.
- Apply a Text filter to add a title.
- Use **Export > Frame** again to export the new thumbnail.



Exports the current timeline as an **EDL (Edit Decision List)** file.

An EDL is a **text-based interchange format** that describes *editing decisions*, not media content.

What an EDL is

An EDL records how a timeline is edited, including:

- Clip order
- Source in/out points
- Timeline in/out points
- Basic transitions (usually cuts only)
- Track assignments (limited)

It does **not** contain video, audio, effects, or rendered output.

An EDL answers the question:

“How the edit is put together?”

not

“What does the final video look or sound like?”

Purpose of exporting an EDL

EDL export is primarily intended for interoperability and offline workflows, such as:

- Moving an edit from Shotcut to another NLE
- Conforming an offline edit to higher-quality media
- Archiving edit decisions separately from media
- Educational or documentary workflows

It is not meant for final delivery.

What is included in Shotcut’s EDL export

Typically included:

- Clip references
- Source timecodes (in/out)
- Timeline placement
- Simple cuts

Typically **not** included:

- Filters or effects

- Transitions beyond simple cuts
- Speed changes
- Keyframes
- Color correction
- Audio filters
- Generators
- Titles
- Compositing or track blending
- Most multi-track relationships

EDL support is intentionally minimal.

Format characteristics and limitations

- Plain text format
- Line-based, rigid structure
- Originally designed for linear tape editing
- Limited metadata capacity
- No extensibility for modern effects

Note:

The EDL format is limited to a single video track and a small number of audio tracks (commonly up to four).

Because of this, EDLs are lossy representations of modern timelines.

Compatibility with other software

EDL is an old but widely recognized format.

Commonly supported (non-exhaustive list):

- Natron
- Blender (VSE)
- Pitivi
- KDenlive
- DaVinci Resolve
- Avid Media Composer
- Adobe Premiere Pro (import with constraints)
- Final Cut Pro (via conversion tools)
- Various conforming and finishing systems

Important:

Each application interprets EDLs slightly differently. Results may vary.

Requirements and expectations

To successfully use an exported EDL:

- Source media must be available to the target application
- File names and timecodes must match
- Reel names (if used) must be compatible
- Expect manual cleanup after import

EDL workflows assume editor intervention, not one-click transfer.

What this export is *not*

- Not a project backup
- Not a render
- Not a universal interchange format
- Not suitable for complex timelines

Recommended use cases

- Rough cut transfer
- Offline-to-online conform workflows
- Archival documentation of edit decisions
- Educational demonstrations of editing structure
- Legacy broadcast workflows

Practical usage note

If your timeline relies heavily on filters, effects, compositing, or generators, exporting to EDL will produce a **simplified and incomplete** representation of the edit.


This is expected behavior, not a bug.

Keyboard Shortcut Ctrl + Shift + 3 (shift + command + 3 on macOS).

The Notes icon on the top toolbar will open/close the panel.

Menu: View → Notes

- Use the new Notes panel to keep notes about your project.
- This panel can moved to other panel sections.



The **Filters Panel** allows you to add and manage effects and adjustments to your video and audio clips. You can apply a variety of filters to enhance your media, such as color correction, size and position adjustments, masks, blurs, audio adjustments and many more. Most filters come with customizable settings.

This panel also lets you:

- Manage the order of applied filters.
- Disable/Enable applied filters.
- Copy and Paste filters.
- Create and save [Filter Sets](#).

Display/Hide the panel

If the Filters panel is not visible in the interface, you can either:

- Click the **Filters** icon on the main toolbar 
- Go to the top menu and select **View > Filters**
(can also be used to hide the panel)

If the panel is visible, but all buttons are grayed out, either:

- Select a **Clip**, a **Track head**, or the **Output track** in the Timeline.
- **Copy** or **Cut** any clip from the Timeline and switch to the **Source** viewer.
- Double-click any item in the **Playlist** to copy it to the **Source** viewer.
- Double-click an image, video or audio file in the **Recent** panel to copy it to the **Source** viewer.

To view a list of all available filters, see the Documentation [Table of Content](#)



The History Panel in Shotcut allows you to manage and navigate through your editing actions efficiently. Here's how to make the most of it:

- **Accessing the History Panel:** To open the History Panel, click the History icon  on the main toolbar or go to the top menu and select **View > History**. This will display a list of all your recent editing actions in chronological order.
- **Using the History Panel:** The History Panel lists each action you've taken in your current editing session, such as adding clips, applying filters, or making cuts. By clicking on any item in the list, you can revert to that specific point in your editing timeline. This can be particularly useful for complex edits where you need to quickly backtrack and make adjustments.
- **Undo and Redo Integration:** The actions in the History Panel are directly tied to Shotcut's Undo and Redo functions. If you undo an action from the History Panel, it will appear in the panel, allowing you to easily redo it if necessary.
- **Managing History:** The panel not only helps you track changes but also allows you to delete specific actions from the history if you wish to simplify your action list.

By incorporating the History Panel into your workflow, you can streamline your editing process and make adjustments with confidence and precision.

The **Jobs** panel is Shotcut's background task manager. It displays the status and progress of intensive processes—such as exporting video, creating proxies, or converting files—allowing you to continue editing or queue multiple tasks without waiting for each one to finish individually.

Opening the Jobs Panel

If the Jobs panel is not currently visible, you can display it using these methods:

- **Toolbar:** Click the **Jobs** icon on the top main toolbar.



- **Keyboard Shortcut:** Press `Ctrl + 0` (Windows/Linux) or `Cmd + 0` (macOS).
- **Menu:** Go to **View > Jobs**.

The Jobs Interface

The Jobs panel typically appears on the right side of the interface. It contains a list of all tasks initiated during the current session.

Job Status Indicators

- **Pending (Empty Bar):** The job is in the queue and will start automatically when previous jobs finish.
- **Running (Percentage/Progress Bar):** The job is currently processing. You can see the elapsed time and an estimate of the remaining time.
- **Finished (Checkmark):** The task completed successfully.
- **Failed (Red X):** The task encountered an error. You can right-click these to view a log for troubleshooting.

Right-Click Menu Options

Right-clicking any job in the list provides several management and diagnostic tools:

Option	Description
Stop	Cancels a running job.
Remove	Removes the entry from the Jobs list (does not delete the actual file).
View Log	Opens a text window showing the technical output of the process. Essential for diagnosing “Failed” exports.
View XML	Displays the MLT XML code that was sent to the background “Melt” engine to process the job.
Open	Opens the resulting file in the Shotcut Source Player (useful for checking a finished export).

Option	Description
Show in Folder	Opens your system's file explorer and highlights the finished file.
Repeat	Restarts a finished or failed job using the original settings.

Common Job Types

- **Export:** Triggered when you click “Export File” in the Export panel.
- **Proxy:** Generated automatically if “Use Proxy” is enabled in **Settings**.
- **Convert to Edit-Friendly:** Created when you choose to optimize a variable frame rate or highly compressed clip.
- **Reverse:** Created when you use the “Reverse” tool in the Properties panel.

Queue Management

- **Sequential Processing:** By default, Shotcut processes one job at a time to maximize system resources. As one job finishes, the next in the list begins automatically.
- **Pausing:** You can right-click in the empty space of the Jobs panel to **Pause** the entire queue. This is helpful if you need to free up CPU power for a complex edit before resuming the export later.
- **Priority:** You can change the background processing priority in **Settings > Job Priority**. Setting this to “Low” ensures that your computer remains responsive for other apps while Shotcut exports in the background. This only affects newly enqueued jobs, not existing running or pending jobs.

Troubleshooting Failed Jobs


If a job fails:

1. Right-click the job and select **View Log**.
2. Scroll to the bottom to look for error messages (often related to “Codecs,” “Disk space,” or “Memory”).
3. Ensure your destination drive has enough free space.
4. If using **Hardware Encoder**, try disabling it and repeating the job to see if the issue is driver-related.

Shotcut includes several built-in layout presets designed to optimize the interface for specific stages of the video production process. These presets automatically open the most relevant panels and arrange them into logical groups, saving you the time of manually opening and docking individual views.

Accessing Layouts

You can switch between layouts at any time by going to **View > Layout**. The current layout is also indicated by the specialized buttons often found at the top right of the Shotcut window.



or



Standard Layout Presets

Logging

- **Focus:** Organizing and reviewing raw footage.
- **Key Panels:** Large **Playlist** and **Properties** panels.
- **Usage:** Best used at the start of a project to inspect clips, set in/out points, and categorize media before moving them to the timeline.

Editing (Default)

- **Focus:** General purpose post-production.
- **Key Panels:** **Timeline**, **Filters**, and **Recent** files.
- **Usage:** The balanced, standard view for most trimming, splitting, and arrangement tasks.

FX (Effects)

- **Focus:** Heavy filter work and keyframing.
- **Key Panels:** Large **Filters** panel and the **Keyframes** timeline.
- **Usage:** Provides the detailed view necessary for precise timing of animated effects and complex filter stacks.

Color

- **Focus:** Color correction and grading.

- **Key Panels:** **Video Scopes** (Waveform, Parade, Vectorscope) and the **Filters** panel.
- **Usage:** Specifically arranged to give you maximum visual feedback while using color grading filters.

Audio

- **Focus:** Sound mixing and leveling.
- **Key Panels:** **Audio Scopes** (Peak Meter, Loudness, Spectrum Analyzer) and the **Timeline**.
- **Usage:** Optimized for adjusting track volume, applying audio filters, and ensuring your project hits target loudness levels.

Player

- **Focus:** Minimalist screening.
- **Key Panels:** Maximized **Project Player**; hides almost all other panels.
- **Usage:** Used for reviewing your final edit without the distraction of technical tools.

Managing Custom Layouts

Beyond the presets, Shotcut allows you to save your own favorite workspace configurations.

- **Add...:** After arranging your panels exactly how you like them, select **View > Layout > Add...** to give your configuration a name. It will now appear in your Layout menu.
- **Remove...:** To delete a custom layout you no longer need, select it from the menu and choose **Remove**. (Note: You cannot remove the built-in system presets).
- **Restore Default Layout:** If your interface becomes cluttered or a panel goes missing, use this option to reset the current preset to its original “factory” state.

Pro Tip: Layouts are “Snapshots”

When you select a layout, Shotcut applies it as a snapshot. If you move a panel after selecting “Color,” Shotcut remembers that change for the remainder of your session. If you want to return to the original “Color” look, simply select the “Color” layout again from the menu.

See also: [Adjusting Panels in the User Interface](#)



The **Recent** panel provides a convenient history of the media files and project files (`.mlt`) you have recently opened in Shotcut. This panel allows you to quickly revisit previous work or re-use assets without having to navigate through your system's file explorer.

Opening the Recent Panel

If the Recent panel is not visible in your current layout, you can open it using the following methods:

- **Toolbar:** Click the **Recent** icon on the top main toolbar.



- **Keyboard Shortcut:** Press `Ctrl + 3` (Windows/Linux) or `Cmd + 3` (macOS).
- **Menu:** Go to **View > Recent**.

Features and Interface

1. The Recent List

The panel displays a vertical list of file paths. By default, these are sorted by the time they were last accessed, with the most recent items at the top. The list includes:

- **Project Files:** Shotcut project files (`.mlt`).
- **Media Files:** Video, audio, and image files that were opened in the **Source** player or added to a project.

2. Opening Items

- **Double-Click:** Double-clicking any item in the list will immediately open that file in the **Source** player. If the item is a project file, Shotcut will prompt you to save your current work before loading the selected project.
- **Drag and Drop:** You can drag items directly from the Recent panel onto the **Source** player or **Playlist**.

3. Searching the List

At the top of the panel is a search bar that allows you to filter the list.

- **Search (F3):** Press `F3` while the panel is focused to jump to the search box. Typing a partial filename or folder path will hide all items that do not match your query, making it easy to find specific files in a long history.

Context Menu (Right-Click)

Right-clicking an item in the Recent panel provides the following management option:


- **Remove:** Removes the selected entry from the Recent list. Note that this **does not delete the file** from your computer; it only clears the shortcut from this list.

Settings

You can manage how Shotcut handles your recent file history through the main settings menu:

- **Settings > Clear Recent on Exit:** When this option is enabled, the Recent list will be completely wiped every time you close the application. This is useful for users on shared computers or those who prefer a clean slate for every session.

This was added in version 23.09 to the **View** menu. It shows all unique and sorted files in the **Playlist** and **Timeline**



This also appears when you drag and drop multiple files to **Source**, **Playlist**, or **Timeline** that might be problematic - as a way to bulk convert them.

Clicking **Convert** shows the same **Convert to Edit-friendly** dialog that you can reach from the **Properties** panel.



If you see an exclamation point icon hover the mouse over the icon to see a tooltip that explains a possible issue.

The Subtitles Panel was added in Shotcut 24.08

The Subtitles Panel can be used to add subtitles to a project. Subtitles can only be added to the Timeline. They can not be added to the playlist or clips. We recommend to save subtitle editing until the timeline editing is complete because moving or cutting clips on the timeline will break the subtitle synchronization.

Subtitle Tracks


To start editing subtitles, click the “add” button and give the subtitle track and name and language.



After the subtitle track is created the subtitle track area will appear in the timeline above the top video track. Multiple subtitle tracks can be added. But only the currently selected subtitle track is shown in the Subtitle Panel and the Timeline.


Adding Subtitles

Click the “Add subtitle” button to make a new subtitle item at the cursor position. Then type in the “Current” text area to add the text.




Importing Subtitles

Subtitles can also be imported. Supported subtitle import formats include: SRT, VTT, ASS, and SSA. Subtitles are imported at the current cursor position. So be sure to place the cursor before importing subtitles from a file.



Exporting Subtitles


Subtitles can be exported to a file. Subtitles are exported in SRT format.



Embedding Subtitles

When exporting a Shotcut project, subtitles will be embedded in the output file if the file format supports it. Formats that commonly support subtitles include MKV, MOV and MP4.

From the **View** menu, choose **Application Log...** to view the log of application messages. The application log can be useful to reveal potential problems or to give developers hints about the application performance.



The dialog shows a window with a scrollable view of the application log.

When requesting support on the forum, it can be helpful to click **Copy** to copy the contents to the system clipboard and then paste the log into a forum post.

Alternately, click **Save** and save the log to a file which can be uploaded to the support forum.

Note: The application log is cleared every time Shotcut starts up. When providing an application log to provide information about a problem, be sure to provide a log from the same session of Shotcut when the problem occurred. Click **Previous** to view the log of the previous session. This is useful if you have to restart or the application crashed.

While the primary panels (Timeline, Filters, Export) are accessed frequently, the **View** menu contains several utility items that toggle specific interface elements or provide diagnostic information about the application.

Interface Toggles

Full Screen

- **Shortcut:** F11
- **Action:** Expands the Shotcut window to cover the entire screen, hiding the operating system's taskbar or dock.

Show Title Bars

- **Action:** Toggles the thin bars at the top of each panel (e.g., "Playlist," "Properties").
- **Usage:** Disabling title bars creates a cleaner, more minimalist look. However, you must enable them if you wish to drag, undock, or close panels using the mouse.

Show Toolbar

- **Action:** Toggles the main horizontal toolbar at the top of the window.
- **Usage:** Hiding the toolbar can provide more vertical space for the Player or Timeline, especially on smaller screens. Note that most toolbar functions are also available via keyboard shortcuts or the main menus.

Show Small Icons

- **Action:** Reduces the physical size of the icons on the main top toolbar.
- **Usage:** If you are working on a laptop or a low-resolution screen, enabling this option frees up valuable vertical space for the **Player** and **Timeline** panels. It is also a popular choice for experienced users who recognize the icons by sight and no longer need large targets to click.

Show Text Under Icons

- **Action:** Toggles the text labels (e.g., "Open File," "Save," "Export") beneath the toolbar icons.
- **Enabled:** Recommended for new users to help learn the interface quickly without hovering the mouse to see tooltips.
- **Disabled:** Recommended for advanced users or those with limited screen space. Turning this off significantly thins the toolbar, providing a more "professional" and minimalist aesthetic.

Utility Panels

History

- **Shortcut:** `Ctrl + 8` (Windows/Linux) / `Cmd + 8` (macOS)
- **Description:** Opens a list of every action taken in the current session (e.g., “Split clip,” “Change filter”).
- **Usage:** You can click on any previous state in the list to “undo” multiple steps at once. The history is cleared every time you close Shotcut.

Notes

- **Shortcut:** `Ctrl + Shift + 3` / `Cmd + Shift + 3`
- **Description:** Provides a simple text area to type reminders, project to-do lists, text to be converted to generated-audio speech, or script notes for voiceovers.
- **Usage:** These notes are saved directly into your `.mlt` project file.

Resources

- **Description:** Displays a list of all unique media files currently used in the project (both in the Playlist and on the Timeline).
- **Usage:** Useful for identifying “missing” files or bulk-converting problematic variable-frame-rate files to “Edit-Friendly” formats.


Diagnostic Tools

Application Log

- **Description:** Opens a text window showing the background technical activity of Shotcut.
- **Usage:** If Shotcut crashes or behaves unexpectedly, developers on the [Shotcut Forum](#) will often ask for this log to help diagnose the issue. You can use the **Save** button to export the log to a text file.
- Use the **Previous** button to view the log of the previous session. For example, if Shotcut crashed, restart, choose View > Application Log, and click **Preview** to view the log of the session that crashed.

This feature provides a unified action search and shortcut editor. Action search makes it quick to find and execute a command. It also provides a keyboard shortcut reference and lets you add or change shortcuts.

The Actions and Shortcuts Window was added in version 22.09.23.



- The window can be accessed by going to Help > Actions and Shortcuts.
- The window can also be accessed by the default shortcut keys ? and /.
- Searches both action name and shortcut.
- Press Return / Enter within the **search** box to move focus to the list of actions.
- Press Return / Enter on a selected **Action** to trigger it and close the dialog.
- Shift, Ctrl, or Alt + Return / Enter on a selected action to trigger it and NOT close the dialog.
- Double-click an action to trigger it but NOT close the dialog.
- The dialog is not modal so you can leave it open and to the side while working.
- Press Esc when the dialog has focus to close it.
- Press F2 when an action is selected to edit its first shortcut.
- Single-click a selected shortcut to enter edit mode using the mouse. This means you typically need two single clicks: one to select, the second to enter into edit mode.
- You can Tab out of edit mode to navigate focus to the apply button and press Space to trigger the apply. It is intentional that you cannot assign Tab as a keyboard shortcut to an action.
- There are some keyboard buttons such as J/K/L that are reserved and cannot be used in a shortcut. You will see an error message when you try to apply it.

- You can define up to 2 shortcuts per action! You may want to leave the default as well as add your preferred shortcut.



Shotcut features a highly flexible, dockable user interface based on the Qt framework. This allows you to rearrange, resize, stack, or even float panels to create a workspace that fits your specific hardware (such as dual monitors) and editing style.

Key UI Components

The interface is composed of several functional panels including the **Playlist**, **Filters**, **Properties**, **Timeline**, and various **Scopes**. You can toggle these on or off using the **View** menu. The main toolbar buttons either open the panel or if it is in a background tab, brings that tab to the front.

Resizing Panels

You can change the amount of screen real estate dedicated to any panel by dragging its edges. Hover your mouse over the border between two panels. The cursor will change to a double-headed arrow. Click and drag to resize both panels simultaneously.

Moving and Reordering Panels

Each panel (except the central player) has a title bar. You can use this to move the panel to a new location.

- **To Move:** Click and hold the title bar of a panel, then drag it toward another area of the screen.
- **Drop Zones:** As you drag, you will see blue “drop zones” appear.
 - **Edge Zones:** Dropping here will snap the panel to the side, top, or bottom of an existing panel.
 - **Center Zone:** Dropping in the center of an existing panel will create a **Tabbed Group**.

Tabbed Groups

To save space, you can stack multiple panels in the same area.

- **Accessing Tabs:** When panels are grouped, tabs appear at the bottom or top of the panel area. Click a tab (e.g., switching between “Playlist” and “Filters”) to bring that panel to the front.
- **Reordering Tabs:** You can click and drag the tabs themselves to change their left-to-right order.

Floating and Undocking Panels

If you use multiple monitors, you may want to move certain panels (like Scopes or the Timeline) to a second screen.

- **Undock:** Click the “undock” icon (two overlapping boxes) in the panel’s title bar, or simply drag the panel completely outside of the main Shotcut window.
- **Redock:** Drag the floating window back over the main Shotcut interface until a blue drop zone appears, then release.

Hiding and Showing Title Bars

For a cleaner look, you can hide the title bars of all panels.

- **Menu:** Go to **View > Show Title Bars** to toggle them.
- **Note:** You must have title bars visible to move or close panels using the mouse.

Closing Panels

If your workspace becomes cluttered, you can close any panel you aren’t using.

- **Method 1:** Click the **X** in the corner of the panel’s title bar.
- **Method 2:** Toggle the panel off in the **View** menu or the main toolbar.
- **Method 3:** Use the dedicated keyboard shortcut (e.g., `ctrl+6` for Filters).

Restoring the Layout

If your interface becomes disorganized or a panel “disappears,” you can quickly return to a known state.


- **Restore Default:** Go to **View > Layout > Restore Default Layout**.
- **Presets:** Shotcut includes several specialized layout presets such as **Editing**, **Audio**, **Color**, and **FX**. These can be found under the **View > Layout** menu.

Introduction

Video Mode is the project **resolution**, **aspect ratio**, frame rate, and **color space**. It is important to set this as you desire at the beginning of your project. If you try to change it later Shotcut will try to retain the timing of edits, but some users have reported problems, and it is risky. Also, at this time Shotcut will not automatically adjust anything for you related to size and positioning of elements like text and picture-in-picture when changing resolution. You can think of this like the page size in a word processor or the canvas size in a drawing program.

Broadcast vs. Non-Broadcast

Shotcut includes a number of video modes that are compatible with broadcast standards and thus the **Blackmagic Design SDI/HDMI** integration if you have that. These are also simply the most popular choices as well:




How to read these:

- HD = high definition (resolution)
- SD = standard (low) definition
- UHD = ultra-high definition (UHD 2160 is the 4K variant of UHD)
- the first number is the vertical resolution (number of lines or rows)
- the second number is the number of frames-per-second (fps)
- p after the first number means **progressive scan**


- i after the first number means [interlaced](#)

If you go to the **Non-Broadcast** sub-menu, there are a number of useful alternatives including trendy square and vertical resolutions:



Custom

In the bottom **Custom** sub-menu, you can **Add...** your own:



If you give your mode a name, it will be saved in the **Custom** sub-menu. Otherwise, you can defined something for the current project only.


The maximum resolution is 8192 x 8192. Only even numbers are accepted because most video codecs only support even values.

Click the down arrow at the end of **Frames/sec** to pick from some popular frame rates. Due to some weird legacy video stuff, most often something that is called 30 or 60 fps is actually slightly less. The exact numbers are achieved through fractions

3000/1001 and 60000/1001. In order to express those in Shotcut, it is important to use enough decimal digits. 29.97 or 59.94 is not considered precise enough because there are some devices and tools that use these exact values instead of the fraction values. Then, this imprecision affects Shotcut's Automatic video mode and export. Some people do not want their source video frame rate uncontrollably altered at export. Thus, Shotcut only really uses 30000/1001 when you use 29.970030 and 60000/1001 when you use 59.940060. These special values are conveniently available in this drop-down menu.

There are many colorspace standards for video, but at this time Shotcut only supports ITU-R BT.601, which is typical for SD, and ITU-R BT.709 for just about everywhere else except HDR.

Of course, after adding a custom video mode, you can use **Settings > Video Mode > Custom > Remove...** to remove it:




Simply click the checkbox next to the name or double-click the row until you see the check appear and click **OK**.

Automatic

Automatic means the resolution and frame rate will be based on the **first** file you **add** to your project. If the first file is not a video clip (for example, image or audio), then the video mode will be **1920x1080p 25 fps**. This is not necessarily the first clip currently in the **Playlist** or **Timeline**. Again, it is the first file you open and then add to the playlist or timeline. Therefore, it is almost always recommended to wisely choose a video clip first when you use **Automatic**. If you are making an image slide show, you may find the default 1080p25 adequate. If not, then set your video mode! If you are only doing audio editing, then it does not matter much except higher frame rate gives you greater precision in edits.

Project vs. Settings

Once you start a project, the current **Video Mode** is adopted (not **Automatic**) or determined (**Automatic**). Then, when you load a project, the **Video Mode** in the **Settings** will be checked if it matches your project. Otherwise, if a match is not found Automatic is checked. (Versions before 20.06 did NOT change its checkmark to indicate the project's mode.) You can also see the video attributes of a currently opened project - assuming it is using the **Timeline** - click **Output** (this was named "Master" in versions before 20.06) in the top left corner and view **Properties**.



You can change the project video properties by choosing a **Video Mode** in the **Settings** menu; However, as mentioned above, be careful with this because Shotcut will not automatically adjust anything for you related to size and positioning of elements.

This sub-menu chooses how Shotcut does its image processing after reading a video or image and before displaying or exporting it. It shows these options:


- **Native 8-bit CPU (fast, flexible)**
- **Native 10-bit CPU (slower, better)**
- **Linear 10-bit CPU (slowest, best)**
- **Linear 10-bit GPU/CPU (Experimental)**

This was added in version 25.12 and replaces the old menu item **Settings > GPU Effects**. Previous versions of GPU Effects did not process the image in linear color when using CPU filters; since version 25.12 it does.

10-bit Video

Bit depth is the number of colors or video brightness levels. 8-bit has only 256 values per color component, for example, red/green/blue. 10 bits has 1024 levels. Previously, 10-bit video sources were only partially supported using **GPU Effects** and only GPU filters. Without GPU Effects (CPU), in order to stay in the 10 bit space, you could only use filters with the `#10bit` tag but no transitions or track blending/compositing. Now, we have ported many CPU video effects to support 10- and 12-bit video sources including transitions and track blending. Also, the `#10bit` CPU video filters also work properly with the experimental GPU processing. Not all of the video CPU filters have been converted; many of them come through another software project that does not support it. However, we were able to convert many for this release or at least mark them as safe for a 10-bit project.

The filters menu **Video** button shows a drop-down to choose a suitable filter for either 10-bit and/or GPU processing pipelines:



Linear Color Processing

Any mixing of pixels with a color space's gamma is not ideal, can cause color distortions, and is technically incorrect. This pixel mixing occurs whenever images are interpolated, for example scaling, blurring, and sub-pixel positioning, as well as with track blending/composition. Previously, only **GPU Effects** and video filters in the **GPU** category operated in linear color. Now, there is an option to use it outside of the GPU processing.

There is no linear option for 8-bit because that naturally causes quantization artifacts. There are not enough bits available to show enough levels for the portions that the human eye is most sensitive.


Non-linear “native” options are still included because there is a significant performance decrease when using linear (more conversions). We have a plan to improve performance for the next release. 8-bit is still included because it is more “flexible” by including more video filters. Of course, you can use an 8-bit video filter in a 10-bit, but the quality of the video clip or track will suffer a reduction in bit depth.



This is a project setting that controls the number of audio channels you want to use in your project. The default is 2 channel stereo, which suffices most people. You can choose **6 (5.1)** for surround sound, but Shotcut does not yet support a variety of surround channel layouts.

For 6 channels, Shotcut uses the FFmpeg default channel layout for 6 channels: Front Left, Front Right, Front Center, LFE, Surround Left, Surround Right

The Preview Scaling setting may be used to reduce the resolution of video processing when you are using a fair amount of filters, transitions, and track blending. This may help with the performance during the editing process.



However, keep in mind that the source video or images must typically be scaled down to the designed preview scale, and that will have a processing cost. See [Settings > Proxy Editing](#) for the remedy to that. Proxy generates substitute clips at the same resolution as preview scale to eliminate the overhead of this scaling.


The submenu lists only the vertical resolution. The horizontal resolution depends on the aspect ratio of your project.

- **360p** is a good choice on a 1080p screen or less
- **540p** is good choice when working on a high density screen or screen with resolution > 1080
- **720p** is a good choice when working on a 4K screen or with 4K UHD footage and you want to get a little more clarity

Preview Scaling was added in version 20.02.

Export

Export > Advanced > Video has the option **Use preview scaling**.



This is intended to provide a faster export for draft review. You should leave the export resolution set to the project resolution unless your preset or video codec dictates a resolution.

Caveats

- The **Stabilize** video filter disables preview scaling for the clip on which it is applied because it is incompatible.
- Preview scaling is a step further away from WYSIWYG (what you see is what you get) and does not provide perfect fidelity and full precision. Some differences are unavoidable especially where blur, softness, and noise are concerned.

Introduction

Proxy editing is the process of creating and using low resolution videos and images in places of the original or optimized (**Convert to Edit-friendly**) files. This provides the advantage that your computer has less work to do in realtime while editing: decoding, scaling, and effects. This is, quite simply, because there are less pixels to compute. Then, when you are ready to export, it will use the original (or optimized) files for full quality.

Performance is optimal when the preview resolution matches the files it is working with. So, this feature is designed to work in conjunction with [Preview Scaling](#). You can still use the proxy mode without **Preview Scaling**, however, in case something is not working absolutely correct with **Preview Scaling**.

Proxy editing was added in version 20.06.

The format of a proxy file can be hotly debated because they typically have a number of goals that is difficult to achieve simultaneously: low resolution, small file size, fast to generate, fast to seek and decode, and decent image quality to do understand your work. Not everyone has these goals, but these are Shotcut's goals, and the format is not configurable. If you do not like how Shotcut generates proxy files, you can still create them with a different process (including Shotcut) if you understand where proxies are stored and how they are named. You are not required to use the MP4 file format for videos, but the name must end with `.mp4` for video and `.jpg` for images. Shotcut can almost always handle a video file whose name does not match the format (but almost never for an image). Fortunately, most people will be happy with JPEG for a proxy image even if not happy with our choice of MP4. A convenience of MP4 is that most media players can play them making it convenient when troubleshooting proxy files.

Settings > Proxy > Use Proxy

This either turns on or off the usage of proxy files for either the current project or the next project you open or start. When you turn this on and open a project, Shotcut looks for proxy files that already exist on your file system and uses them automatically and dynamically. However, it does not automatically generate a proxy file at this point if there are any missing. Any files added to the **Playlist** or **Timeline** will, however, generate a proxy file if one does not exist according to the following rules:

- It is a video (optionally with audio) or an image file.
- It is not an image sequence.
- It does not have an alpha channel (transparency).

- The video is not only **cover art**.
- It does not already have a proxy.
- Proxy for this file has not been disabled in **Properties**.
- Both the image width and image height are more than 1.3X the preview scaling resolution (or 540 if preview scaling is off).
- A proxy job for this file was not yet created.

It generates a proxy by queuing a job in **Jobs**. You can continue working with the original at this point. When this proxy job completes, it uses the **Replace** command to asynchronously update matching clips in **Source**, **Playlist**, or **Timeline**. In doing so, it tries very hard to retain all changes thus far including trimming, filters, and transitions. The vertical resolution of the proxy will be the same as your current **Preview Scaling** resolution. If **Preview Scaling** is set to **None**, then it uses **540p**. The width of the proxy will be whatever matches the display aspect ratio for the target height.

If you turn **off** proxy with a project opened, Shotcut automatically reopens the project. This causes it to no longer look for proxies during reopen, but it does take time for the project to reload, it clears undo history, and it resets selected clips. It does this without requiring you to save your current project, however.

If you turn **on** proxy with a project opened, Shotcut also automatically reopens the project as well while locating whatever proxies are currently available. Then, it prompts if you want to generate proxies for files that do not currently have them subject to the same rules above. These proxy jobs will **not** do a replace operation as they complete successfully. As the Replace command is still immature, we do not want someone to open a big old project, generate proxies for nearly everything and have it murder your project! Rather, once all of the jobs are done, reopen the project to pick up the proxies.

Settings > Proxy > Storage

Proxy files can be stored in one of two locations: a global folder or a project folder. When opening a project, it looks in both places. The global folder defaults to a sub-folder named “proxies” of your App Data Directory. The project folder is a sub-folder named “proxies” of a project folder created with **New Project > Start** on Shotcut’s startup screen.

Choose **Settings > Proxy > Storage > Set** to change your global project folder to a new location.

Choose **Settings > Proxy > Storage > Show** to open a folder view from your operating system’s file manager to see which folder is currently in use. This is where newly generated proxies are saved. The result of this action depends on whether you have a project currently opened or started and whether it has a project folder.

Choose **Settings > Proxy > Storage > Use Project Folder** to control whether you want proxies to be generated in a project folder, if in use, or always in the global folder.

Settings > Proxy > Hardware Encoder

You have the option to use your configured hardware encoder to generate proxies. The lead developer of Shotcut has not found much of an advantage in using it for these low resolution files if you have a strong CPU. Some systems with a weak CPU and compatible GPU may benefit from using it. For most users, do not expect it to significantly improve the speed. Since most hardware encoders tend to create larger files for the same quality as a software encoder, it uses HEVC to keep the file size reasonable. With that said, do not expect them to be much smaller than the software-generated H.264. In fact, they are going to be slower to decode. Still, it is there for your option and experimentation.

Properties

The **Properties** panel for video clips and images display **(PROXY)** next to the resolution when it is using a proxy file. You may also notice that **Properties** shows information about the proxy file instead of the original/converted except for the clip name, duration, frame rate, aspect ratio, and color range. Properties should continue to reflect any overrides where possible.

The **Properties** panel also has a **Proxy** menu button with the following options:

- **Make Proxy** - forces a proxy to be generated without adhering to the rules above. If you choose not to batch convert a project to proxies, use this to selectively generate proxies. The proxy job this generates will do a replace operation.
- **Delete Proxy** - Unfortunately, this is not available on Windows due to file locking prevents it from working (file is likely in use). Otherwise, this does what it says.
- **Disable Proxy** - prevents a proxy from being generated for this file. It only applies within the current project and not globally. Also, replaces all matching clips in the project with the original if it using a proxy. Re-enabling does not create or replace with proxy; you need to either reopen the project (if proxy exists) or choose Make Proxy.
- **Copy Hash Code** - shows a dialog with a 32 character alpha-numeric code that has also been copied to the system clipboard. This is helpful to track down problems with proxy files or to generate proxy files externally.

Export

Normally when you export, the proxy clips are replaced with their original or converted. However, if you go into **Advanced** mode and turn on **Video > Use preview scaling**, not only does export use the preview scaling resolution, but also

it uses proxy clips and images to further speed up export! This is intended to more quickly make a rough draft for review, not as a final output, of course.

Known Problems

1. A proxy can hide a problem with the actual source media, for example, frame accurate seeking on video.
2. There is no obvious way to add a proxy to a clip-only project. It is by design not to load or create a proxy until you add it to the playlist or timeline since the Source player can be used to quickly preview clips to decide whether to use one. However, it is possible. If you know the proxy already exists, you can save the project and reopen it to pick it up. Otherwise, choose **Properties > Proxy > Make Proxy**.
3. **File > Export Frame...** exports using the proxy up-scaled to project resolution it since uses the current image from the player.
4. It does not work with **File > Open MLT XML As Clip**. It does not replace the clips in the sub-project with proxies.



Scrub Audio was added in version [15.09](#)

In Shotcut, the audio is muted by default when you:

- Drag (scrub) the Playhead in the Timeline
- Move the Playhead forward or backward frame by frame (with the left/ right arrow keys, or the mouse wheel on the [Timecode](#))
- Increase the playback speed (forward or backward)

The Scrub Audio feature allows you to hear the audio while navigating through your timeline. This can be useful for precise audio editing and synchronization.

To enable/disable **Scrub Audio**:

- Go to **Settings > Player**.
- Check/uncheck the **Scrub Audio** option.



Settings > Player > Use JACK Audio

JACK (Jack Audio Connection Kit) is a low-latency audio system widely used on Linux in media workflows. It is designed for applications that require precise timing, flexible audio routing, and reliable synchronization.

When **Use JACK Audio** is enabled, Shotcut can:

- **Participate in a JACK-based audio environment**
Shotcut's audio output becomes available in the JACK graph, allowing it to be connected to other JACK-compatible tools such as digital audio workstations, audio processors, or monitoring utilities.
- **Provide advanced audio routing control**
JACK offers a centralized way to manage audio connections across applications and devices. With a JACK control tool (for example, QjackCtl), users can explicitly route audio to specific physical outputs or processing chains, which can be difficult to achieve with standard desktop audio systems.
- **Provide transport synchronization with other applications**
Transport control is playback control such as pause, play, rewind, fast forward, and seek. JACK lets applications that opt into this to control each other's playhead and playback. Obviously, this only works well when there are similar projects in each application. For example, one can do the video editing in Shotcut without audio do and all of the audio work for the same movie in a digital audio workstation (DAW). Run both at the same time, possibly on different screens, and load the respective projects for the same super-project (your movie). Now, when you seek, scrub, or play the video in Shotcut the corresponding audio project in the DAW plays at the correct time as well.

This option is intended for Linux and macOS users who already have the JACK audio system installed and running.

Purpose

- Enables Shotcut to connect to the JACK audio server if your JACK setup does not allow sharing the audio device.
- Allows precise audio routing and transport control between Shotcut and other JACK-aware applications!
- May provide lower latency compared to standard system audio backends.

Behavior

When **Use JACK Audio** is enabled:

- Shotcut sends and receives audio through JACK
- Audio ports appear in JACK patch bays (e.g. qjackctl, Helvum)
- Audio routing is handled externally by JACK
- System audio devices are no longer managed directly by Shotcut

When disabled:

- Shotcut uses the default system audio backend (e.g. PulseAudio / PipeWire)

Typical use cases

- Professional audio workflows
- Synchronizing Shotcut with DAWs or audio tools (e.g. Ardour)
- Advanced audio routing and monitoring
- Low-latency audio playback and capture

Requirements and notes

- A running **JACK server** is required
- If JACK is not running, enabling this option may result in **no audio**
- Audio connections must be managed manually using a JACK control tool
- Not required for standard editing or playback

Limitations

- Not available on Windows
- No effect unless JACK is installed and running
- Increased complexity compared to default audio handling
- Not intended for casual or beginner workflows



This option affects playback and controls two things:

1. when ON: drop some video frames to try to keep the audio continuous and keep the video running near real-time

This can make the video to appear choppy or slower frame rate than expected when it is on. Even when this is on it is possible that video processing is so intense that audio becomes choppy as well because Shotcut will not drop video frames indefinitely. In fact, currently it will only drop 5 consecutively (subject to change). Turn this off if you want to see every video frame regardless of how slow it goes or choppy the audio becomes.


2. when OFF: use multiple threads for image processing (filters, transitions, and track blending/compositing). Turn this setting off if you want to see if multiple threads can reduce some choppy playback or lag.

Additional Notes

- When this is option is off it turns on the same thing as **Export > Advanced > Video > Parallel processing**.
- Shotcut will use up to 4 background threads to render 4 video frames at the same time.
- Sometimes, this is the cause of a malformed image, which should be reported so it can be debugged.
- This is in addition to image slice-threaded processing used in many effects. The number of these threads is typically the same as the number of CPU threads. These threads do not consume nearly as much memory.
- The rule for the number of frame threads:

# CPU threads	# Shotcut threads
> 4	4
4	3
3	2
2	1
1	1

Controls how **interlaced video** is presented in the Shotcut player.



This option affects **preview display only**. It does not modify clips, timeline data, or export output.

The default is ON.

What “Progressive” means

Progressive video displays each frame as a complete image drawn in a single pass.

This contrasts with interlaced video, where each frame is composed of two alternating fields captured at different times. Interlacing was historically used in analog television systems to reduce bandwidth and improve perceived motion on CRT displays.

Common interlaced formats included:

- 480i (NTSC)
- 576i (PAL / SECAM)
- 1080i (HD broadcast)

Modern displays are progressive by nature and must convert interlaced material before showing it.

What the **Progressive** option does

When **Progressive** is enabled:

- Interlaced footage is shown as full progressive frames
- Field separation and combing artifacts are reduced or removed
- Motion appears cleaner and easier to evaluate

Note:

When Progressive is enabled, the **Deinterlacer** setting determines how interlaced video is converted to progressive frames for preview.

When **Progressive** is disabled:

- Interlaced footage may be displayed closer to its raw field-based structure
- Field artifacts (combing, line flicker) can be visible on motion
- Preview reflects the interlaced nature of the source more directly

Important scope clarification

This setting affects player preview only.

It does not:

- Change source media
- Alter clip or timeline properties
- Force progressive export
- Apply deinterlacing to the final output

Export behavior is controlled separately.

Visual characteristics

Progressive enabled

- Cleaner motion
- No visible field separation
- More readable preview on modern displays

Progressive disabled

- Possible combing on motion
- Visible interlacing artifacts
- Closer representation of raw interlaced fields


Recommended usage

- Enable **Progressive** for most modern workflows
- Especially useful when working with:
 - Interlaced broadcast recordings
 - PAL / SECAM / NTSC archival footage
 - Legacy television material
- If you disable **Progressive** to more directly view the interlacing, then you might also want to set the player zoom to 100% to prevent fields from mixing when scaling.

Limitations and notes

- This is a preview-only setting
- It does not replace proper deinterlacing during export
- It does not convert interlaced footage to progressive media
- Final output quality depends on export configuration

Selects the **algorithm** used to convert interlaced video to progressive for display in the Shotcut player.



This setting affects **preview quality and performance only**. It does not modify clips, timeline data, or export output.

Historical context

Interlacing originated in early analog television systems as a way to reduce bandwidth usage while maintaining acceptable motion smoothness on cathode-ray tube (CRT) displays. By transmitting alternating fields instead of full frames, broadcasters could double the perceived refresh rate without increasing signal bandwidth.

As display technology evolved, flat-panel screens (LCD, LED, OLED) became standard. These displays are inherently progressive and cannot show interlaced fields directly. As a result, interlaced video must be converted into progressive frames for proper viewing.

Deinterlacing exists to bridge this gap between legacy interlaced formats (such as PAL, SECAM, NTSC, and 1080i broadcast video) and modern **progressive** displays.

What Deinterlacing is

Interlaced video stores each frame as two fields captured at different times. Deinterlacing reconstructs those fields into a single progressive frame suitable for modern displays.

Different deinterlacing methods trade quality, motion accuracy, and performance.

Scope and behavior

- Applies to player preview
- Option is always selectable, regardless of the Progressive setting
- Has no effect on exports
- Does not alter source media

The selected method defines how interlaced content is processed when displayed.

Deinterlacer options

One Field (fast)

This is the default. It uses only one field to build each frame.

- Very fast
- Half vertical resolution
- No motion blending

Visual result

- Stable motion
- Softer image
- Reduced detail

Use when: performance is critical or hardware is limited.

Linear Blend (fast)

Blends the two fields together into one frame.

- Fast processing
- Simple averaging

Visual result

- Reduced combing
- Motion blur on movement
- Loss of sharpness

Use when: quick preview with fewer visible artifacts is acceptable.

YADIF temporal only (good)

Motion-adaptive deinterlacing using temporal information (differences between frames).

- Preserves more detail
- Balanced quality and performance

Visual result

- Cleaner motion
- Fewer artifacts than blending

Use when: regularly previewing interlaced footage.

YADIF temporal + spatial (better)

Uses both temporal and spatial analysis.

- Improved edge reconstruction
- Higher processing cost

Visual result

- Sharper edges
- Cleaner motion

Use when: preview quality is more important than performance.

BWDIF (best)

Advanced deinterlacing algorithm designed for high-quality reconstruction.

- Best motion handling
- Highest detail preservation
- Most computationally expensive

Visual result

- Smooth motion
- Minimal artifacts

Use when: inspecting interlaced footage carefully, archival work, or quality-critical review.

Performance considerations

Method	Quality	Performance
One Field	Low	Very fast
Linear Blend	Low–Medium	Fast
YADIF (temporal)	Medium	Moderate
YADIF (temporal + spatial)	High	Slower
BWDIF	Highest	Slowest


Limitations

- Preview-only setting
- Does not replace export deinterlacing
- Higher-quality methods may reduce playback performance

- No effect on progressive sources



Controls how Shotcut **resamples images** during player preview when video is scaled, repositioned, or rotated.



This setting affects **preview quality only**. It does not change clips, timeline data, or exported files.

What interpolation is

Interpolation is the computation of new pixel values based on **neighboring pixels** when an image does not align exactly with the pixel grid.

In practical terms, it affects:

- Scaling (zooming in or out)
- Sub-pixel positioning
- Rotation
- Motion that places pixels between pixel boundaries

Because pixels cannot be partially drawn, interpolation determines how smooth or sharp the result looks. This is sometimes described as *sub-pixel rendering*.

Scope and behavior

This setting applies only to the player preview.

It does **not**:

- Affect export quality
- Change render output
- Modify source media
- Apply to text, SVG, or HTML rendering

A separate interpolation setting exists in **Export > Advanced > Video**.

Interpolation options

Nearest Neighbor (fast)

No interpolation is performed.

- Uses the closest existing pixel

- No averaging or smoothing
- Fastest option

Visual result

- Sharp but blocky
- Visible pixel stepping
- Jagged edges when scaling or moving

Use when: performance is critical or when pixel-perfect preview is desired.

Bilinear (good)

This is the default. It interpolates using the four nearest pixels.

- Simple averaging
- Moderate smoothing
- Low computational cost

Visual result

- Softer than Nearest Neighbor
- Reduced jagged edges
- Slight blur

Use when: general editing with good performance.

Bicubic (better)

Interpolates using a larger neighborhood of pixels.

- More complex averaging
- Better edge preservation
- Higher processing cost

Visual result

- Sharper than Bilinear
- Fewer artifacts
- More natural scaling

Use when: preview quality matters more than speed.

Lanczos (best)

High-quality interpolation using a wider sampling window.

- Preserves fine detail
- Best edge reconstruction
- Highest processing cost

Visual result

- Very sharp
- Minimal blur
- Possible ringing on high-contrast edges

Use when: inspecting image quality or doing precision work.

Performance considerations


Method	Quality	Performance
Nearest Neighbor	Low	Very fast
Bilinear	Medium	Fast
Bicubic	High	Slower
Lanczos	Highest	Slowest

Usage notes

- Higher-quality interpolation improves visual clarity, not export quality.
- On slower systems, Lanczos may reduce playback responsiveness.
- This setting does not affect:
 - Text anti-aliasing
 - SVG rendering
 - HTML content



Adjusts **audio/video playback alignment** in the Shotcut player to compensate for synchronization differences during preview.



This setting affects **player playback only**. It does not modify clips, the timeline, or exported files.

What synchronization means

Audio/video synchronization (often called *lip-sync*) describes how accurately sound and image are aligned in time.

When synchronization is off, viewers may notice:

- Speech not matching mouth movement
- Sounds occurring slightly before or after visual events

Such offsets can be caused by:

- Audio hardware latency
- Video decoding delay
- Driver or system timing differences

Parameter

Video Offset (-250 ms to +250 ms)

Shifts the **video playback timing** relative to the audio during preview.

- **0 ms**
Audio and video play as decoded, with no adjustment.
- **Negative values (-)**
Video is displayed **earlier** relative to audio.
Use when video appears late.
- **Positive values (+)**
Video is displayed **later** relative to audio.

Use when audio appears late.

Important note:

The adjustment is expressed in **milliseconds (ms)**.

How to use this setting

1. Play a clip with clear audio cues (speech, claps, impacts).
2. Observe whether sound or image leads.
3. Adjust **Video offset** until audio and visual events align naturally.

This is especially useful when monitoring on systems with known latency.

Scope and behavior

This setting affects player preview only.

It does not:

- Alter clip timing
- Shift audio or video on the timeline
- Affect exports or renders
- Correct sync issues in the source media

Any adjustment made here is temporary and non-destructive.

Recommended usage

- Correcting preview sync issues caused by audio interfaces
- Compensating for Bluetooth or external audio device latency
- Improving lip-sync accuracy during editing

If sync problems persist in exported files, they must be addressed on the timeline or during export, not here.

Limitations


- Preview-only adjustment
- Global setting (not per clip)
- Does not fix incorrectly synced source media
- Not a replacement for proper timeline alignment

See also: [Audio/Video Properties](#) to learn how to adjust a sync problem with a clip

The **Rectangle Selection** tool, introduced in version [23.11.29](#), allows you to use your mouse to select a single clip or a group of clips in the Timeline.

The tool is enabled by default.

To disable or enable it, go to **Settings > Timeline**



There are a few supported actions for multiple selection:

- Copy
- Cut
- Lift
- Move to another place in the Timeline
- Ripple Delete
- Split (version [24.01](#) and up)
- Apply Copied Filters

Shortcuts

- When the tool is **disabled**, hold Shift to enable it temporarily
- When the tool is **enabled**, hold Shift to disable it temporarily



This timeline option was added in version 23.11. This option automatically adds an empty audio or video track to have ready so you do not need to add them manually. An empty track also helps to provide some space in order to select by dragging a box or rectangle or to provide more area for scrubbing.

It only adds tracks when you add something to the timeline or move a clip between tracks. Also, it does not automatically add a track when you open a project with no empty tracks.

In addition, this can be helpful to new or infrequent users who do not know you need to add an audio track or how to do that to add music. Or that you need to add a video track to overlay an image or video on top of another.



This is a sub-menu in **Settings > Timeline** that was added in version 23.09 with the following options:

- **Center the Playhead** - you can still move the playhead before the center point. But during playback, when the playhead reaches the center, it stops moving and the timeline starts scrolling in a smooth fashion.
- **No** - no automatically scrolling
- **Page** - when the playhead reaches the end of the timeline panel, the timeline suddenly scrolls by nearly the width of the timeline panel and stops scrolling until the playhead reaches the end again. This seems to be the most common among video editing tools - except perhaps by the page size.
- **Smooth** - the traditional behavior of Shotcut: when the playhead reaches the end of the timeline panel, it remains there and the timeline starts scrolling smoothly.

Scroll to Playhead on Zoom is a checkbox separate from the above options. This controls whether Shotcut scrolls the timeline to keep the playhead within view after changing the timeline level of zoom.



This setting is only available on Linux.

Shotcut uses OpenGL technology to draw its user interface and display video. OpenGL is typically implemented in the GPU of your computer and its driver in the operating system. These two pieces (Shotcut and the UI library it uses) and the OpenGL implementation need to communicate using a complex protocol. Sometimes, they do not communicate well and have compatibility or interoperability problems. Thus, there are some workarounds:

- use only **software** using a library called Mesa 3D. While this is known to be very compatible, it is much slower since it is not hardware-accelerated.
- on Windows, convert the OpenGL protocol to **DirectX** (Direct3D) using a library called **ANGLE**. While this is typically hardware-accelerated, Microsoft provides an automatic software fallback of its own called WARP.

If you choose **Automatic**, Shotcut's UI library chooses one of the options based on the model of your GPU, and usually that is OpenGL. You can see what Shotcut is using by looking in **View > Application Log...** from the main menu. Then, after about 100 lines, look for the lines start with

```
[Info ] Mlt::GLWidget::initializeGL
```

If you see the following, then it is using DirectX:

```
[Info ] Mlt::GLWidget::initializeGL OpenGL vendor "Google Inc."
```

```
[Info ] Mlt::GLWidget::initializeGL OpenGL renderer "ANGLE..."
```

If you see the following, then it is using Software:

```
[Info ] Mlt::GLWidget::initializeGL OpenGL vendor "VMware, Inc."
```

```
[Info ] Mlt::GLWidget::initializeGL OpenGL renderer "Gallium 0.4 on llvmpipe  
(LLVM 3.4, 256 bits)"
```

This setting corresponds to the key `opengl` in the [configuration file or registry](#).

See also

https://wiki.qt.io/Qt_5_on_Windows_ANGLE_and_OpenGL#ANGLE_Project



The App Data Directory is where Shotcut stores:

- automatically saved project files in the `autosave` sub-folder
- builtin and saved generator, filter, and export presets in the `presets` sub-folder
- custom video modes in the `profiles` sub-folder
- thumbnails and waveforms in the `thumbnails` sub-folder
- non-project-specific proxy videos and images in the `proxies` sub-folder
- favorited transition wipes in the `transitions` sub-folder
- `db.sqlite3` is a database where thumbnails and audio levels for waveforms are saved. (This database provides a more cross-platform way to order things by access/read time than a file system.)
- `shotcut-log.txt` which is a log file of messages from the Shotcut programming code, also viewable within the application through **View > Application Log...** This is overwritten on every run of Shotcut. Thus, if Shotcut is failing to start, there might be a clue in this file.
- **optionally** `shotcut.ini` if you used **Settings > App Data Directory > Set...** or **command line option** `--appdata` to store Shotcut's various settings or **configuration**. The idea here is that the app data directory can be placed in a portable location such as USB stick or network file share. Also, some people find it easier to locate and edit than the registry on Windows, plist on macOS, or `~/.config/Meltytech/Shotcut.conf` file on Linux.

Use **Set...** (*Settings > App Data Directory > Set...*) to change the location of the app data directory to make the Shotcut configuration and helper files portable (USB stick) or usable from multiple systems (network file share).


Use **Show...** (*Settings > App Data Directory > Show...*) to open the folder in your OS file manager (Explorer on Windows, Finder on macOS).



Settings > Job Priority

Low: Default (Recommended)

Normal: Option to improve performance on Intel 12th generation CPUs with E-cores.




Introduced in version [22.01.30](#).



This option removes all the entries in the **Recents** panel every time you cleanly exit Shotcut (not if it crashes). This prevents other users on the computer from seeing what you were working on.

There is also a command line option `--clear-recent` that triggers this.



If you don't see the **Playlist** panel in your interface, to open it you can either

- Use the `Ctrl` + `6` keyboard shortcut (`command` + `6` on MacOS)
- Click on the **Playlist** icon on the top toolbar.
- Open the **View** menu and click on the **Playlist** option.

The **Playlist** is a panel in the Shotcut interface where you can create a list of media files (such as video or audio clips) to be used in your project. This playlist panel allows you to easily organize and manage the media assets you plan to use in your editing timeline. You can add media files to the playlist, reorder them, and preview them before adding them to your project timeline for editing. This feature provides a convenient way to manage your media assets and streamline your editing workflow in Shotcut.


In Shotcut, the playlist feature can also be utilized independently of the timeline. This means you can use the playlist to organize and manage media files for projects that don't necessarily require a timeline-based editing approach (see [Editing](#) and [Exporting](#) in the Quick Start Guide). This flexibility allows users to leverage the playlist functionality for various editing needs beyond traditional timeline-based editing.

Slideshow Generator makes it easy to add more than one image or clip to the timeline and automatically:


- fill the frame,
- make a smooth zoom animation for still images, and
- add a transition between each item.

Steps to Follow

1. Open the **Playlist** panel
2. Add your selected images
3. Drag within **Playlist** (or use sort in its menu) to the desired order




4. Select the images or videos to add to the slideshow: hold **Shift** to select a range or **Control** (**Command** on macOS) to select individually
5. Click the menu button at the bottom of the **Playlist** panel
6. Choose **Add selected to Slideshow**
7. The **Slideshow Generator** dialog opens



8. Adjust the settings and click **OK**

9. Your images are added to the current track on the **Timeline** (a video track is added automatically if it was empty)




Suggestions

- **Image duration** only applies to still images and **Audio/Video duration** applies to everything else. Audio/Video duration defaults to 4 hours until you change it. Shotcut remembers the last value you used and defaults to that next time.
- **Audio/Video duration** still respects the source duration and its in & out points; duration here is a *maximum*.
- After it creates use the usual timeline editing methods to adjust the images and durations as required.
- If trying to fit a number of slides to an existing audio track:
Estimate the slide duration (Length of track / number of slides) and use this in the add dialog. If its too short or long, undo and repeat with a modified timing. Adjust the duration of the first or last clips for small adjustments.
- Set **Zoom effect** to 0 to turn it off. This is advised if most of the clips are videos or music.

Ideas

- Show off holiday or travel photos.
- Create a video montage.
- Make a video loop for an information screen.
- Use music files and set **Audio/Duration duration** very high to make a music bed track that cross-fades between each song.

When we write about the “player” in Shotcut, we are taking about the part that looks like below.



When you open a media file (or generator or device in Open Other), it opens in the **Source** player. It is called “Source” because it can be the source of edit operations in the **Playlist** or **Timeline**. The **Project** player is the preview and playback controls for the **Timeline** (or **Playlist** if you are not using the timeline or for advanced users).

“Open” can be any of:

- **File > Open**
- **File > Open MLT XML as Clip**
- **File > New generator** or **New Generator** on the main toolbar
- double-click a non-project (not .mlt) file in **Recent** or **Files**
- double-click or drag-n-drop a **Playlist** item
- drag-n-drop a non-project file from outside Shotcut to the player

Source is also a Shotcut clipboard viewer. Whenever, you **Copy** or **Cut** a *single* clip in the **Timeline**, it is also put into the **Source** player. That way you can use drag-n-drop or any timeline editing operation to place the clipboard contents into the timeline—not only **Paste** but also **Append** and **Overwrite**. Likewise, as mentioned above, you can add the Source clip to the **Playlist**.

While all other parts of the UI are panels that can be moved or made to float in a separate window, the player cannot currently. Also, currently, the two players are only available as tabs and not side-by-side like some other video editors. You can adjust the size, however, by adjusting the surrounding panels and the size of the main Shotcut window.


The **Loop** feature was introduced in Shotcut version [24.01.28](#)

Introduction


The Shotcut preview player has a loop feature that allows the player to loop around a region when it reaches the end.

Enabling Looping

The loop mode can be accessed from the Player menu



It can also be accessed by clicking on the loop button in the player controls




When looping is enabled, the loop button in the player controls will be highlighted




Click the loop button again to disable looping

When player looping is enabled, a highlight region appears in the player timeline to show the looping area



If the looped region is in the timeline, then the highlight will also appear in the timeline ruler



Loop Ranges

Shotcut provides the following options to select the region to be looped:

- Loop All - loop the entire clip or timeline (depending on what is selected)
- Loop Marker - loop around the area specified by a range marker that is currently under the cursor. If the cursor is not under a range marker in the timeline, this option will not work.
- Loop Selection - loop around the selected clips in the timeline. If no clips are selected in the timeline, this option will not work
- Loop Around Cursor - set a loop that starts one second before the current cursor position and ends one second after the current cursor position



The **Timeline** is the core of the Shotcut editing workflow. It is a multi-track, non-linear editing area where you arrange your media clips, apply transitions, and mix audio to create your final project. Unlike the **Playlist** (which is a simple list of files), the Timeline allows for complex layering and precise timing of video and audio elements.

Opening the Timeline

If the Timeline panel is not visible in your interface, you can open it using any of the following methods:

- **Keyboard Shortcut:** Press `ctrl + 5` (Windows/Linux) or `cmd + 5` (macOS).
- **Toolbar:** Click the **Timeline** icon on the top main toolbar.
- **Menu:** Go to **View > Timeline**.

The User Interface

1. Tracks

The Timeline is divided into horizontal rows called tracks.

- **Video Tracks (V1, V2, etc.):** Used for video, images, and transparent generators. Higher tracks (e.g., V2) will overlay tracks below them (e.g., V1) based on the **Blend Mode** (default is “Over”).
- **Audio Tracks (A1, A2, etc.):** Used for music, sound effects, and voiceovers.
- **Track Header:** Located at the far left of each track. It contains controls to **Mute** (audio), **Hide** (video), and **Lock** (preventing any edits to that track).
Hold Alt (`option` on macOS) while clicking **Mute** to affect all of the *other* tracks effectively making the track **solo**. The same can be done for **Hide/Show**.

While some video editors have somewhat complicated “arm” or “record” toggles on tracks, Shotcut simply has the concept of a *current track*, which has a gold color background. This sets the target for various operations. Clicking or selecting a track head or clip changes the current track automatically. There are also convenient keyboard shortcuts Up or Down.

2. The Playhead

The vertical white line with a “head” at the top indicates the current frame being shown in the **Project Player**. You can move the playhead by clicking on the time ruler at the top or by dragging the playhead itself.

3. Timeline Toolbar



The icon at the far left is a menu button. Click it to browse all of the possible Timeline actions. The rest of the toolbar contains a subset of these:

- **Cut** (scissors): This copies the the selected clip(s) to the clipboard, then removes them including the space they occupied, and moves all of the following clips to the left (ripple).
- **Copy** (2 pages): This copies the selected clip(s) to the clipboard. If only one clip is selected it is also in the **Source** player if you switch to it because the **Source** player in Shotcut is also an internal clipboard viewer.
- **Paste** (clipboard icon):** Insert the clip(s) in the clipboard memory at the current playhead time position on the current track.
- **New Generator (plus document icon)**: Opens a menu to insert a synthetic clip to an empty or new track at the current playhead. A generator or synthetic clip are things like text, solid background color, simple beep sound, or even a new drawing or 2D animation you can make with the [Glaxnimate](#) app!
- **Append (+)**: Append the current **Source** clip to the end of the current track.
- **Ripple Delete (-)**: Remove the selected clip and moves all of the following clips to prevent a gap.
- **Lift (^)**: Removes the selected clip(s) but leaves a gap.
- **Overwrite (v)**: Adds the current **Source** clip to the current track at the playhead but will not move anything out of the way or move any clips.
- **Split (|)**: Slices the selected clip at the playhead position. If no clip is selected, it uses the clip on the current track. If there is no clip on the current track, it looks for the topmost clip.
- **Marker**
- **Previous Marker (<)**
- **Next Marker (>)**
- **Snap (magnet icon)**: When enabled, clips will “stick” to the edges of other clips or the playhead for precise alignment.
- **Scrub (spaceship icon)**: Allows you to hear audio while dragging the playhead.
- **Ripple (circles icon)**: When enabled, moving or trimming a clip will cause all subsequent clips on that track to shift accordingly.
- **Ripple All Tracks (circle with * icon)**: When enabled, rippling affects clips across all tracks to try to keep things synchronized. Be careful with this! It often causes undesirable or unexpected changes. Therefore, we recommend to leave it off and turn it on only for possibly a single or short series of operations.
- **Ripple Markers (circle with bookmark icon)**: When enabled and an action causes clips to ripple, also move the markers after the start of the clip.
- 3 buttons and a slider to **Zoom** in and out the timeline.
- **Record Audio (microphone icon)**: Start recording audio at the current playhead on an empty or new audio track. This is especially useful for making a voice-over (narration or spoken word).

Basic Operations

Adding Media

1. **Drag and Drop:** Drag a clip directly from the **Playlist** or the **Source Player** onto a track.
2. **Append (A):** Adds the clip from the Source player to the end of the track.
3. **Overwrite (B):** Places the Source clip on the timeline, covering any existing media.
4. **Insert (V):** Places the Source clip and pushes existing media forward to make room.

Trimming and Moving

- **To Move:** Click and drag the center of a clip.
- **To Trim:** Hover over the left or right edge of a clip until the cursor changes, then drag to shorten or lengthen the clip.

Creating Transitions

To create a transition (such as a cross-dissolve), simply drag one clip so that it overlaps another clip on the **same track**. A “transition object” (marked with an ‘X’) will appear in the overlap area. You can select this object and use the **Properties** panel to change the transition type (e.g., Wipe, Iris, or Clock). You can use trimming to create a transition by dragging the start or end of a clip across a neighboring clip.

Advanced Controls

- **Track Operations:** Right-click the track header area to add new tracks, delete tracks, or change track height.
- **Output (Master) Track:** The very bottom header labeled “Output” allows you to apply filters (like a watermark or color correction) to the **entire project** at once.
- **Zooming:** Use the slider at the bottom right of the timeline, or the keyboard shortcuts **+** and **-**, to zoom in for detail or out for a project overview.

Common Keyboard Shortcuts

Action	Shortcut
Append to Track	A
Split at Playhead	S
Split all (unlocked) tracks	Shift+S
Lift	Z
Ripple Delete	X
Copy	C
Insert	V
Overwrite	B
Zoom In / Out	+ / -
Fit Timeline to Screen	0 (zero)



Shotcut supports standard clipboard operations for managing clips on the Timeline. These functions allow you to move media between tracks, duplicate clips, or transfer media from the **Source** player into your project.

Track Selection

Before performing any clipboard operation, you must ensure the correct track is selected.

- The current track is highlighted in the **Track Header** (gold colored background).
- **Paste** operations will always occur on the current track at the position of the **Playhead**.

Copying Media

Copy (`Ctrl+C` / `Cmd+C`) places a reference of the selected clip or transition onto the clipboard.

- You can copy a clip from the **Timeline** to duplicate it elsewhere.
- You can copy a clip from the **Source Player** to prepare it for insertion into the Timeline.
- **Note:** Copying a clip also copies any filters applied to that specific clip instance.

Cutting Media

Cut (`Ctrl+X` / `Cmd+X`) removes the selected clip from the Timeline and places it on the clipboard. Cut is a ripple action: it automatically shifts all following clips on that track to the left to close the gap.

Pasting Media

Paste (`Ctrl+V` / `Cmd+V`) places the contents of the clipboard onto the selected track at the playhead's current position. This too is a ripple action: it inserts rather than overwrites.

Timeline Toolbar Icons

While keyboard shortcuts are recommended for speed, these actions are also available as icons on the Timeline toolbar:

- **Cut:** Represented by the “scissors” icon.



- **Copy:** Represented by the “two pages” icon.



- **Paste:** Represented by the “clipboard” icon.



Advanced Paste: Append, Overwrite, and Replace

While standard pasting (`Ctrl+V`) is useful for general movement, Shotcut provides three “Advanced Paste” operations that offer more surgical control when moving media from the **Source player** (or the clipboard) onto the **Timeline**.

- **Append (A):** the fastest way to build a rough cut. Instead of worrying about playhead placement, this command ignores the current playhead position and places the clip at the very **end** of the current track.
- **Overwrite (B):** places the clip at the current playhead position, covering any media already on that track for the duration of the new clip. It does not shift any other clips on the timeline (regardless of Ripple settings). It acts like laying a new piece of physical film directly on top of the old one, “eating” into the existing footage.
- **Replace (R):** is a specialized tool used to swap an existing clip on the Timeline with a new clip from the **Source** player while preserving the exact duration and position of the original clip.

Select a clip on the timeline, open a new clip in the Source player, adjust its **IN** (starting) point, and press `R`. The timeline clip is swapped out. If the new clip is longer, it is automatically trimmed to the same duration. If the new clip is shorter, the replaced clip is shorter and either leaves or extends a gap.


Best used for swapping out a placeholder or temporary clip for a final version, or updating a clip while keeping all applied filters and transitions intact.

Summary of Shortcuts

Action	Windows/Linux	macOS
Cut	<code>Ctrl + X</code>	<code>Cmd + X</code>
Copy	<code>Ctrl + C</code>	<code>Cmd + C</code>
Paste	<code>Ctrl + V</code>	<code>Cmd + V</code>
Append	<code>A</code>	<code>A</code>
Overwrite	<code>B</code>	<code>B</code>
Insert	<code>V</code>	<code>V</code>
Replace	<code>R</code>	<code>R</code>




Output appears above the track headers in the **Timeline**:



(In versions before 20.06, this was named “Master.”)

Output refers the sum total output of the timeline - all tracks mixed (audio) and blended or composited (video). If your clips are the inputs to the timeline, this represents the output. It is very similar to the master output on [audio mixing consoles](#). You can click this to select it and then see **Properties** of the timeline (the attributes of your project’s Video Mode) or **Filters** that affect the entire timeline. When you have added a filter to **Output**, a funnel icon (same as the main toolbar icon that opens the Filters panel) appears in its block:



Often when people report all their video is black or shifted or all their audio is silent or too loud, they have added a filter on this accidentally and forgot about it.

Ripple All Tracks has the icon  in the **Timeline** toolbar and menu. Here is how it works:

- Only applies when **ripple** mode (**Ripple trim and drop**) is also on: 
Or when using an edit operation that ripples: **Cut** or **Remove** (Ripple Delete).
- Only affects the clips after the moved clip on the *same* track.
- Affects footage after the original *start position* of the moved clip on *other* tracks (I use the term “footage” when I mean there is a time range and not strictly whole clips, and that often includes a split.)
- When moving a clip right-to-left it simply pushes including a split on another track per the moved clip’s original *start position*.
- Moving left-to-right is more complicated and usually has 2 actions:
 1. Remove the footage from the moved clip’s original start-to-end positions including other tracks.
 2. Insert the moved clip to the new position, which shifts the clips on the same track to the right but also shifts footage on other tracks the same duration equal to the moved clip’s duration.

The keyboard shortcut is Ctrl + Alt + R (option + command + R on macOS).

Description

The **Nudge** tool lets you make precise adjustments to the position of a clip in the Timeline without having to rely solely on dragging with the mouse, which can sometimes be challenging. Each *nudge* moves the selected clip one frame to the left or to the right.

Introduced in Shotcut version [24.01.28](#).

How to use

There are three ways to access the Nudge tool:

TIMELINE MENU

- Select a clip in the Timeline
- Go to: **Timeline** menu > **Edit**
- Click on **Nudge Forward** or **Nudge Backward**


RIGHT-CLICK

- Right-click on a clip in the Timeline
- Click on **Nudge Forward** or **Nudge Backward**

KEYBOARD SHORTCUTS

This is the most efficient way to use the Nudge tool.


- Select a clip in the Timeline
- Press the , (coma) keyboard key to nudge the clip to the LEFT
- Press the . (dot) keyboard key to nudge the clip to the RIGHT



(Animated GIF - Click on the image to begin playback)

Examples

The Nudge tool is especially useful when precise alignment between a clip and a specific moment on an audio waveform is needed.



(Animated GIF - Click on the image to begin playback)

It is also very handy when you need to move a clip in a cluttered timeline.



(Animated GIF - Click on the image to begin playback)


Limitations

- The Nudge tool will not work when multiple clips are selected.
- You cannot nudge a clip located at position 00:00:00:00 on the timeline.
- It will not work on a **Group** of clips.
- You cannot use the Nudge tool to create a transition.

Nudging multiple clips

While you can't use the Nudge tool when multiple clips are selected, if you enable **Ripple** or **Ripple + Ripple All Tracks** you will be able to nudge multiple clips.


- When only **Ripple** is enabled: The selected clip and ALL subsequent clips on the same track will nudge.
- When **Ripple** and **Ripple All Tracks** are enabled: The selected clip and ALL subsequent clips in the Timeline will nudge. Be cautious, however. Depending on the position of the selected clip, with **Ripple All Tracks** enabled, you run the risk of splitting the clips on the tracks above and/or below.



(Animated GIF - Click on the image to begin playback)

Markers are like time-based bookmarks for your project to help you remember and quickly go to sections.

Timeline Markers were added in version 21.10.31.



- Press **M** or click toolbar button the first time to add a marker at the playhead (current position) with no dialog.




- Markers appear in the time ruler/track/row at the top of the **Timeline**.
- Drag a marker to move it.
- There are several ways to seek to a marker: Click directly on a marker, press **<** or **>**, or click the seeking buttons on the timeline toolbar.




- If snapping is turned on then dragging clips or their edges snaps to markers and dragging markers snaps to clip edges or the playhead.
- Press **Alt + R** (**option + R** on macOS) or click on the **Ripple timeline markers with edits** button on the **Timeline** toolbar to have the timeline markers shift according to the direction that the clips are dragged to in the timeline. Also works in conjunction with either or both of the other two ripple functions (**Ripple trim and drop** and **Ripple edits across all tracks**). Press **Alt + Shift + R** (**option + Shift + R** on macOS) to turn on and off all three ripple icons at the same time.




- Click toolbar button or press **M** when the playhead is on the exact start time of a marker to open the **Edit** dialog where you can change the name, color, start, and end times.





- A color change is remembered for all new markers until a new color is chosen.
- Right-click a marker to open its context menu to **Edit**, **Delete**, **Choose Color...**, or **Choose Recent Color**.




- Press Ctrl + Shift + M when the playhead is on the exact start time of a marker to **Delete** it.
- Ctrl +drag (command on macOS) a marker to changes its duration (a marker with a duration longer than 1 frame is also called a “range”).
- When the mouse cursor is over a marker then a tool tip appears with the name and time. If the marker is longer than 1 frame (range) then the tool tip will also show its end time and duration.




- **View > Markers** or press Ctrl + Shift + 6 (shift + command + 6 on macOS) to open the Markers panel where all markers are listed and can be categorized by several columns (**Color**, **Marker**, **Start**, **End**, and **Duration**). Clicking on a marker in the list will seek to that marker in the timeline. Click on a column header to change a column’s sorting to be either ascending or descending. Under the markers list there are buttons to **Add a marker**, **Remove the selected marker** and **Deselect the marker**. Click on the hamburger menu to **Remove All Markers** and to remove or add any of the columns. A search bar is also available to filter the list by name. The rest of the panel offers the same options as the **Edit** dialog for markers in the timeline where the name, color, start, and end times can be changed.



- **Export > From** lists all range markers (duration longer than 1 frame). Selecting a marker here allows the exporting of just that section of the timeline.



- **File > Export > Markers as Chapters...** outputs a .txt file in the format used to make chapters for YouTube. Only timeline markers with a duration of 1 frame (default) are exported as chapters.




- Markers are saved to the project file in a XML tag named `<properties name="shotcut:markers">`.
- Marker operations support **Undo** and **Redo**.




1. Setting your microphone: (This only needs to be done once for Shotcut on Windows or macOS)

Open Other - Audio/Video Device or **Settings > Audio/Video Device...**




Click on Open

2. Move the **Timeline** play head to the time where to start the recording.
3. Click **Record Audio** (microphone icon) on the **Timeline**.



4. If you did not create a project folder using **New Project > Start** a file save dialog appears to name the audio file. Click **Save** for the file name.




5. It uses the *current* track if it is an *audio* track and *empty* at the play head and *beyond*. Otherwise, it adds an audio track.
6. Recording starts
7. During recording the project plays without audio and shows a growing audio clip.
8. When done, click **Record Audio** again to stop recording.



The Align To Reference Track tool allows multiple clips to be aligned to the audio in a reference track. This can be useful to align clips that were recorded at the same time from multiple camera angles.

Align To Reference Track was added in version 22.06.23.



Limitations


- This tool can only be used to align clips with audio that is similar to the audio in clips on another track. Both the clip to align and the reference track must have audio that is similar so that the algorithm can detect the alignment.
- The tool does not use timecodes for alignment

Usage


Place clips with similar audio on separate tracks. For example, put all the clips from each camera source on its own track.

Choose one track to be the “Reference Track”. Typically, this will be the audio from the best audio source. For example, use the camera that had the microphone closest to the presenter.


Select the clips to be aligned (do not select any clips on the reference track)




Right-click on the selected clips and choose from the menu: More->Align To Reference Track




In the dialog, select the reference track. Then click “Process”




If the results look good, click “Apply”.




Alternatively, both “Process” and “Apply” can be run as one action by clicking “Process + Apply”.



The selected clips will be moved on the timeline to be aligned to the reference.



Speed drift adjustment



If the selected clips can not be aligned, or if the alignment drifts over time, re-run the tool and select “Speed adjustment range” and click “Process”. With this option enabled, the alignment algorithm will try to detect alignment by speeding or slowing the clips. The clip speed difference can be detected by up to +/-0.5%.

If the tool is able to detect alignment with the “Speed adjustment range” option, after clicking “Apply”, the clips will be moved and their speed will be changed to make the alignment.



The Apply Copied Filters option was added in the 24.04 release.

“Apply Copied Filters” is an option in the Timeline context menu which allows a filter to be applied to many timeline clips at the same time. It can also be used to update existing filters with the same value.

Applying filters for the first time

This video shows how to create a filter and then apply it to multiple timeline clips

Updating previously applied filters

This video shows how to easily modify the filter on multiple clips. First the filter is modified on one clip. Then, the filter is copied and applied to the other clips. Note: if the filter already exists, Apply Copied Filters will replace the filter with the new filter. It will not add the filter multiple times.



There is a black or gray line that appears across timeline clips that have audio if **Settings > Timeline > Show Audio Waveforms** is on.



Adjust the volume of a clip by dragging this line up or down. Doing so adds a **Gain/Volume** audio filter to the clip's **Filters**. Double-click it to reset it back to 0 for no change (or you can remove the added filter).

Use **Settings > Timeline > Adjust Clip Gain/Volume** to turn this feature off.

When the line has not been moved, it represents -1 dBFS. With it you can visualize when the audio level is dangerously close to clipping. However, keep in mind that there can be audio clipping that occurs due to other Gain/Volume filters or mixing with clips on other tracks.

This feature was added in version 25.07.

See also [Gain / Volume Audio Filter](#)

Many actions for clips in the timeline require that you select one or more of them. There are also some actions that only work when a single clip is selected. Click the menu button in the top left of the **Timeline**: there is a **Selection** menu.

Selected clips on the timeline have a red outline unless it is part of a group. When a Group is selected, the clips have a white outline.

Select Single Clip

Usually, you simply click a clip with the left mouse button to select it. However, you can also select the clip under the playhead/cursor using the keyboard combination **Ctrl** + **Space** (**command** + **space** on macOS). If there is more than one clip under the playhead, it selects the top clip.

You can select the neighbor of a selected clip (move the selection) by holding **Ctrl** (**command** on macOS) with left, right, up, or down arrow keys.

Multiple Selection

Here are some of the actions you can perform with multiple clips selected:

- Copy
- Cut
- Lift
- Ripple Delete
- Move
- Split
- Apply Copied Filters
- Align To Reference Track

To select a range (multiple consecutive clips) click the first (or last) clip, press and hold **Shift** and click the last (or first) clip.

To select or deselect multiple clips individually hold **Ctrl** (**command** on macOS) while clicking the clips.

Rectangle Selection

You can also use the mouse to drag out a rectangle to select multiple clips. If **Settings > Timeline > Rectangle Selection** is on simply click an empty area of the timeline near a corner where you want to start. Then, holding the left mouse button down drag to select. Release the left mouse button when you are done.

If that setting is turned off, hold **Shift** to draw the rectangle (without Shift clicking seeks).

To facilitate this feature, it is helpful to have empty space above or below the clips. You can start the rectangle by clicking below the existing tracks. To have an empty track at the top turn on **Settings > Timeline > Automatically Add Tracks**.

Group/Ungroup

A group is a multi-selection that is saved with the project. To make one, select multiple clips, right-click any member, and choose **Group/Ungroup** from the context menu. Alternatively, use keyboard shortcut **Ctrl + G** (**command + G** on macOS). Now, whenever you click a member of a group, it selects all of the members and show a white outline instead of red. You can have multiple groups but not overlapping groups.

Frequently you need to work with only one member of a group. You could ungroup and re-group later, which is tedious, or:

- If the group is already selected **Ctrl**-click the one to use.
- If the group is not selected **Alt**-click to select a single member of a group.

Select Track

Did you know that you can not only select clips but also a track? With a selected track, you can add Filters to it that affect all clips on that track! There is a special track at the top of the track headers labeled **Output**, and by selecting that you can add a filter to everything in the project. However, it is better to think about that as putting an effect between the timeline and the output–preview or the export.

However, there is also a *current* track. It has the gold color background. Selecting a track or clip also makes it current, but it also possible to use the keyboard to change the current track: **Ctrl + Alt + Up** or **Down** (**option + command** on macOS).

Some actions depend on the current track. See the **Track Operations** in the timeline's menu button:

- Insert Track
- Remove Track
- Move Track Up
- Move Track Down
- Show/Hide
- Lock/Unlock
- Mute/Unmute



Use the **Freeze Frame** Timeline action to create an effect where the video pauses for a short time. This is not for playback control during editing. This action was added in Shotcut version 25.05.

To use this move the playhead (also known as the “cursor” or rather grandiloquent “current time indicator”) to the time position where you want to hold. Then, right-click the clip and choose **Freeze Frame**. This creates a one second clip with only that frame. To change its duration, either turn on **Ripple** mode on the timeline toolbar or press and hold Shift. Now, drag the right edge of the clip to adjust it.

Feel free to make your own keyboard shortcut for this action in **Help > Actions and Shortcuts**.

The **Replace** action is a powerful timeline utility designed to swap an existing clip with a new one while maintaining the structural integrity of your edit—specifically preserving filters and track positioning.

Replace first appeared in version [20.04.12](#)


Here is a comprehensive breakdown of how it works, its behavior regarding clip length, and best practices for managing keyframes during the process.

Overview

The **Replace** action allows you to swap a clip on the **Timeline** with whatever is currently loaded into the **Source** player. Unlike a “Lift and Paste” operation, **Replace** is non-destructive to your filter stack; it swaps the underlying media while keeping any applied effects or corrections intact.

How to Execute a Replace

1. **Prepare the Source:** Load your new clip into the **Source** player. You can do this by:
 - Opening a file directly from your computer.
 - Dragging a clip from the **Playlist**.
 - Copying a clip from elsewhere in the **Timeline**.
2. **Target the Timeline:** Right-click the clip you wish to swap out on the Timeline.
3. **Apply:** Select **Replace** from the context menu.



The **Replace** command is not limited to existing media; it can also be used to fill empty spaces on a track without disturbing the rest of your edit.

- **How it works:** Right-click any empty (blank) area between two clips and select **Replace**.
- **Duration Constraint:** The blank space acts as a strict boundary. If the gap is 2 seconds long, Shotcut will take the first 2 seconds of the Source clip (starting from its current **In point**) and drop it perfectly into that slot.
- **Layout Integrity:** Crucially, this operation does **not** trigger a “ripple” effect, even if **Ripple** and/or **Ripple All Tracks** are enabled. The clips to the right of the gap will stay exactly where they are, ensuring your timeline synchronization remains intact.

Clip Length & Ripple Behavior

The behavior of the Timeline after a replacement depends entirely on the duration of the Source clip relative to the original Timeline clip.

Case A: Source Clip is Shorter

If the incoming clip is shorter than the original:

- The timeline clip's duration will shrink to match the source clip.
- **Warning:** If the original clip had keyframes (e.g., a size/position/rotate movement), they may be cut off or misaligned because the clip's tail has been truncated.

Case B: Source Clip is Longer

If the incoming clip is longer than the original:

- The timeline clip will initially maintain its original duration.
- **Expansion:** You can manually extend the clip by dragging the right trim handle to reveal the extra footage from the source clip, up to its maximum length.


Key Features to Remember

Feature	Behavior
Filters & Effects	All filters on the original clip are preserved and applied to the new footage automatically.
Transitions	Replace may affect existing transitions if the new clip length is shorter than the transition point.
Keyframes	Filters are kept, but if the clip length changes significantly, you should check your keyframe timings to ensure they still align with the new footage.




In Shotcut, a transition is a specialized object created when two clips on the same track overlap. While many video editors require you to drag an effect onto the junction of two clips, Shotcut uses a more intuitive overlap method to create cross-fades, wipes, and other visual transitions.

Drag and Drop



To create a transition such as a video dissolve with an audio cross-fade between two adjacent clips on the same track:

1. **Position the Clips:** Place two clips next to each other on a track.
2. **Overlap:** Drag the second clip toward the first clip so that it overlaps the end of the first clip.
While you are dragging, a popup appears with a time value to show you how much you moved the clip and thus the duration of the transition.
3. **The Transition Object:** You will see a new rectangular object appear in the overlap area with a purple “X” shape inside it. This indicates a transition has been created.



4. **Playback:** When you play the timeline, the first clip will now smoothly fade into the second clip.


In the screenshot above, this created a gap in the track because **Ripple** is off. If you turn it on, there will be no gap and all of the clips following the clip you dragged are moved.



You can also drag the first clip to the right to make a transition; however, **Ripple** must be off for that. Otherwise, ripple drag to the right is used to push the clips on the track.

Trimming

Another way to make transitions if there is enough non-revealed footage before or after the clip (due to trimming creating so-called “handles”) is by trimming:



Drag the right edge of the first clip or the left edge of the second clip over the other clip.

Adjusting a Transition

- **Duration:** To make the transition longer or shorter simply drag the edges of a neighbor clip or the transition object, or move the clips to change the amount of overlap.
- **Selection:** Click the transition object to select it. When selected, it will have a red border, and its properties will be available in the **Properties** panel.


Removing a Transition

Immediately after making a transition, possibly by accident, you can simply **Undo** (Ctrl+Z, command+Z on macOS). Otherwise, there are two behaviors: **Lift** and **Ripple Delete**.

If you do not want clips to move choose **Lift**, the transition is removed, and each clip is automatically trimmed to fill the gap.

If you want to retain the full duration of the clips and not change their trimming, clips must move because two things were occupying the same space-time. Use **Ripple Delete** for that.

With a transition object selected, open the **Properties** panel (`Ctrl+2`) to change how the transition looks:



Parameters

Video

A dropdown menu containing various transition types (Dissolve, Horizontal/Vertical Wipe, Iris, Barn Door, etc.). Dissolve is the default. The others are kinds of wipe. A wipe is an effect that does not distort the clips at all; they only change how one blends with the other using some shape or pattern.

Choose **Cut** to not do any effect. This is useful to have and audio cross-fade with a hard change from clip A to clip B. Choosing this changes the slider to **Position** to control when the cut occurs.

Custom opens a whole new world of wipe patterns. See [Creating a Custom Video Wipe](#) and a large supply of community-made ones in the *Transitions* section on the [Shotcut - Resources](#) Web page. Choosing this opens a file dialog and you pick a file that you downloaded or made. To save a custom file you have selected to the list of wipes for convenient, future use: click the star button to add it as a favorite.

Invert Wipe

Reverses the direction of the wipe or animation.

Softness

Adjusts the blurriness of the edge for wipe transitions.


Audio

Choose between **Cross-fade** (the default: smoothly mixing the sound over time) or **Mix** (adding the audio levels together).


Video Preview

The embedded viewer provides a quick preview of the wipe without having to play it in the main player. It starts automatically when you choose a wipe. You can also seek and scrub the mini player. After the preview has played once click **Play** to play it again.

In the **Properties** panel, you can select **Custom...** from the **Video** list-box. This allows you to load any grayscale image or video file to use as a transition map (wipe pattern).




- White areas of the image appear first. Black areas appear last. For example, `Subspots.png` :



- The duration of a video wipe is not automatically adjusted to the same duration as the transition. It is better to use one that is similar duration or adjust the duration of the transition, but often close is good enough.
- The file you choose is displayed below the list-box.
- Click the favorite button (star icon) to save this custom file as a favorite that will appear at the bottom of list-box for future transitions. This copies the file to your **App Data Directory** so you will not need to keep the current file in the future. Hint: To make the current transition use the custom file in App Data, un-select the transition, re-select it, and choose it at the bottom of the **Video** list box.

- You can download many custom wipes from [Shotcut - Resources](#).

Video Multi-track Transitions



Multi-track transitions are created when **two clips on different tracks overlap in time**.

Shotcut blends the overlapping portion to transition from one clip to another.

This technique is commonly used for fades, crossfades, compositing, and layered transitions.

Unlike same-track transitions, multi-track transitions rely on the **track stacking order**: clips on higher tracks appear above clips on lower tracks.

Basic concept

In the timeline:

- Tracks are layered vertically
- The **top track visually overlays** the tracks below
- When clips overlap in time, the top clip normally hides the one underneath

To create a transition, the top clip must become **partially transparent** during the overlap.

Method 1: Track-level fade controls

Shotcut provides **direct fade controls on clips in the timeline**, allowing quick fade-in and fade-out adjustments without adding filters.


Each video clip has small **fade handles** at the **top corners** of the clip in the timeline:

- **Left corner** — controls the **fade-in**

- **Right corner** — controls the **fade-out**

To create a fade:

1. Move the mouse over the top-left or top-right corner of the clip.
2. When the fade handle appears, **click and drag inward**.
3. The shaded area shows the duration of the fade.



During playback, the clip gradually transitions from **transparent to visible (fade-in)** or **visible to transparent (fade-out)**, revealing tracks below if present.

These fades are equivalent to applying video fade filters but provide a **faster, visual method** for common transitions.

Notes:

The fade duration can be adjusted at any time by dragging the handle again.

Method 2: Using the Opacity filter

This is the most common method with the [Opacity filter](#).

Steps

1. Place two clips on **different tracks**.
2. Overlap the clips on the timeline.
3. Select the **clip on the upper track**.
4. Add the **Opacity** filter.
5. Add [keyframes](#) to animate opacity.

Example fade:

- Start of overlap → Opacity **100%**
- End of overlap → Opacity **0%**

Result: the upper clip gradually reveals the lower clip.

Method 3: Using fade filters

Shotcut provides dedicated fade filters that automate opacity animation.

Common filters include:

- [Fade In Video](#)
- [Fade Out Video](#)

Example crossfade

1. Place clips on two tracks.
2. Overlap them.
3. Apply **Fade Out Video** to the upper clip.

The top clip fades away while the bottom clip becomes visible.

Method 4: Using mask filters

Mask filters can define **which parts of a clip are visible** during the overlap.

Examples:

- [Mask: From File](#)
- [Mask: Simple Shape](#)

This allows transitions such as:

- wipes
- reveals
- geometric transitions
- custom image-based transitions

The mask determines which pixels of the top clip reveal the clip below.

Method 5: Using blend modes

Shotcut includes different [blend modes](#) that combine the pixels of two tracks.

Examples include:

- Add
- Multiply
- Overlay
- Screen

Instead of fading between clips, these modes **mathematically combine colors** from both layers.

This technique is often used for:

- creative transitions
- visual effects
- compositing

Method 6: Using animated transforms

Transitions can also be created by **moving or scaling the upper clip**.

Common examples:

- slide transitions
- zoom reveals
- picture-in-picture transitions

Typical workflow:

1. Place clips on two tracks
2. Apply **Size, Position & Rotate**
3. Animate position or scale with keyframes

This causes the top clip to move away and reveal the clip below.

Track order behavior

Track stacking determines visibility:

- Higher track = rendered above
- Lower track = background layer

If the top clip is fully opaque, the clip below will not be visible.

Typical examples

Common multi-track transitions include:

- crossfade
- fade-to-black
- wipe transitions
- split-screen reveals
- animated slide transitions

These are all variations of **revealing the lower track over time**.

Usage notes


- Overlapping clips is required.
- The top clip usually carries the transition filter.
- Keyframes allow precise timing control.
- Complex transitions may combine multiple filters.

Limitations

- Requires multiple tracks

- Manual setup compared to same-track transitions
- Incorrect opacity or masking may hide lower clips completely

Audio Multi-track Transitions



Audio transitions can also be created when **clips on different tracks overlap in time**.

During the overlap, the sound from one clip gradually decreases while the other increases. This is commonly called an **audio crossfade**.

Unlike video, track order does **not affect audio**. All audible tracks are mixed together during playback.

Basic concept

When two audio clips overlap:

- Both clips play simultaneously
- Their relative levels determine which one is heard more strongly
- A transition is created by **changing the volume over time**

Method 1: Using Fade filters

The simplest method is to apply **audio fade filters**.

Common filters:

- [Fade In Audio](#)
- [Fade Out Audio](#)

Steps

1. Place clips on **different tracks**.
2. Overlap them on the timeline.
3. Apply **Fade Out Audio** to the first clip.

4. Apply **Fade In Audio** to the second clip.

Result: one sound fades out while the other fades in.

This produces a **smooth audio crossfade**.

Method 2: Using the Gain / Volume filter

Transitions can also be created using the **Gain / Volume** filter with keyframes.

Steps

1. Add **Gain / Volume** to a clip.
2. Add keyframes to the **Level** parameter.

Example fade:

- Start of overlap → 0 dB
- End of overlap → very low level (e.g. -70 dB)

On the second clip, perform the reverse animation.

This method allows **precise control** over fade timing.

Method 3: Using track-level gain

Audio transitions can also be managed at the **track level**.

Instead of adjusting individual clips:

1. Use track gain or track filters.
2. Animate levels across the overlapping region.

This approach is useful when multiple clips share the same transition behavior.

Method 4: Combining with audio effects

Transitions can also include **audio processing effects** during the overlap.

Examples:

- **EQ filters** for tonal change
- **Reverb** to simulate distance
- **Delay** effects
- **Compression** adjustments

These effects are optional but can help create more complex transitions.

Typical examples

Common audio transitions include:

- dialogue crossfades
- background music transitions
- ambient sound blending
- smooth scene changes

These transitions avoid abrupt audio cuts.

Usage notes

- Audio crossfades usually sound best with **moderate overlap**.
- Very short overlaps can sound abrupt.
- Music transitions often require longer fades than dialogue.
- Use headphones or speakers to verify the result.

Limitations

- Audio mixing depends on track levels.
- Incorrect gain settings may cause clipping.
- Too many overlapping tracks may reduce clarity.



The Speech To Text tool can analyze the audio for a project and generate text in the [View > Subtitles](#)

Speech To Text was added in version 24.10.29.


Using The Tool

1. Place your video in the Timeline.
2. In the **Subtitle** panel, click on the **Detect speech...** button
3. Wait for the 2 jobs to complete (the **Speech to Text** job might take a while to complete depending on the length of your video).

About Speech To Text

Shotcut's Speech To Text feature uses AI based on [OpenAI's Whisper](#), courtesy of the [whisper.cpp](#) project.

Our builds include a basic model that has decent speed and accuracy but not a big size. (You can think of the model as the brain.) You can [download](#) a bigger and better brain (model) in `ggml` format and configure it in the **Speech to Text** dialog, but it will be slower.



The dialog creates two jobs that appear in the **Jobs** panel: one to export audio and another to convert to text. The results are added to the **Subtitles** panel as a new top-level Subtitle Track.

Currently, the only GPU our build supports is Apple Silicon. Otherwise, it is heavily multi-threaded on the CPU.

Known Quirks:

- Subtitle items sometimes start earlier than expected. Timing is provided by the model and tool, and we lack the skills and resources to improve this.
- Expect there to be occasional errors. Like humans and non-ideal conditions, it is not perfect. We will not take action on bug reports about some piece of audio not converting to the expected text.


OpenAI has made some [warnings](#) about the usage of their Whisper models:

In particular, we caution against using Whisper models to transcribe recordings of individuals taken without their consent... We recommend against use in high-risk domains like decision-making contexts, where flaws in accuracy can lead to pronounced flaws in outcomes.

Text to Speech, for example, converts the text “Hello, my name is Heart. What do you want me to say?” to audio:

0:00 / 0:03

It is available since version 25.10 in the **Notes** panel as well as **Subtitles**. Choosing it opens this dialog:



Unless, you do not yet have Docker installed on your computer. Then, it shows a dialog about Docker:

This feature requires **Docker**, which provides an installer, and the automatic download of an **13.2 GB** Docker image.

If you already installed Docker it could not be found at the expected location:


Click **OK** to continue and locate the `docker` program on your system.

Click the  button next to the Voice to hear a preview of it. The symbol  or  at the beginning of the voice name indicates gender.

Clicking **OK** opens a file save dialog if you did not yet specify the output file name. Then, it creates a job in the **Jobs** panel. It can take quite a long time especially if you have not yet downloaded the large Docker image. The first time you use this in an app session it always generates a “docker pull” job to get the latest version. But that goes quickly if there are no updates (it rarely updates). The first time ever it will take a long time to download that much data, and it is normal and expected that the progress % appears stuck because `docker` is not able to show a good progression. Finally, a second job to actually do the conversion runs, and that is not very fast either. When, it completes, Shotcut opens the generated WAV file in the **Source** player.

- This uses **Docker** like a plugin framework. The engine for this is **Kokorodoki**, and the model is Kokoro—both of which are not made by us. Do not ask us for more languages or voices.
- There are Docker installers for Windows and macOS from **docker.com**. For Linux, it is usually preferable to get it from your distribution but ensure you get the real docker and not podman or the desktop icon dock bar. On Debian-based systems, it is the `docker.io` package.
- Docker has an engine (service) that must be running to use this feature. If you are on Windows or macOS, this might not be running after a reboot. Open the **Docker Desktop** app where you can turn it on, and in its Settings there is an option **Start Docker when you sign in to your computer**. On Linux, it is usually automatically started unless you or something disabled it in systemd.
- The quality with subtitles is heavily dependent upon the timing and duration of each item. If it sounds choppy or cut-off, you either need to increase the speech speed and/or the item durations. Also, multi-line subtitle items are discouraged because that introduces a pause as it thinks it is like a new paragraph.
- This is not available in the Linux Flatpak.

The **Export** panel (also **File > Export Video...**) is used to create a new video or audio file from your project because **File > Save** saves a project file. Think of this like a photo editor that saves its own format (e.g. .psd or .xcf) versus saving a JPEG or a word processor that saves to its own format (e.g. .doc or .odt) versus a PDF.



The most important thing to know is to click **Export File** to export your project as a new video file. It opens the standard save file dialog for your operating system where you must choose a folder and give it a name.

Here are the other things to know about the Export panel in its basic mode:

- Click **Advanced** to show many more controls and options. However, by clicking **Advanced** you should understand that Shotcut expects you to be advanced and know what you are doing. There are simply too many combinations of settings and options that Shotcut's advanced export mode lets you create invalid combinations and does not protect you from yourself.
- The **Reset** button is used in advanced mode to reset all the options and settings to their default values.
- The list of many things on the side are **Presets**. There are many presets provided by Shotcut in the **Stock** category, and you can add your own that will be shown under the **Custom** category. The **+** and **-** buttons at the bottom of the list are used to add and delete custom presets.
NOTE: It is not necessary to choose a preset! Shotcut comes with default values that creates a high quality H.264 MP4 file at a reasonable size and speed that automatically adapts to the resolution, frame rate, and visual complexity of your project. Clicking **Reset** or the preset named **Default** restores the panel to its default values.
- The **From** field lets you choose what to export. You can export more than just the timeline. In fact, you do not even need to use the Timeline in Shotcut! Here


are the options:

- **Timeline:** This is the default if you have added anything to the Timeline.
- **Current Playlist Bin:** Exports the **Playlist** as a single, sequential file. Since a playlist has optional bins (categories), it depends on which bin is currently chosen. If you are not using bins then it is the entire playlist.
- **Each Playlist Bin Item:** Exports a separate file for each item in the playlist or, optionally, the current bin if using playlist bins.
- **Source:** You can export a single clip that has been trimmed and filtered in the **Source** player.
- **Marker...:** Exports a range **Marker**, which is a marker with a duration.
- **Use hardware encoder** is option to use the hardware-accelerated encoder in your NVIDIA or AMD GPU or CPU (Intel Quick Sync on most processors since 2012, but it may depend on your motherboard as well). You will need to tell Shotcut which you have, but if it has not been configured Shotcut tries to automatically detect it when you click the checkbox to turn it on. Pay attention to the status message area just below the player controls for several seconds after clicking the checkbox to see what it reports. This feature currently only supports the H.264 (aka AVC) and HEVC (aka H.265) codecs.
- **Configure...** opens a dialog for the configuration of the hardware encoder. This is useful in case there was a problem with automatic detection, you want to see what automatic detection found, or you want to change it if you have more than one available on your system.

Next, you might need to understand some basic, fundamental concepts of video and audio compression to understand why it takes long or why the file size is not what you expected. Please see this article

<https://vimeo.com/blog/post/video-compression-basics/>

Controls whether Shotcut uses **hardware-accelerated video encoding** instead of software (CPU-based) encoding during export.



Hardware encoding can significantly **reduce export time** and **CPU usage**, at the cost of some flexibility and, in certain cases, slightly reduced compression efficiency.

What hardware encoding is

Hardware encoding uses dedicated circuitry inside the **GPU** (Graphics Processing Unit) or integrated graphics processor to compress video.

This is different from software encoding, which uses the CPU only.

Key characteristics:

- Faster exports
- Lower CPU load
- Fixed-function encoders optimized for speed
- Limited to codecs and formats supported by the hardware

What a GPU is (in this context)

A **GPU** is a processor specialized for graphics and parallel workloads.

Modern GPUs often include **dedicated video encoding blocks**, separate from 3D

rendering units.

These blocks are designed to encode formats such as H.264 or HEVC efficiently and with low power consumption.

How Shotcut detects hardware encoders

Shotcut **automatically scans the system** and lists all hardware encoders it can access on the current operating system.


What appears depends on:

- Your GPU vendor (Intel, AMD, NVIDIA)
- Installed drivers
- Operating system support
- Enabled APIs (VAAPI, NVENC, etc.)

Unchecked options mean:

- The encoder exists in Shotcut
- But is **not available on your system**

Common hardware encoder types explained



VAAPI (Linux)

VAAPI (Video Acceleration API) is commonly used on Linux systems.

Examples:

- **h264_vaapi** — H.264 encoding using VA-API
- **hevc_vaapi** — HEVC (H.265) encoding using VA-API
- **av1_vaapi** — AV1 encoding using VA-API (if supported)

Typically available on:

- Intel GPUs
- AMD GPUs
- Some newer integrated graphics

NVENC (NVIDIA)

NVENC is NVIDIA's proprietary hardware encoder.

Examples:

- **h264_nvenc**
- **hevc_nvenc**
- **av1_nvenc** (newer GPUs only)

These appear only if:

- An NVIDIA GPU is present
- Proper NVIDIA drivers are installed

If unchecked, it usually means no NVIDIA GPU is detected.

Codec prefixes explained

- **h264** — Widely compatible video codec
- **hevc** (H.265) — More efficient compression, slower to decode
- **av1** — Very efficient, newest standard, limited hardware support

Suffixes:

- **_vaapi** — Linux VA-API backend
- **_nvenc** — NVIDIA hardware encoder

What Shotcut uses hardware encoding for


When enabled:

- Video compression is offloaded to the GPU
- Audio encoding remains CPU-based
- Filters and effects may still be processed on the CPU or GPU depending on settings

Export speed and performance gains

Typical improvements:

- **2× to 10× faster exports**, depending on:
 - Resolution
 - Codec
 - GPU capabilities
- Dramatically reduced CPU usage
- Lower system power consumption



Actual gains vary by hardware and project complexity.

Quality considerations

Hardware encoders prioritize **speed and consistency**.

Compared to software encoders:

- Compression efficiency may be slightly lower
- File sizes may be larger at the same quality target
- Advanced tuning options are limited

For most practical use cases, quality remains very good.

Important nuance:

Hardware encoding is usually less compression-efficient, not lower quality. The visual quality difference is often small or invisible; the difference is usually in bitrate efficiency, not obvious artifacts.

Recommended usage

- Enable hardware encoding for:
 - Long exports
 - High-resolution video
 - Preview or draft renders
 - Laptop or low-power systems
- Prefer software encoding for:
 - Maximum compression efficiency


- Archival-quality exports
- Advanced codec tuning

Limitations

- Available encoders depend on hardware and drivers
- Fewer fine-grained quality controls
- Some export presets may not support hardware encoders
- Not all codecs are available on all GPUs

Export presets allow you to **save a complete set of export settings** and reuse them consistently across projects.

They are especially useful when you repeatedly export to the same platform, resolution, or delivery standard.



A preset stores format, codecs, resolution, frame rate, quality, and advanced options exactly as configured.

What export presets are

Export presets are:

- A snapshot of export parameters
- Applied instantly with one click
- Non-destructive and reusable
- Stored under the **Custom** category

Export presets are not:

- Project templates
- Guaranteed-safe configurations
- Automatically compatible with all codecs or hardware encoders

Using presets assumes you understand the implications of the selected settings.


How to create a custom export preset

1. Open the Export panel

Click **Export** in the top toolbar.

2. Switch to Advanced mode

Click **Advanced** at the bottom of the Export panel.



This exposes all export parameters, including codec selection and low-level options.

3. Configure export settings

Adjust settings across the available tabs. Common examples:

Format / Codec

- Container: mp4 , mov , mkv , etc.
- Video codec: H.264, HEVC, AV1, etc.
- Audio codec: AAC, Opus, PCM, etc.

Video

- Resolution (e.g. 1920×1080, 3840×2160)
- Aspect ratio (e.g. 16:9)
- Frame rate (e.g. 24, 25, 30, 60 fps)
- Scan mode (progressive/interlaced)

Quality / Rate control

- Quality percentage
- Constant quality vs bitrate-based modes

Not all codecs support the same quality or rate-control options.

4. Save the preset

Click the **+** (**plus**) button at the bottom of the preset list.

A text editing dialog opens where you can make additional changes. Often, at this point, you want to remove lines that you do not want to be included. And often, those are things you want to change along with the project settings such as resolution, aspect ratio, frame rate, colorspace, and number of audio channels. The preset is really a list of thing to change from the defaults. The defaults are all the things you see in Advanced when you start Shotcut and do not change anything except for the things managed by Settings: Video Mode, Processing Mode, and Audio Channels. You can also get back to defaults by clicking **Reset** or choosing the **Default** preset.


5. Name the preset

Provide a **descriptive name** that reflects its purpose, for example:

- YouTube Higher Quality H.264
- Archive ProRes 422
- Preview Low Quality
- Mobile 720p Fast Export

Click **OK** to save.

The preset appears under **Custom**.




Using custom presets

- Select the preset from the **Custom** category to apply all saved settings.
- The preset remains selected for future exports until changed.

- You can modify settings after selecting a preset, but changes are **not saved** unless you overwrite or re-save it.

Deleting presets

- Select a custom preset
- Click the – (**minus**) button
- The preset is removed permanently



This does not affect exported files.

Example preset scenarios

Fast preview export

- Container: MP4
- Codec: H.264 (hardware encoding enabled)
- Resolution: 1280×720
- Quality: 45–55%

Used for:

- Quick reviews
- Draft uploads
- Testing edits

Online platform delivery

- Container: MP4
- Codec: H.264

- Resolution: 1920×1080
- Frame rate: Match source
- Quality: 60–70%

Used for:

- YouTube, Vimeo, general web delivery

Archival / master export

- Container: MOV or MKV
- Codec: ProRes or high-quality H.264
- Resolution: Match source
- Quality: High or lossless

Used for:

- Long-term storage
- Re-editing later
- Maximum quality retention

Important warnings

Codec compatibility

Not all combinations are valid.

Examples of common pitfalls:

- Choosing a codec not supported by the selected container
- Enabling hardware encoding with unsupported formats
- Mixing incompatible pixel formats or profiles

These issues can:

- Cause export failure
- Produce unplayable files
- Result in silent audio or missing video

Hardware encoding interaction

Presets may include hardware encoders. However, the **Use hardware encoder** checkbox will still affect which codec is used unless you choose a hardware codec that Shotcut does not manage, for example ProRes.

If the preset is used on a system where the hardware encoder is unavailable:

- The export may fail
- Shotcut may fall back to software encoding
- Results may differ from expectation

Knowledge level required

Creating custom presets ranges from **basic to advanced** usage.


- Simple presets (resolution + quality) are generally safe
- Advanced presets require understanding:
 - Codecs and containers
 - Bitrate vs quality modes
 - Hardware encoding limitations

Use caution when copying presets between systems.

Best practices

- Name presets clearly and specifically
- Test presets on short clips before long exports
- Avoid overloading a single preset with too many assumptions
- Keep separate presets for preview, delivery, and archive
- Presets are saved as files in the `presets > encode` folder in the app data directory (**Settings > App Data Directory > Show**). You can share these with others. You must restart Shotcut after manually adding preset files in order for Shotcut to see them.

Shotcut's export panel has an advanced mode that you can access by clicking the **Advanced** button at the bottom of the panel. By default the next time you restart Shotcut, the advanced mode is turned off again. However, you can pin it to the on or enabled state by clicking the checkbox next to **Advanced**.



You should also be familiar with [Export Basics](#).

Video

Codec

Audio

Other

See MLT's [documentation](#) for information about the syntax and parameters. MLT's [consumer properties](#) also apply here. Most essential parameters are written by the presets and various form fields of **Export > Advanced** but you can override them here (must use the correct name, of course). Some parameters are appended to the generated ones. For example, you can supply a `x265-params`, and the **Codec** tab will prepend values needed for rate control. Prepending lets you specify overrides. You might also find [FFmpeg's codecs and formats documentation](#) more useful here as it is more organized. Please keep in mind that Shotcut is not a `ffmpeg` command line front end except in the rare cases of **Properties > Convert**



and Reverse and proxy media generation. Thus, there are many things you cannot copy from ffmpeg examples are possible here. Only options from FFmpeg that affect the encoder or muxer.


Custom Export Presets

You can save your own custom preset by clicking the **+** button at the bottom of the Export panel.

To add a new Custom Export Preset:


- Set your preferred parameters in the Video, Codec, Audio and Other tabs.
- Click the **+** button.
- Give a name to your preset and click OK





To delete a Custom Export Preset:

- Select the custom preset you want to delete
- Click the  button at the bottom of the Export panel
- Click OK to confirm



See also: [Export Basics](#)



Certain formats such as MP4 and MP3 (and many others but not all) support embedded metadata attributes such as Artist and Title. See the [MultimediaWiki FFmpeg Metadata](#) page for more information about which formats supports which attributes.

In Shotcut, you can embed these attributes into your exports by using the **Other** tab in the **Advanced** mode of the **Export** panel. You can add one or more of the following lines with the value you want on the right side of the = sign:

- meta.attr.artist.markup=
- meta.attr.comment.markup=
- meta.attr.copyright.markup=
- meta.attr.date.markup=
- meta.attr.description.markup=
- meta.attr.genre.markup=
- meta.attr.title.markup=

For audio language, you can set along to an ISO 639 language code. For example:

```
aLang=eng
```




Introduction

It is fairly easy to run an export on the command line because Shotcut essentially does the same thing. For example, do an export in the UI and then **View > Application > Log...** and scroll to the end. You will see the command line printed; however, the file name is fully encoded to workaround potential character set problems.

Quick Start

The general basic form is:

```
LC_ALL=C melt project.mlt -consumer avformat:export.mp4
```

Of course, you substitute your own file names possibly with a path prefix for `project.mlt` and `export.mlt`. `LC_ALL=C` sets that environment variable to mitigate problems with numeric strings in some locales. That command line generates a H.264/AAC MP4 using software encoder x264 with `melt` & FFmpeg's default encoding options.

Encoding Options

You can add some encoding options at the end of the command line like

```
LC_ALL=C melt project.mlt -consumer avformat:export.mp4 movflags=+faststart vcodec=
```

Notice that each option is not preceded by one or two dashes. Rather, the `-consumer` specifies a consumer "object", and the options that follow until the next dash add properties to that object. There are an extreme number of `avformat` (named after FFmpeg's `libavformat` library) options and values! Moreover, they depend on operating system and which hardware encoders are available. There is a rough [reference of them](#) on the MLT web site. If you know or find some `ffmpeg` command line examples, you can adapt **some** of its options but **not** in the same order or syntax! Melt only accepts the encoder and muxer options here, but only **some** of `ffmpeg`'s options are for the encoder and muxer. Moreover, `ffmpeg` command lines do not usually clearly organize the options. So, you need to be a fairly advanced user. See also [Advanced Export](#).

You can use Shotcut's UI to prepare these option=value pairs. Basically, setup and export, and click the  button at the bottom of the **Export > Presets** list. That opens a dialog with all of the option=value pairs in a text box. You can copy and paste these, but you need to convert the new lines into spaces.

Project vs. Export Job XML

The above command line uses a `-consumer` option, but if you looked at Shotcut's log you see it does not. The difference is that Shotcut adds the consumer and its options into the MLT XML. That is the main difference between them (except for a clip-only project). Basically, Shotcut adds a fourth line to the project XML:

```
<consumer mlt_service="avformat" movflags="+faststart"... />
```

So, an easier way to do all of the above is:

1. Open a project in Shotcut
2. Click **Export**
3. choose export preset or configure export options
4. Click **Export > File**
5. Right-click the new job in **Jobs** and choose **Stop**
6. Right-click the job and choose **View XML**
7. In the dialog that opens, choose **Save**
8. Choose a location and filename, and click **Save**
9. At the command line run `LC_ALL=C melt job.xml`

Platform Notes

- With Linux Flatpak, run melt using

```
LC_ALL=C flatpak run --command=melt org.shotcut.Shotcut
```
- With Linux Snap, run melt using

```
LC_ALL=C shotcut.melt
```
- With Linux AppImage, I think the only way is to **mount** the image and then use the reported `/tmp` directory followed by `/usr/bin/melt`, for example

```
LC_ALL=C /tmp/.mount_shotcuD5hIPg/usr/bin/melt
```
- With Linux portable, run melt using

```
LC_ALL=C path/to/Shotcut.app/melt
```
- With macOS, simply run


```
LC_ALL=C /Applications/Shotcut.app/Contents/MacOS/melt
```
- On Windows (outside of a package system like msys2), `cd` to the directory where Shotcut is installed and run `melt` in `CMD.EXE` or `./melt` in a shell. Alternatively, provide the full path to the `melt.exe`:

```
LC_ALL=C 'C:\Program Files\Shotcut\melt.exe'
```

💡 Pro Tip: Shotcut includes a few other executables you can run in the same manner:

- `ffmpeg`
- `ffplay`
- `ffprobe`
- `glaxnimate`

When an audio/video clip is selected, the **Properties** panel shows the audio and video properties.



- The heading shows the **title** of the clip. The default is the file name, but you can click into it or press **F2** to rename the clip. Renaming here only affects how it appears in the current project. It does not change the file name or its name in other or future projects.
- **Comments:** is where you can add your own description, keywords, or tags. These are searchable in the **Playlist** panel for playlist clips only. Remember, in Shotcut, **Playlist** and **Timeline** are separate containers for clips. So, you need to open (double-click) a playlist item to see that clip objects' comments that are searchable.
- The **Speed** setting allows the playback speed of the clip to be increased or decreased. There is a down arrow button that has some commonly used speed presets such as 0.5 (slower) and 2.0 (faster).
- **Pitch Compensation** affects how the audio sounds for speed-changed clips. When it is off, the frequency or pitch changes according to the speed. When it is on, Shotcut tries to maintain the correct pitch, but it can only do so much. It works ok for small speed changes but for large speed changes you probably need to mute the audio.
- **Duration** only overrides the source duration, not the duration of the playlist item or timeline clip. If you are using Time Remap or Speed filters to slow down the video you might need to increase this duration because these filters do not have


access to the source clip to make the changes. **Speed** here within **Properties** can and does automatically adjust **Duration**.

- **Timeline duration** is always disabled because you cannot adjust the duration of a clip numerically at this time. You need to use the actions available in the Timeline and playback controls.

Video Tab

- **Track** lets you change which of multiple embedded streams or tracks within a multi-track file. Or you can use it to disable the video on this selected clip.
- **Aspect ratio** lets you override the wideness or tallness of the video in case it was mis-detected or for an effect. For example, you can force a 4:3 video to be stretched to 16:9 to make it fill the screen.
- **Scan mode** lets you override whether the clip is progressive or interlace in case it was mis-detected.
- **Color range** lets you override whether the clip uses the full numeric range of values or the more common (since MPEG-1) limited range. Normally, you should not need to change this and only change it if you know for sure. If this shows **Broadcast Limited** but the **Video Waveform** scope shows values below 0 or above 100, then your video is most likely actually full range, and you will experience clipping if not corrected.
- **Rotation** is another override. Sometimes a video is shot upside-down or with the camera turned, but the camera did not specify the correct orientation in the video file.

Audio Tab




- **Track** lets you change which of multiple embedded streams or tracks within a multi-track file. Or you can use it to disable (mute) audio on this selected clip.
- The **Sync** slider appears if the clip also has video. Use this to fix the audio/video synchronization for the clip. However, before doing so, you should also verify your Shotcut's system level playback audio/video synchronization in [Settings > Player > Synchronization](#)

Metadata Tab

This shows non-editable metadata that is embedded in the file. You can enter your own metadata in **Comments**, but that only saves with the Shotcut project. Shotcut does not modify source media assets.

Button Bar



Menu

The first button in the Properties button bar is a menu button:

- **Reset**
- **Show In Files**
- **Show In Folder**
- **Copy Full File Path**
- **More Information**
- **Start Integrity Check Job**
- **Convert to Edit-friendly**
- **Extract Sub-clip**
- **Extract Subtitles**
- **Set Creation Time**
- **Set Equirectangular**
- **Measure Video Quality**
- **Export GPX**
- **View Bitrate**

Convert

This is the same as the **Convert to Edit-friendly** menu item above.


Reverse

Proxy

Open With and Reload

- **Open With** shows a menu of other tools with which you can open this video or audio-only file.
- Choose **System Default** to let the operating system decide which tool to use—the same as if you double-clicked the file in its file manager.
- Choose **Other** to add an application that is on your computer.
- Choose **Remove** to remove something you added with **Other**.
- You can think of this as “Edit With” especially useful for images and audio files.
- There is a file watcher upon opening with another tool as long as selection (Properties) does not change. If it does, you can use **Reload**. This does not yet reload—whether manual or automatic—every clip object based on this file.
- This is not available in the Linux Flatpak.
- This was added in Shotcut version 25.10

Shotcut can not “pass through” or copy codecs without re-encoding. The closest thing available to that is Properties > menu > Extract sub-clip, which uses the in and out points of a trimmed clip.




Extract sub-clip will find the key frame closest to the in/out points and copy the codecs into a new container without re-encoding. Any filters applied to the clip will not be applied to the sub-clip because that would require re-encoding.

This feature has some limitations and may not always provide the expected result:

- The extraction will round to the nearest codec keyframe - which could be more than a second away from the selected in/out points
- The extraction tries to pack the media into the same container format as the original. But it may not always be able to get the packing or timing to match the original

When an image is selected, the **Properties** panel shows information and options.




Fields

- The heading shows the **title** of the clip. The default is the file name, but you can click into it or press **F2** to rename the clip. Renaming here only affects how it appears in the current project. It does not change the file name or its name in other or future projects. There is a tooltip for this field that shows the full path and name of the file.
- **Comments:** is where you can add your own description, keywords, or tags. These are searchable in the **Playlist** panel for playlist clips only. Remember, in Shotcut, **Playlist** and **Timeline** are separate containers for clips. So, you need to open (double-click) a playlist item to see that clip object's comments that are searchable.
- **Resolution** shows the source image dimensions.
- **Duration** shows the amount of time this image will appear in either the **Playlist** or **Timeline**. You can only edit the duration of the image in the **Source** viewer—whether a playlist item or not. You can also adjust the duration by changing the out point in the **Source** viewer. The only way to change the duration of an image on the **Timeline** is by dragging its right edge (trimming). The format of the time value is controlled by **Settings > Time Format**.
- When you open an image or drop one into **Playlist** or **Timeline**, it gets a default duration. The default is 4 seconds. Click **Set Default** to make the current duration displayed become the new default. That does not change the duration of any images that are already in the **Playlist** or **Timeline**—only newly added ones.
- **Pixel aspect ratio** lets you change the aspect ratio in case your image was exported directly from a video or tool such that it is not square pixels. Images normally default to assume 1:1, meaning square pixels.

- Click the **Image sequence** checkbox if this image is a member of an image sequence and you want to make Shotcut play the sequence instead of the single image for many frames. An image sequence is a collection of image files that represent frames of video saved with a counter number in the file name.
- The frame rate of an image sequence normally follows the current **Video Mode**. However, you can increase **Repeat** to make it slower in integer multiples by repeating each image in the sequence. If you need something slower or more precision, you need to convert the image sequence to a video. You can easily do that in Shotcut. Simply open the image in the **Source** player, check **Image sequence** if not already, and choose **Export > From > Source**. You might want to use an `intermediate` or `lossless` export preset for that to preserve the most quality.


Button Bar



Menu

The first button in the Properties button bar is a menu button:

- **Show In Files** opens the integrated Shotcut **Files** panel to the file's folder and selects the file.
- **Show In Folder** opens your system's default file manager to the file's folder and—on Windows and macOS only—also selects the file.
- **Copy Full File Path** copies the full path and file name to the system clipboard.
- **Set Creation Time** opens a dialog that shows the detected capture date and time, if found, or the file's creation time where you can change it. This is useful in the **Playlist** where you can re-order items by date.



- **Reset** reverts any changes you have made to aspect ratio or image sequence.

Proxy

This opens a menu:

- **Make Proxy**
- **Delete Proxy**
- **Disable Proxy**
- **Copy Hash Code**


See [Settings > Proxy Editing](#) for more information.

Open With and Reload

- **Open With** shows a menu of other tools with which you can open this video or audio-only file.
- Choose **System Default** to let the operating system decide which tool to use—the same as if you double-clicked the file in its file manager.
- Choose **Other** to add an application that is on your computer.
- Choose **Remove** to remove something you added with **Other**.
- You can think of this as “Edit With”.
- There is a file watcher upon opening with another tool as long as selection (Properties) does not change. If it does, you can use **Reload**. This does not yet reload—whether manual or automatic—every clip object based on this file.
- This is not available in the Linux Flatpak.
- This was added in Shotcut version 25.10



The **Color** can be used to create a color clip. After choosing the menu option, the dialog will appear.




Click on the “Color...” button to choose a color. After the color is chosen, you can click “Open” to open the color clip in the source player. Or, you can click “Add To Timeline” to add the color clip directly to the timeline.

Use this to add text to your **Playlist** or **Timeline** as a standalone clip object with its own duration and effects (filters) that affect only the text. Otherwise, when using a text filter on a video or image clip, any additional filters you add also affect the image and not only the text!

This exists in several places:

- The main menu at **File > New > Text**
- The main menu at **File > Open Other...**
- The 2nd icon in the main toolbar

- The 5th icon in the **Timeline** toolbar with the same icon



1. Enter the text you want to display in the main, big, empty text box. Change the **Background color** only if you want to simply display the text on top of a solid color.
2. Use **Preset** only if you want to save your text snippet along with its **Background color** to use again at a later time.
3. Choose a particular text generator:
 - **Simple** is probably the most common and obvious choice. It can display metadata from variables, but all of the text must have the same font, size, and color, including outline.
 - **Rich** lets you change the fonts, colors, and do other formatting within a single text clip; however, outline requires the addition of the separate **Outline** filter.
 - **Typewriter** animates the text character-by-character or word-by-word.

4. Click **Add To Timeline** to immediately add the text to the Timeline at the current playhead. If there is already a clip on the current track at this position it adds a new track holding the new text object. Otherwise, click **Open** to open the text clip in the **Source** player where you can preview it before adding it to the **Playlist** or **Timeline**.

NOTE: The **Add to Timeline** does not show when you choose to trigger this from the **Timeline** toolbar. Instead, **OK** appears and clicking it adds it to the **Timeline**.


Use the Drawing/Animation Generator to open a simple drawing program included with Shotcut to draw something *new* on top of a video or a solid color background. The program is named **Glaxnimate**, and it creates a structured (vector) drawing rather than painting pixels. Structured means it consists of shapes and lines that you can adjust at any time. You can also animate them. This includes text, emojis, and raster images as well.

If you want to include a vector drawing or animation that you have downloaded you can import it like any other piece of media. In addition to its own `rawr` format, Glaxnimate also supports SVG, Lottie, Rive, and After Effects animations. However, these are limited and can have interoperability problems. Sometimes you can fix the issues in Glaxnimate.

Getting Started

You can trigger this from several places:

- The main menu at **File > New > Drawing/Animation**
 - The main menu at **File > Open Other...**
 - The 2nd icon in the main toolbar
- 
- The 5th icon in the **Timeline** toolbar with the same icon



Preset is not very useful here but included because all of the generators have a preset system.


1. It is a good idea to set the **Duration** you want ahead of time, but you can change it later from **Properties** if you need to.

2. Click **Add To Timeline** to immediately add the text to the **Timeline** at the current playhead. If there is already a clip on the current track at this position it adds a new track holding the new text object. Otherwise, click **Open** to open the text clip in the **Source** player where you can preview it before adding it to the **Playlist** or **Timeline**.

NOTE: The **Add to Timeline** does not show when you choose to trigger this from the **Timeline** toolbar. Instead, **OK** appears and clicking it adds it to the **Timeline**.

3. If you did not use **New Project > Start** to start this project, a file save dialog opens for you to choose a location where to save the drawing/animation, which is a separate file that ends with the `.rawr` filename extension. Otherwise, when using a proper project folder, this file is generated automatically in your project folder.
4. Glaxnimate opens. See the [Glaxnimate Web site](#) for more help with using it.

Properties




Click **Edit** to launch Glaxnimate to edit your animation. If your drawing is in the Timeline with an image or video clip on a track below a drawing with a transparent background then you will see a preview of the Shotcut video as the canvas background in Glaxnimate. Also, as long as this clip remains selected, any changes you **save** in Glaxnimate are automatically detected and updated in Shotcut. However, if you change selection or do something that affects it, then click **Reload** to manually refresh the appearance of the drawing inside Shotcut.

Also, like most clips, you can change the clip name in the header and enter comments.




Choosing **Image/Video from HTML** from the menu **File > New** or a toolbar **New Generator** button shows presets/templates in **Properties**:




Choosing a preset that is a template shows a UI similar to below but also with some fields such as **Line 1**, optionally up to **Line 3**. Enter your custom text. You can use the hint in the template name to click either **Generate Image** or **Generate Video**.

Choosing **(custom)** shows a form like this:



This can be used to create something like the text in this image:



Expand each section to show a show text box where you can enter or paste each block of code.

Notes

- This requires Google Chrome or Chromium.
- It was added in version 25.10
- Please see the stock **Presets** for examples.
- It defaults to a transparent background!
- The stock presets also demonstrate a template facility for up to 3 lines of text. You can make your own templates by using %1 , %2 , or %3 as placeholders in the Body field.
- Expand each section to reveal a text input box.
- **Body** means the `<body>` tag. Shotcut generates the full HTML. (You can include `<script>` blocks inside the body if you find the need to.)
- It is designed to make it easy to copy from codepen.io, but Shotcut does not include pre-preprocessors for things like SCSS or TypeScript. Therefore, in codepen.io click the  button in the top right corner of the edit block to choose **View Compiled** before copying.
- This does not support WebGL or embedded video.
- **Generate Video** is limited to 15 frame-per-second for performance reasons. A somewhat modern or fast computer and SSD hard drive are recommended.
- Generate automatically opens the result in the **Source** viewer so you can preview it with its HTML for revising if needed. Once you add it to **Playlist** or **Timeline**, **Properties** now reflects the image or video and no longer shows the HTML inputs.
- This is not available in the Linux Flatpak.

See [#Resources](#) filtered to the [#HTML](#) tag to get more ideas.


A Note Regarding Animation and Video

The animation cannot be recorded properly frame-by-frame because the web APIs lack that comprehensively. There is something called Web Animation API that does support that but requires coding the animation a certain way that is limiting and does not work with the vast majority of things on codepen. So, this works like a low performance screen recorder. If you need something with higher frame rate and consistency you need to use a web browser and a screen recorder. However, then you must use a background color and chroma key, which is not nearly as good as true alpha and lacks translucency. So, Generate Video in Shotcut is a compromise. You can slow down the animation in Shotcut's **Properties > Speed** if you want and can even correct some of the timing with a time filter. Or, simply use the image option and you can do some other animation in Shotcut or Glaxnimate.

The **Lissajous** generator creates a Lissajous clip.

A Lissajous curve is a pattern made by combining two wave motions: one side to side, and one up and down.

After selecting, this dialog box will appear.




You can adjust the dials to set the curve's properties (values range from 0.00 to 1.00):

- “X Ratio” controls the horizontal wave frequency.
- “Y Ratio” controls the vertical wave frequency.

After setting your desired curve, you can either click:

- “Open” to open the lissajous clip in the source player. Where you can adjust and preview the curve before adding to timeline.
- “Add To Timeline” to add lissajous clip directly to the timeline.
- “Save” to “Preset”.

If you selected “Open”, drag from project player to add clip to timeline.




You can continue adjusting the curve after adding it to the timeline by selecting the clip – this will reopen the dialog box in Properties.



This generates a **synthetic noise video with matching noise audio**, similar to what was historically seen on analog televisions tuned to a channel with no broadcast signal.

The generator produces a predefined clip consisting of visual noise, also known as “**snow**,” and broadband noise sound.



This is a generator, not a filter. The result is a self-contained clip that can be placed on the timeline from the preview player or directly placed on the timeline from the generator.

What TV noise is (historical context)

In the analog television era, a TV not receiving a valid broadcast signal would display random noise:

- **Visual noise (“snow”)**
Caused by thermal noise and random electromagnetic interference picked up by the tuner.
- **Audio noise**
Broadband hiss resulting from the absence of a modulated audio carrier.

This noise was not generated intentionally by broadcasters; it was the raw output of the receiver when no signal was present.

Viewers commonly saw and heard this:

- After broadcast hours
- On unused channels
- When reception was poor
- When an antenna was disconnected

Why this noise existed

Analog TV receivers amplified whatever signal they could detect. When no structured signal was present, they amplified random electrical noise, making it visible and audible.

This made “static” a familiar indicator of:

- No transmission
- Signal loss
- Tuning between channels

Generator behavior in Shotcut

- Produces a fixed-duration clip (video + audio)
- Video and audio noise are continuous and uncorrelated
- The clip can be trimmed, looped, or extended like any other media
- The generated noise is synthetic and repeatable

Visual characteristics

- High-frequency black-and-white noise
- Fine-grain texture
- May appear slightly patterned or structured at close inspection

Note:

Adding a very small **Blur: Gaussian** ($\approx 1\%$) removes visible patterning and produces a more natural analog-style noise appearance.

Audio characteristics

- Broadband noise (hiss-like)
- No pitch, rhythm, or tonal structure
- Constant level over time

This closely resembles analog TV static rather than digital white noise used in audio testing.

Contemporary use

Today, noise generators are used almost exclusively for historical, symbolic, or stylistic purposes.

Common modern uses include:

- Intro or outro visuals referencing analog television
- Transitions indicating loss of signal or interruption

- Documentary or archival storytelling
- Retro or broadcast-era aesthetics
- Simulating “dead air” or system failure

They are not used for technical testing in modern digital workflows.

Recommended use cases

- Creating a “no signal” or “broadcast interruption” effect
- Evoking analog-era television
- Background texture for glitch or degradation effects
- Sound design for interruption or uncertainty

Limitations

- Not true analog noise (digital approximation)
- No parameter control within the generator itself
- Fixed initial duration
- Not suitable for audio calibration or measurement



This generates an **animated procedural plasma pattern** based on mathematical interference functions.

The animation is continuous and non-repeating, and all controls interact with each other.

This type of imagery originates from early computer graphics and demoscene visuals, where smooth color gradients and flowing motion were generated mathematically rather than from filmed footage.



In Shotcut, Plasma is a generator, not a filter.

All parameters are adjusted in the **Properties** tab of the generated clip.

Behavior Note:

Plasma must be evaluated during playback.

A single frame is misleading and does not represent the effect.

Small numeric changes can result in qualitatively different motion and structure.

How the Plasma generator actually behaves

The Plasma generator combines several evolving wave fields.

The controls do not correspond to simple axes or isolated properties.

Instead, they influence:

- Pattern topology (rounded vs stretched)
- Temporal evolution (static vs pulsing)

- Directional drift
- Rate of internal phase change

Parameters

All parameters range from 0.00 to 1.00 and are evaluated continuously over time.

Speed 1

Controls the **overall temporal evolution** of the plasma field.

- Lower values
Slower internal evolution; patterns feel more stable.
- Higher values
Faster internal changes; increased pulsing and morphing.

Important:

Speed 1 strongly affects whether the plasma appears **static**, slowly drifting, or actively pulsing.

Speed 2

Influences the **spatial structure bias** of the plasma.

- Values below 1.00
Tend to produce elongated or stretched structures.
- Values at or near 1.00
Tend to produce rounder, more isotropic patterns.

This parameter affects **shape character**, not direction alone.

Speed 3

Works in conjunction with Speed 2 to define **pattern geometry**.

- When Speed 2 $\neq$ Speed 3
The plasma exhibits directional stretching.
- When Speed 2 $\approx$ Speed 3
The plasma becomes more symmetrical and rounded.

Key point:

Speed 2 and Speed 3 form a pair. Their relative values matter more than their absolute values.

Speed 4

Controls the **rate of internal phase cycling**.

- Lower values
Colors and structures change slowly.
- Higher values
Faster internal cycling; more visible pulsing.

This parameter contributes to the “alive” feeling of the plasma.

Move 1

Controls **global drift with deformation** along one axis.

- Lower values
Minimal drift.
- Higher values
Slow directional movement combined with compression/expansion.

This is **not** a pure translation.

Move 2


Controls **global drift with deformation** along the second axis.

As with Move 1:

- Movement includes squishing and stretching
- Direction and speed depend on interaction with Speed parameters

Behavior note:

When Speed controls (1, 2, and 3) are set to similar values (with Move controls at high values), the plasma pattern may collapse into a **full-frame color cycling** effect, where structure largely disappears and the image appears as flashing and smoothly changing colors.



Visual characteristics

- Abstract, organic motion

- Strong dependence on parameter interaction
- No fixed orientation or stable geometry
- Behavior changes qualitatively, not linearly

Historical context

Plasma effects were widely used in:

- Early computer demos (1980s - 1990s)
- Demo scene productions
- Software intros and loading screens
- Low-power real-time graphics demonstrations

They showcased mathematical creativity rather than realism.

Contemporary use

Today, plasma generators are used primarily for stylized and abstract visuals:

- Background textures
- Motion graphics layers
- Retro or demo-scene aesthetics
- Abstract transitions
- Visual fillers or placeholders

They are not used for realism or calibration.

Use with other filters

Plasma is often combined with other filters for creative results:

- **Blend Mode**
Overlay plasma on footage for stylized color motion.
- **Opacity**
Subtle background animation behind titles or graphics.
- **Blur**
Further smooth gradients or reduce harsh transitions.
- **Color Grading / Saturation**
Emphasize or mute the vivid colors.
- **Crop / Mask**
Use plasma only in specific regions.

Usage notes and tips

- Always preview **in** motion.
- Small parameter changes can have large visual effects.


- Plasma works best as a supporting visual, not a focal subject.

Limitations

- No deterministic control over exact shapes
- Highly sensitive to parameter interaction
- Not suitable for precise motion design
- Single-frame inspection is misleading

Generates **standardized test images** known as *color bars*, historically used to calibrate and verify video signal quality.

Color bars were not decorative graphics; they were **reference patterns** designed to check color accuracy, brightness, contrast, saturation, and signal integrity.



In Shotcut, the **Color Bars** generator produces a static image that can be added to the timeline or previewed in the player.

Because this is a generator and not a filter, the color bar type is adjusted in the **Properties** tab rather than in the Filters panel.

What color bars were (historical context)

Before digital video and HDMI, television signals were **analog** and prone to distortion during transmission and recording.

Color bars were used to:

- Calibrate television receivers
- Align broadcast equipment
- Verify tape recorders, cameras, and monitors
- Diagnose signal problems (color shift, phase errors, clipping)

When a TV channel was off-air, during maintenance, or before program start, viewers often saw color bars instead of regular programming.

Why color bars mattered

Color bars allowed technicians to adjust:

- Brightness (black level)
- Contrast (white level)
- Color saturation
- Hue / phase
- Channel balance

Because everyone used the same reference patterns, equipment could be aligned consistently across studios, broadcasters, and countries.

Contemporary use

Although originally designed for calibration and broadcast testing, color bars are now primarily used as a historical or stylistic visual reference.

In modern workflows, they are commonly used for:

- Intro or outro segments referencing broadcast-era television
- Archival, documentary, or educational content
- Intentional retro or institutional aesthetics
- Visual cues indicating “technical context” or system boundaries

In these cases, color bars function as a symbolic visual language, not as a calibration tool.

Behavior

- Produces a single still image
- No animation or time-based behavior
- The image can be:
 - Added directly to the timeline
 - Previewed in the player and dragged later
- Changing the **Type** updates the generated pattern

Type


Selects which standard color bar pattern is generated.

Each option corresponds to a real-world broadcast standard.

100% PAL color bars

Full-amplitude PAL color bars.


- Uses maximum color saturation
- Often used for equipment stress testing
- Less common in broadcast due to risk of clipping



100% PAL color bars with red

Same as standard 100% PAL bars, with an additional red reference bar.


- Used for more precise color phase and chroma checks
- Helpful when aligning PAL-specific equipment



95% BBC PAL color bars

A BBC-specific PAL variant.

- Slightly reduced amplitude compared to 100%
- Designed to stay within safer broadcast limits
- Widely used by the BBC and UK broadcasters




75% EBU color bars

One of the most widely used European standards.

- Reduced color saturation (75%)
- Minimizes signal clipping
- Standardized by the European Broadcasting Union (EBU)

This became a practical broadcast default across Europe.



SMPTE color bars

Standard used primarily in North America and Japan.

- Defined by the Society of Motion Picture and Television Engineers
- Includes additional reference areas for black and white levels
- Still widely recognized today




Philips PM5544

A complex, iconic test pattern generated by Philips equipment.

- Combines color bars with geometry, grayscale, and alignment markers
- Used extensively in European broadcasting (PAL/SECAM)
- Designed for full-system calibration, not just color


Often remembered as a classic “TV test card”.



FuBK

German broadcast test pattern.


- Used by ARD/ZDF and other German networks
- Includes color bars, grayscale, and resolution references
- Designed for studio and transmitter alignment



Simplified FuBK

A reduced version of the FuBK pattern.

- Retains essential color and grayscale references
- Easier to read on consumer displays
- Less detailed than the full FuBK pattern



Visual characteristics

- Static, high-contrast image
- Precisely defined colors and bars
- No artistic variation

Recommended use cases

- Intro or outro visuals referencing broadcast-era television
- Intentional retro or institutional aesthetics
- Creating material for archival or educational content
- Emulating broadcast-era aesthetics
- Teaching video standards and history

Limitations

- No animation or motion
- Not intended for creative imagery
- Does not replace modern color management tools
- Interpretation depends on display calibration

Acronyms Reference

Summary

Label	What the letters mean
PAL	Phase Alternating Line
BBC	British Broadcasting Corporation
EBU	European Broadcasting Union
SMPTE	Society of Motion Picture and Television Engineers
PM5544	Philips test generator model 5544
FuBK	Funkübertragungs-Betriebs-Kontrolle
SECAM	Séquentiel Couleur A Mémoire

PAL

PAL = Phase Alternating Line

- Analog color TV system developed in Germany
- Fixes NTSC hue errors by alternating color phase every line

100% PAL color bars

- Color bars encoded using the PAL system
- 100% = full chroma amplitude (maximum color saturation)

100% PAL color bars with red

- Same as above
- Explicit red field or red-dominant bar added
- Used to isolate the R-Y (red difference) component

BBC

BBC = British Broadcasting Corporation

95% BBC PAL color bars

- PAL color bars defined by BBC engineering standards
- 95% = chroma limited to 95%
- Slightly below maximum to stay within UK broadcast safety margins

EBU

EBU = European Broadcasting Union

- Organization that defined common technical standards across Europe

75% EBU color bars

- Color bars standardized by the EBU
- 75% = reduced chroma level
- Became the European broadcast reference

SMPTE

SMPTE = Society of Motion Picture and Television Engineers

- USA standards body

SMPTE color bars

- Color bars defined by SMPTE
- Designed primarily for NTSC (later adapted to digital)
- Include PLUGE and chroma reference elements

Philips PM5544

PM5544 = Philips Model 5544

- **PM** = Philips Measurement (internal product line)
- **5544** = model number

This is not an acronym, but a specific test signal generator model whose pattern became a de-facto standard.

FuBK

FuBK = Funkübertragungs-Betriebs-Kontrolle

German compound word:

- **Funk** = radio / broadcast
- **Übertragung** = transmission
- **Betriebs** = operational
- **Kontrolle** = control / monitoring


Meaning: *Broadcast Transmission Operational Control*

This refers to a German broadcast test pattern, not a device.

Simplified FuBK

- Same origin as FuBK
- Reduced test elements
- Keeps only the core operational controls

This generates a **continuous test tone** at a fixed frequency and level.
This generator produces **audio only** and is commonly used for testing, calibration, and reference purposes.



Audio Tone is a generator, not a filter.
All controls are adjusted in the **Properties** tab of the generated clip.

What an audio tone is

An audio tone is a pure, steady signal at a specific frequency, traditionally used in audio engineering to:

- Test signal paths
- Calibrate equipment
- Identify problems such as distortion or imbalance
- Provide a known reference signal

Unlike music or noise, a tone contains no variation over time unless modified by other filters.

Parameters

Frequency (20 - 9 999 Hz)

Sets the **pitch** of the generated tone.

- **Low values (20 - 100 Hz)**
Very low frequencies; felt more than heard.
Often used to test bass response or subwoofers.
- **Mid values (≈440 - 1 000 Hz)**
Clearly audible tones; commonly used as reference frequencies.
1000 Hz is a standard test tone in audio engineering.
- **High values (>3 000 Hz)**
Sharp, piercing tones; useful for testing high-frequency response.

Note:

The tone is a single, fixed frequency with no harmonics.

Level (0 to -90 dB)

Sets the output level of the tone.

- **0 dB**
Full level (relative to project headroom).
- **Lower values**
Progressively quieter tones.
- **-90 dB**
Nearly inaudible.

This allows precise control over how loud the tone is relative to other audio.

Auditory characteristics

- Constant pitch
- Constant level
- No modulation, rhythm, or dynamics
- Can become fatiguing at higher frequencies or levels

Recommended use cases

- Testing audio signal chains
- Checking speaker balance and routing
- Verifying channel assignments
- Calibrating monitoring levels
- Creating reference tones for tutorials or demonstrations
- Generating beeps or reference signals for synchronization

Usage notes and tips

- Use moderate levels to avoid listener fatigue.
- High-frequency tones can be uncomfortable; use caution.
- Combine with [Pan](#), [Balance](#), or [Swap Channels](#) to test routing.
- Add filters ([Fade](#), [Gain](#), [Limiter](#)) if you need controlled start/stop behavior.


Limitations

- Produces only a single pure tone
- No built-in modulation or envelopes
- Not suitable for musical content
- Continuous tone unless trimmed or faded



This generates a **countdown or count-up leader** traditionally used in film and television before program start.

This generator recreates classic **film leader / broadcast countdown** visuals, including numbers, timing marks, and optional reference sounds.



Count is a generator, not a filter.

All controls are adjusted in the **Properties** tab of the generated clip.

What a count leader is (historical context)

Before digital playback, films and tapes required manual cueing and synchronization.

Countdown leaders were used to:

- Signal that a program was about to start
- Give projectionists and broadcasters time to cue playback
- Synchronize sound and picture
- Provide a reference point (“start at zero”)

These leaders were commonly seen:

- At the beginning of cinema reels
- Before television programs
- On broadcast tapes and test transmissions

The rotating clock, quadrant background, and large numbers are all part of this legacy.

Generator behavior

- Produces an animated video
- Duration is explicitly defined by the user

- Animation and numbers depend on style and direction
- Must be evaluated during playback to understand timing

Parameters

Direction

Defines whether the count decrements or increments.

- **Down**
Classic countdown toward zero.
- **Up**
Count-up / chronometer behavior.

Style

Controls how time is represented visually.

Seconds

Displays only whole seconds.

- Does not display the starting value
- Displays 0 at the end
- Example (Duration = 10 s, Down):
9 → 8 → ... → 1 → 0

Seconds +1

Displays the starting second value.

- Displays the starting number
- Does not display 0 at the end
- The animated needle continues through the final second

Example (Duration = 10 s, Down):

10 → 9 → ... → 1 (no visible 0)

Frames

Displays frame numbers based on the project's frame rate.

- Uses the project FPS
- The counter spans the entire duration of the generated clip, starting from the total number of frames (countdown) or from zero (count-up).
- Useful for frame-accurate timing

Timecode

Displays time in HH:MM:SS:FF format.

- FF = frames
- Matches standard video timecode representation

Clock

Displays time in `HH:MM:SS` format.

- No frame display
- Acts as a visual timer

Sound

Controls reference audio cues.

2-Pop

Plays a short beep exactly one second before the end.

- Classic synchronization cue
- Widely used in film and broadcast workflows

Silent

No sound.

Frame 0

Plays a short beep every second.

- Useful as a rhythmic timing reference

Background (Clock | None)

Controls the visual style of the background.

- **Clock**
Animated film-leader style background with rotating sweep and quadrant divisions.
- **None**
Plain background with numbers only.

Drop Frame (checkbox)

Controls whether drop-frame timecode is used.

- Relevant only when using **Style > Timecode**
- Applies to NTSC-derived frame rates
- Does not affect visual countdown styles

Duration (time input)

Sets the total length of the generated count.

- Format: HH:MM:SS.mmm
- Determines how many seconds, frames, or units are shown

Visual characteristics

- Large, high-contrast numbers
- Rotating clock-style sweep (when enabled)
- Quadrant background reminiscent of film leaders
- Designed for readability and timing clarity



Contemporary use

Today, count leaders are used primarily as historical or stylistic elements, rather than technical necessities.

Common modern uses include:

- Intro or outro sequences
- Documentary or archival context
- Retro or broadcast-era aesthetics
- Clear visual timing before content starts
- Creative transitions referencing cinema or television history

Use with other filters

Count generators are often combined with filters for creative or practical purposes:

- **Blur**
Soften the image for stylized transitions.
- **Grain / Noise**
Enhance film-era authenticity.
- **Color grading**
Convert to aged tones or colorize.

- **Blend modes**
Overlay the count on footage as a transitional element.
- **Fade In / Fade Out**
Smooth entry into the main content.

Usage notes and tips

- Choose **Seconds +1** if you want the starting value to be visible.
- Use **2-Pop** only when an audio reference is meaningful.
- Match **Duration** carefully to your intended cue point.
- Preview playback to verify timing behavior.


Limitations

- Not intended for modern synchronization workflows
- Visual behavior depends on selected style
- No per-frame visual customization
- Purely representational, not functional timing control



This generates a **synchronized audio beep (“bip”)** and **white visual flash** on a black background at a fixed interval.

Both the sound and the flash occur **at the same time**, repeating according to the selected frequency.



Blip Flash is a generator, not a filter.

Its control is available in the **Properties** tab of the generated clip.

What a blip/flash signal is (historical context)

Synchronized audio and visual cues have long been used in film, television, and broadcast engineering as timing references.

Historically, similar signals were used to:

- Verify audio - video synchronization
- Provide clear temporal markers
- Test signal paths and monitoring chains
- Cue operators during playback or transmission

A well-known example is the “**2-pop**” in film workflows:

a short beep and visual cue placed exactly one second before program start.

The **Blip Flash** generator generalizes this idea into a repeating reference signal.

Generator behavior

- Produces an animated video clip with audio
- Black background with a brief white flash
- A short audio beep occurs exactly in sync with the flash
- The pattern repeats at a constant interval

The clip can be trimmed, looped, or extended like any other media.

Parameters

Frequency (1 - 10 seconds)

Defines the **interval** between each bip/flash event.

- **1 second**
Flash and beep occur every second.
- **Higher values (e.g. 5 - 10 seconds)**
Flash and beep occur less frequently.

This value controls timing, not pitch or volume.

Visual characteristics

- Full-frame white flash
- Black background between flashes
- No additional graphics or motion
- High contrast for immediate visibility

Audio characteristics

- Short, percussive beep
- Broadband, attention-grabbing sound
- No pitch variation
- Occurs only at the flash moment

Contemporary use

Today, **Blip Flash** is used mainly as a technical or stylistic reference, not as a consumer-facing element.

Common modern uses include:

- Audio/video synchronization testing
- Verifying delay or latency in playback systems
- Timing reference during editing or demonstrations
- Visual and audio cueing in tutorials
- Intentional minimalist or “technical signal” aesthetics

Use with other filters

Blip Flash is often combined with filters depending on intent:

- **Gain / Volume**
Adjust beep loudness.
- **Color grading**
Change flash color or the black background for creative use.
- **Blend modes** like “Screen”
Overlay **Blip Flash** on footage as a flashing element.

Usage notes and tips

- Use short frequencies for sync testing.
- Use longer frequencies for cueing or demonstrations.
- Because the signal is repetitive, trim the clip to the required length.
- Always preview playback to confirm timing.

Limitations

- No control over beep pitch or waveform
- No control over flash duration
- Fixed visual style
- Not a substitute for precise measurement tools




Choose **Screen Snapshot** from the menu **File > New** or a toolbar **New Generator** button to generate an image of your computer screen. How it behaves depends on which operating system you use. This feature was added in Shotcut version 25.10, and it is not currently available in the Linux Flatpak (we hope to include it in the next version).

First of all, if you did not create a project folder it opens a file save dialog for you to choose the image file name and location. After the image is created, it opens in Shotcut's **Source** player. From there, you can add it to Playlist or Timeline.

Windows and Linux/X11

Next, you see the following toolbar either at the top or centered on the screen:



Ironically, there is no **Window** button when running on Windows!

If you do not want to include Shotcut's window, minimize Shotcut *before* choosing something from the toolbar. Shotcut will restore itself.

If you choose **Rectangle**, you need to use mouse or touch to draw a rectangle over the part of the screen to capture.

If you choose **Window** you need to click an application window to grab its image.

Linux/Wayland

On Wayland, there is no toolbar; it only does fullscreen capture and does not minimize Shotcut. You can edit the image (configure an **Open With** helper) or use a Shotcut crop video filter to isolate an area.

macOS

On macOS, there is no toolbar or fullscreen option. Rather, it goes directly into rectangle selection mode. To switch between rectangle and window selection press space before clicking anything; you will see the mouse pointer change.

Also, it always minimizes Shotcut in version 25.10. We plan to add a toolbar with an option whether to minimize Shotcut in a future version. Meanwhile, if you want to grab an image of Shotcut use shift + command + 5.

Choose **Screen Recording** from the menu **File > New** or a toolbar **New Generator** button to generate a video of your computer screen and record audio from your microphone (technically, the default audio input, which is usually the microphone). How it behaves depends on which operating system you use. This feature was added in Shotcut version 25.10, and it is not currently available in the Linux Flatpak (we hope to include it in the next version).

First of all, if you did not create a project folder it opens a file save dialog for you to choose the image file name and location. After the video is created, it opens in Shotcut's **Source** player. From there, you can add it to Playlist or Timeline.

Windows

On Windows, this simply launches the **Snipping Tool** that comes with Windows, which shows a toolbar at the top of the screen with a drop-down button where you can choose **Rectangle** (default) or **Window**. When done with the mouse or touch to draw the rectangle or pick the window, the toolbar changes and shows a **Start** button. Click it. When you are done click the Stop button in its toolbar. When done, a file may yet be created, and Shotcut will not open it. Rather, the full Snipping Tool opens. It has a save button that looks like . There is also a menu button  with  **Settings** where you can configure it to always save, and to which folder. Finally, you need to open the saved file yourself in Shotcut. A good way to do that is to use the integrated **Files** panel and its **Videos** preset location.


macOS

On macOS, there is no toolbar. Rather, it goes directly into fullscreen recording after choosing a save file. Make sure you choose a unique file name; the job will fail if you do not. It creates a job in the **Jobs** panel. You can minimize Shotcut yourself if you do not want to record it. (Trim the recorded video in Shotcut to remove the minimizing and restoring.) Right-click the job in Shotcut and choose **Stop This Job** and the recorded video should open in the **Source** player. If you minimized Shotcut do not forget to restore its window in order to stop the recording job.

P.S. This uses the `screencapture` command line utility that comes with macOS.

Linux/X11

Next, you see the following toolbar either at the top or centered on the screen:



If you do not want to include Shotcut's window, minimize Shotcut either before or after choosing something from the toolbar. Shotcut will restore itself.

If you choose **Rectangle**, you need to use mouse or touch to draw a rectangle over the part of the screen to capture.

If you choose **Window** you need to click an application window.

A new job appears in Shotcut's **Jobs** panel. It shows an updating duration of the recording. When you are done, right-click this job and choose **Stop This Job** and Shotcut opens the recorded video in the **Source** player. If you minimized Shotcut do not forget to restore its window in order to stop the recording job.

It is normal and expected if a dialog opens suggesting that you can convert this video because it is variable frame rate. You can if you want to, or click **Cancel** to dismiss it. You can always convert it later from **Properties** if you experience issues with timing.

Linux/Wayland


On Wayland, there is no toolbar, and the behavior depends on whether you are running GNOME, KDE, or something else. On GNOME or KDE, it creates a job in Shotcut's **Jobs** panel and launches `ScreenCapture` in GNOME Shell or `Spectacle` in KDE.

In **GNOME**, when done recording, right-click Shotcut's job and choose **Stop This Job**. The recorded file should open in the **Source** player.

In **KDE**, a job is created in **Shotcut**, but that is not where you stop it. Rather, in order to start it, you need to click the start button in the Spectacle panel that pops up. Likewise, to stop recording there is a Spectacle stop button in the KDE system tray. The job in Shotcut should stop automatically at this point and the recorded file opens in **Source** but not always.

If **something else** this is an OBS Studio app launcher if `obs` is in your `PATH`. Otherwise, it opens a file dialog for you to choose an executable that it saves to Shotcut's config file. Now, it is simply an app launcher for that tool.

Opens a **network-based media stream** by providing its URL.



This feature is used to **receive and play streams**, not to broadcast or stream to others.

What a network stream is

A network stream is audio/video data transmitted over a network using a streaming protocol instead of being stored as a local file.

Common examples include:

- Live streams (TV, cameras, radio)
- Network cameras
- Media servers
- Test streams
- Local network streams (LAN)

Shotcut uses FFmpeg to open and decode these streams.

What this feature does (and does not)

It does:

- Open and play network-based audio/video streams
- Treat the stream as a live or linear source in the **Source** player
- Allows recording the stream using the **Export** panel; however, this is a transcode operation. Shotcut cannot directly save streams.

It does not:

- Stream video to another person or service
- Act as a streaming server
- Provide chat, conferencing, or live broadcasting tools

This is a **receiver**, not a sender.

Parameter

URL

Specifies the **network stream address** to open.

The URL defines:

- The protocol
- The server or device
- The stream path

Examples of supported protocols (depending on build and FFmpeg support):

- `http://` / `https://` – progressive or streaming media
- `rtsp://` – IP cameras and live feeds
- `rtmp://` – live streaming servers
- `udp://` / `tcp://` – raw network streams
- `srt://` – secure, low-latency streaming

Not all URLs or protocols will work; support depends on MLT, FFmpeg, and the stream itself.

Typical use cases

- Capturing a live IP camera feed
- Recording a network TV or radio stream
- Ingesting a live stream for later editing
- Monitoring or archiving network-based video
- Testing streaming workflows

Behavior in Shotcut

- The stream opens as a **live source**
- Duration may be unknown or infinite
- Seeking may be limited or unavailable
- Stability depends on network conditions
- Reception continues until stopped

Usage notes and tips

- Ensure the URL is reachable and accessible.


- Authentication (username/password) may be required for some streams.
- Network latency and packet loss can affect playback quality.
- Not all online video websites provide direct stream URLs.
- DRM-protected streams cannot be opened.

Limitations

- Playback depends on network reliability
- No authentication mechanisms other than what you can put in the URL
- No built-in stream discovery
- No broadcasting or sending capability
- Limited control over stream parameters
- Some formats or protocols may not be supported

Opens a **professional video capture device** from **Blackmagic Design** that provides **SDI or HDMI inputs**. These devices are either PCIe cards installed in a desktop machine or Thunderbolt peripherals to be used with Macs and high end Windows laptops.

This allows Shotcut to record from external sources instead of importing media files.



What SDI / HDMI capture is

SDI (Serial Digital Interface) and HDMI are standards used to transmit **uncompressed digital video** from cameras, mixers, or playback devices.

Typical sources include:

- Broadcast cameras
- Video switchers
- External recorders
- Playback decks
- Professional capture cards (PCIe or external)

Shotcut accesses these devices through the system's capture framework and underlying drivers.

Scope and behavior

- Requires * **Blackmagic Design SDI/HDMI capture hardware**
- Device availability depends on:
 - Installed capture card

- operating system driver support
- If no supported device is detected, the **Device** field may appear inactive.
- Opens as a live video in the **Source** player
- At this point you can confirm things are working, audio/video levels, and microphone & camera placement
- Use the **Export** panel to record to a file
- This cannot be used as the input for **Timeline > Record Audio**

Controls

Device

Selects the SDI/HDMI capture device.

- Displays only supported capture cards
- May be inactive if:
 - No compatible card is installed
 - Drivers are missing or inactive

This control is unrelated to display screens or monitors.

Signal mode

Defines the **expected video signal format**.

Options include:

- **Detect Automatically**
- **HD formats** (720p, 1080i, 1080p)
- **SD formats** (NTSC, PAL)
- **UHD formats** (2160p / 4K)

Each entry specifies:

- Resolution
- Scan type (progressive or interlaced)
- Frame rate

Signal mode explained

Detect Automatically

Shotcut asks the capture device to detect the incoming signal format.

This is the most convenient option, but:

- Not all cards support auto-detection
- Detection may fail or misidentify the signal

Manually selected modes

When a specific mode is chosen:

- Shotcut expects the input signal to **exactly match**
- Mismatched settings can result in:
 - No video
 - Garbled image
 - Capture failure

Manual selection is often required for older or simpler cards.

Warning note:

Not every card model supports automatic signal detection, and not all cards support all signal modes.

This is a hardware limitation, not a Shotcut limitation.

Typical use cases

- Live capture from professional cameras
- Ingesting output from video mixers
- Recording SDI feeds from broadcast equipment
- Capturing HDMI output from external devices
- Studio or live production workflows

Usage notes and tips


- Verify the exact output format of the source device.
- If auto-detect fails, select the signal mode manually.
- Ensure the capture card supports the chosen resolution and frame rate.
- SDI and HDMI audio handling depends on the capture card.
- Test capture before long recordings.

Limitations

- Requires dedicated capture hardware
- Hardware and driver dependent
- Incorrect signal mode prevents capture
- Not suitable for consumer webcams or displays

Shotcut can record from your webcam or other audio and video inputs on your system.

This only appears on Windows and macOS. See [Open Other > Video4Linux](#) on Linux.




1. Choose your input devices. You can also choose None for one if you want an audio- or video-only capture.
2. Click **Open** to start playing these device inputs in the **Source** player. Shotcut is not yet recording to a file!
3. Use **Properties** to make more input changes if needed. It simply shows the Video input and Audio input drop-downs.
4. Adjust your camera framing, lighting, and audio/video levels at this point where these make sense.
5. To start recording to a file, use the **Export** panel, choose an appropriate preset, click **Capture File**, and enter a name in the file save dialog that opens.
6. When you are done recording, click **Stop Capture** in **Export**. (This button replaces **Capture File** while a recording is in progress.)
7. The captured file opens now in the **Source** player and is added to the **Recent** panel.

NOTE: This is not the recommended way to record audio to the timeline because it is not so convenient. Use **Timeline > Record Audio** for that. However, you can use this function (without file capture) to verify your audio device selection is working OK in Shotcut.



Opens a **video capture device** using the **Video4Linux (V4L2)** subsystem. This allows Shotcut to capture live video (and optionally audio) from supported devices instead of importing a media file.

This feature is **Linux-specific**.



What Video4Linux is

Video4Linux (V4L2) is the standard Linux framework for video capture devices. It is used by webcams, USB capture cards, TV tuners, and other video input hardware.

Shotcut uses V4L2 to access these devices directly and treat them as live video sources.

Scope and behavior

- Available **only on Linux**
- Captures **live video** from supported devices
- Appears as a source clip that can be:
 - Previewed
 - Recorded to a file in **Export**
- Device behavior depends heavily on kernel support

Controls

Device (device selector)

Selects the video capture device exposed by Video4Linux.

When compatible hardware is connected, this field lists available devices such as webcams, USB capture cards, or TV tuners. If no supported device is detected, the list may be empty.

Typical examples (when devices are present):

- Built-in webcam
- USB webcam
- HDMI / composite capture card
- TV tuner device

If no compatible device is connected, this list may be empty.

How devices usually appear

On Linux, Video4Linux devices are exposed as character devices under `/dev/`.

Common examples:

```
/dev/video0  
/dev/video1  
/dev/video2
```

In some cases, Shotcut may display both the path and a short description, depending on the driver and desktop environment, for example:

```
/dev/video0 (Integrated Webcam)  
/dev/video1 (USB Video Capture)
```

But the path is always the authoritative identifier.

Resolution (input, pixels)

Specifies the **capture resolution**.

- Must be supported by the selected device
- Unsupported resolutions may:
 - Be ignored
 - Fall back to a default
 - Fail to start capture

Example values:

- 640x480
- 1280x720
- 1920x1080

Frame rate (0.00 - 1000.00)

Sets the **capture frame rate** in frames per second (fps).

- Lower values reduce CPU usage
- Higher values require device and driver support

Note:

The actual frame rate may be limited by the hardware.

TV Tuner section

Only relevant if the selected device includes **TV tuner functionality**.

Standard

Selects the **analog television standard**:

- **Automatic**
- **NTSC**
- **PAL**
- **SECAM**

These standards correspond to historical analog broadcast systems and are relevant only for legacy TV tuner hardware.

Channel (0 - 99)

Selects the **analog TV channel number**.

- Channel numbering depends on regional broadcast standards
- Only meaningful for analog TV capture devices

Audio Input

Selects the audio capture backend.

Common options include:

- **None:** no audio capture
- **PulseAudio:** default audio system on most Linux desktops, also works with PipeWire
- **ALSA:** low-level Linux audio interface

The available audio input options depend on the Shotcut build and the system's audio configuration.

Important note:

Audio capture is handled separately from video and depends on system audio

routing.

Typical use cases

- Capturing webcam footage directly into Shotcut
- Recording from USB or HDMI capture cards
- Digitizing analog video sources (VHS, camcorders)
- Capturing legacy TV tuner signals
- Live preview of external video sources

Usage notes and tips


- Always test capture settings before long recordings.
- Use supported resolutions and frame rates to avoid failures.
- Audio sync depends on system audio configuration.
- Some devices expose limited or non-standard controls.

Limitations

- Linux-only feature
- Device support depends on drivers
- Analog TV features are legacy and hardware-dependent
- No built-in channel scanning for tuners
- Audio and video may require manual sync adjustment

Opens a **live audio input** using **PulseAudio** as the capture backend. This allows Shotcut to receive audio from the Linux desktop audio system instead of specifying a low-level hardware device.

This feature is **Linux-specific**.



What PulseAudio is

PulseAudio is a user-space audio server commonly used on Linux systems. It sits above **ALSA** and provides:

- Device abstraction
- Software mixing
- Input/output routing
- Per-application audio control

Most desktop Linux environments route microphones and audio inputs through PulseAudio or **PipeWire**, which is PulseAudio compatible.

Behavior in Shotcut

- This device has no configurable controls inside Shotcut.
- Audio source selection is handled entirely by the system.
- Shotcut receives whatever input PulseAudio is currently providing.

How audio input is selected

When using PulseAudio:

- Microphone or input device selection is done via:

- System sound settings
 - Desktop audio control panels
 - External tools such as `pavucontrol`
- Shotcut simply connects to the active PulseAudio input stream.

If the wrong microphone is used, it must be corrected outside Shotcut.

Scope and behavior

- Linux-only feature
- Audio capture only
- Uses the desktop's active audio routing
- No channel count or device selector in Shotcut
- Appears as a live audio source

This is the default and recommended audio capture path on Linux.

Typical use cases

Use this **Open Other** option to prepare for **Timeline > Record Audio** by verifying the audio input, audio input levels, and microphone placement.

- Voice-over recording
- General microphone recording
- Simple, system-managed audio input

Advantages

- Simple setup
- Works with most microphones
- Handles device switching automatically
- Compatible with desktop audio routing
- No ALSA device knowledge required

Limitations

- No direct device control in Shotcut
- Relies entirely on system configuration
- Not suitable for low-level or multi-channel capture
- Latency depends on PulseAudio/PipeWire configuration

Relationship to other capture options


- **PulseAudio**
 - High-level, desktop-managed
 - Recommended for most users

- **ALSA**
 - Low-level, hardware-specific
 - Intended for advanced or specialized workflows

Opens a **raw audio capture device** using **ALSA (Advanced Linux Sound Architecture)**.

This allows Shotcut to capture audio directly from an ALSA PCM device instead of importing an audio file.

This feature is **Linux-specific** and intended for advanced or low-level audio capture.



What ALSA is

ALSA is the low-level audio subsystem of the Linux kernel.

It provides direct access to sound hardware through PCM (Pulse Code Modulation) devices.

Unlike **PulseAudio** (or PipeWire), ALSA:

- Operates closer to the hardware
- Provides minimal abstraction
- Does not handle mixing, routing, or device management automatically

Shotcut accesses ALSA through FFmpeg.

Important context

On Linux, **Timeline > Record Audio** uses PulseAudio (or PipeWire), not ALSA. Microphone selection and routing are expected to be configured at the desktop/system level. Shotcut does not provide internal ALSA or PulseAudio configuration panels.

Controls

PCM Device (input)

Specifies the ALSA PCM device name to open.

- Default value: `default`
- This refers to the system's default ALSA capture device

Advanced users may enter explicit ALSA device identifiers, such as:

- `hw:0,0`
- `plughw:1,0`

If an invalid or unsupported device is specified, capture will fail.

Channels (0 - 8)

Specifies the **number of audio channels** to capture.

- 1 — Mono
- 2 — Stereo (default)
- Higher values — Multi-channel capture (if supported by the device)

The device must support the requested channel count.

Scope and behavior

- Linux-only feature
- Captures audio only
- Appears as a live source
- No internal level controls
- No automatic device selection

This capture path bypasses desktop audio routing.

Typical use cases

- Capturing audio from:
 - USB audio interfaces
 - Professional sound cards
 - External mixers
- Low-latency or direct hardware access
- Advanced workflows where PulseAudio/PipeWire is not setup
- Technical testing or diagnostics

Usage notes and warnings

- ALSA device names are system-dependent
- Incorrect PCM device strings will cause silent failure
- No automatic resampling or mixing is provided

Relationship to PulseAudio / PipeWire

- PulseAudio (or PipeWire) sits on top of ALSA
- **Timeline > Record Audio** uses PulseAudio
- ALSA device capture is a direct, bypass path


If recording fails via ALSA, a common workaround is to:

1. Close any other process using ALSA because often the device cannot be used by more than one process, which is a big reason why PulseAudio, PipeWire, and JACK exists.
2. Record audio using a dedicated ALSA-capable application
3. Import the resulting audio file into Shotcut

Limitations


- No device discovery UI
- No level meters
- No internal routing or mixing
- Requires ALSA knowledge
- Behavior depends on kernel driver and hardware support

In Shotcut, as well as in video **compositing** software, a keyframe is a frame used to indicate the beginning or end of a change made to a parameter. For example, a keyframe could be set to indicate the point at which audio will have faded up or down to a certain level.



Or, **Keyframing** is the simplest form of animating a change to a parameter.

Keyframing in a compositing software:-



Example of Result of keyframing in shotcut:-




Image source from [Zoom Image and then Pan - #2](#) by samth



Keyframe Types

Hold

A hold keyframe makes the value not change until possibly the next keyframe. Also, there is NO transition between the previous video frame's values and the time on this keyframe. Essentially, it introduces a hard cut to a new value and sustains it until otherwise. If you think about it in terms of video transitions; this is the hard or jump cut.

Linear


Linear is a straight line between keyframes like / or . Sometimes this is called a ramp. It is a simple way to adjust values by a consistent amount from one video frame to the next.

Smooth

Smooth keyframes draw the smoothest possible line between multiple keyframes. As a result, the values change can accelerate and decelerate. This is okay sometimes to smooth things, but for more control you probably want to use an ease curve.

Track Symbols


The symbol of the keyframe on the track in the **Keyframes** panel represents its type. However, it is double sided because the type actually applies between keyframes as explained below in Easing. So, you end up with shapes like these:



(Auto Bezier = Smooth)

Tool Tip

Hover the mouse over a keyframe for a few seconds shows a tool tip that shows the time position, the type, and its value on the second line.



Easing


Easing was added in version 23.11.

See the [Easing Functions Cheat Sheet](#) web page for an animated graphical overview of the Easing options. The icons in the menu are also a quick reminder of the aggressiveness of the acceleration or the trick (back, elastic, bounce).

The behavior applies *between* keyframes. Remember this: when thinking about “in” or “out”, for Shotcut “in” refers to the beginning of something - a clip, filter, or parameter between keyframes. And “out” refers to the end of something. So, this is not easing *into* a keyframe and easing *out of* a keyframe. Rather, it is easing *into* or *out of* the change/segment/span/tween of a parameter *between* keyframes.


To Next

You can choose between **Ease Out** or **Ease In/Out** to affect the behavior with the next keyframe:



From Previous

Choose the behavior between the previous keyframe and this one:




Simple keyframes are a quick way to animate a filter over the course of a clip. Simple keyframes only support two keyframes:

- The “start” keyframe is relative to the beginning of the clip and defines how the filter animates in.
- The “end” keyframe is relative to the end of the clip and defines how the filter animates out

The space between the start and the end keyframes is not animated

Simple keyframes are activated by dragging the circles in the upper left and right of the keyframes panel



- To set the starting values for the filter, move the cursor to the left side and then set the filter parameters in the filter panel.
- To set the middle values for the filter, move the cursor between the start and end keyframes and set the filter parameters in the filter panel.
- To set the ending values for the filter, move the cursors to the right side and then set the filter parameters in the filter panel.


Here is a demonstration that sets the brightness filter to animate in brightness, hold the brightness, and then animate out the brightness.

The Audio Loudness scope provides a real-time calculation of the program loudness for the preview player.


The meter displays 6 metrics as defined in [EBU-R128](#).

These metrics include:

- Momentary loudness
- Short term loudness
- Integrated loudness
- Loudness range
- Peak
- True peak



The display can be customized by clicking on the hamburger menu to disable the display of specific meters.




The meter readings could become inaccurate if the user seeks around the timeline. To get accurate readings for a segment of a clip or timeline, follow these steps:

1. Seek to the position that you want to begin measurement
2. Click the “Reset” button in the Audio Loudness scope
3. Click “Play” in the player.
4. Observe the Momentary, Short Term, Peak and True Peak values while the player plays (these are dynamic metrics that change over time)
5. When the player reaches the position where you want to stop, click pause
6. Observe the Integrated loudness and Loudness range. These values apply for the duration of the segment that was played.



The Audio Peak Meter Scope provides an illustration of the instantaneous peak audio value for each channel.

In this screenshot, the **Audio Channels** setting is set to 6 channels.




The scope also shows a “recent peak” indicator as a small line above the level bar.



The Audio Spectrum Scope provides an illustration of the instantaneous audio level for the range of frequencies of the audio.

The left axis represents the level of the frequency in dBFS. The bottom axis represents the frequency for each bar.




The scope also shows a “recent peak” indicator as a small line above each level bar.



The Audio Surround Scope was added in the 24.04 release.

The Audio Surround Scope provides a visualization of the multichannel sound.

In this figure, the Audio Surround Scope is shown with the [Audio Peak Meter Scope](#).



The center of the display represents the position of the listener.


The blue lines extending from the center of the display visualize the signal strength for each channel. In the illustration above, you can see the relationship between the level for each channel in the Audio Peak Meter and the Audio Surround Scope.

The purple lines connecting the blue lines help to illustrate the width of the sound in the room.

The green concentric circles illustrate the “focal point” of the sound. In the display above, the focal point is in front of the listener and slightly to the right.


The display adjusts to the number of channels in the [Audio Channels](#) setting. In this screenshot, the setting is set to Stereo (2 channel). So the C, LS and RS

channels are grayed out.




The Audio Vector Scope was added in the 24.04 release.


The Audio Vector Scope provides an illustration of the phase relationship between two channels. A typical display for two stereo channels is shown below:




To show the phase relationship between different channels, choose the channels from the drop down boxes:



If the vector scope shows a vertical line without any horizontal separation between the points, that means the two channels are perfectly in phase. Typically caused by a mono signal that is split to left and right:




If the vector scope shows a horizontal line without any vertical separation between the points, that means the two channels are perfectly out of phase. This means one channel is an inverted version of the other channel. This can be caused by microphones that are wired incorrectly. The problem with such an audio signal is that if they are summed together (such as converting stereo to mono), the channels will cancel out and the result will be silence. The following figure shows the vector display for two channels that are inverted. It also shows the waveform display where you can see that the waveforms are the exact opposite.







The Audio Waveform Scope provides an illustration of the waveform of each audio channel.



Passing the mouse over the display will show a bubble that indicates that value for each audio sample.




View > Scopes > Video Histogram opens a panel to display a histogram of the image luma and the red, green and blue color channels. The histogram updates with every image in the preview player.







The Video RGB Parade Scope shows a waveform for each of the R, G & B channels in a row.




The Video RGB Waveform Scope shows the waveforms for each of the R, G & B channels displayed on top of each other.



The Video Vector Scope provides a vectorscope display. This scope can be used to evaluate the hue and saturation in the image. The presence of colors are mapped as a circle containing Red, Magenta, Blue, Cyan, Green and Yellow demarcations. The more samples that exist for a given value, the brighter the sample will appear in the display.




When the mouse is passed over the scope, a bubble appears to display the U & V values for the sample under the mouse.



The scope also includes a skin tone line, located at about the 11 o'clock point between the red and yellow demarcation points. It is common for colorists to center the skin tones around this line.



Shotcut provides a video waveform scope to aid in setting video levels.



See [Video Scopes](#) for general information about video scopes.

The waveform scope displays graticules for 100 and 0 IRE. These graticules are hard fixed to a Y value of 16 = 0 IRE and a Y value of 235 = 100 IRE. These assignments are not user configurable.

While IRE is technically an analog video construct, it is included in this scope as a familiar and useful indicator for video levels.


Shotcut provides a video zoom scope to aid in inspecting images.



See [Video Scopes](#) for general information about video scopes.

The video zoom scope receives YUV 4:2:0 images from the preview player and converts the image to a zoomed in window as RGB. The mouse cursor allows the user to select individual pixels and shows the x,y position, YUV values, and converted RGB values.

NOTE: For RGB sources, some chroma bleeding occurs due to the conversion from 4:2:0 to RGB. This can be demonstrated by opening File->Open Other->Color Bars. The source color bar image is generated with no chroma bleeding. But the images is converted to YUV 4:2:0 for the preview display which results in vertical chroma bleeding as shown below:



This chroma bleeding may or may not appear in the exported file depending on the image format requested by the export parameters.




Filters in Shotcut are the primary way to apply effects, color grading, and audio adjustments to your media. The **Filters** panel allows you to manage these effects for individual clips, entire tracks, or the whole project.

Adding a Filter

To add a filter, first select a clip on the **Timeline** or open (double-click) an item in the **Playlist**.

1. Click the **Plus (+)** button at the bottom left of the panel or simply press `F`.
2. A menu will appear with several tabs:



- **Favorites (Star icon):** Your most frequently used filters.
- **Video (Monitor icon):** All available video effects and color tools.
- **Audio (Speaker icon):** All available audio processing tools.
- **Time (Stopwatch icon):** All available time and speed tools.
- **Sets:** Builtin and custom groups of filters you have saved.

3. Click a filter name to apply it to the selected item.

Searching for Filters

Instead of browsing through the long Video and Audio lists, you can use the search bar at the top of the “Add Filter” menu.

- Simply start typing the name of the filter (e.g., “Size” for *Size, Position & Rotate*).
- You can also enter a keyword for a filter such as “transform” because many effects are known by other names in other tools. Keywords contain synonyms and categories.
- You can also enter a “tag” such as `#10bit` for filters compatible with a 10-bit video processing pipeline or `#gpu` for filters that run on a GPU.
- The list will filter in real-time. Use the arrow keys and **Enter** to quickly select and apply a filter without using your mouse. Simply pressing `Enter` or the `Down` key

moves focus to the results.

Favoriting Filters

If you use certain filters regularly, you can add them to your **Favorites** tab for quicker access.

- Click the **Star** icon next to the filter name.
- Once starred, the filter will appear in the first tab (the Star tab) every time you click the **Plus (+)** button.

Removing a Filter

To remove an effect you no longer need:

1. Select the clip on the Timeline.
 2. In the Filters panel, click on the name of the filter you wish to delete.
 3. Click the **Minus (-)** button at the bottom of the panel.
- *Note: This will permanently remove the filter and all its current settings/keyframes.*

Reordering Filters

The order of filters in the list is the order in which Shotcut processes them. This is known as the “Filter Stack.” For example, a “Blur” filter placed *above* a “Mask” filter will behave differently than one placed *below* it.


- **To Reorder:** Click a filter in the list to highlight it, then use the **Up Arrow** and **Down Arrow** buttons at the bottom of the panel to move it.
- You can also drag-n-drop to reorder.
- **General Rule:** Usually, you want corrective filters (like White Balance) at the top and creative effects (like Vignette) near the bottom.

Disabling vs. Removing

You do not have to remove a filter to see how your video looks without it.

- Uncheck the **checkbox** next to the filter name in the list to temporarily disable it.
- Double-click the name of the filter in the list also toggles whether a filter is active.
- This allows you to toggle an effect on and off to compare the “Before” and “After” without losing your settings.

Shotcut allows you to copy the entire “stack” of filters from one clip and paste them onto another. This is an essential feature for maintaining visual consistency across multiple shots, such as applying the same color grade or audio normalization to several clips in a row.



How to Copy Filters

1. Select the clip on the **Timeline** that has the filters you want to copy.
2. Open the **Filters** panel (`Ctrl+6`).
3. Click the **Copy Filters** icon (two overlapping pages) located in the toolbar at the bottom of the panel.
4. A menu opens:
 - **Copy Current**
 - **Copy Enabled**
 - **Copy All**
5. Choose one.

How to Paste Filters

1. Select the destination clip on the Timeline where you want to apply the effects.
2. In the **Filters** panel, click the **Paste Filters** icon (the clipboard icon) in the bottom toolbar.
3. The copied filters will be added to the bottom of the existing filter list for that clip.

Selective Copying

If you have multiple filters on a clip but only want to copy *some* of them, use the following workflow:

- **Uncheck** the filters you do not want to copy in the filter list.
- Click the **Copy Filters** button.
- **Re-check** the filters on your original clip.
- Select the new clip and **Paste**. Only the filters that were checked at the moment of copying will be applied.

Applying Filters to Multiple Clips (Bulk Paste)

As of version 24.04, Shotcut includes a specialized feature to update many clips at once.

1. Copy the filters from your source clip as described above.
2. Select multiple clips on the Timeline (hold `ctrl` or `shift` while clicking).
3. **Right-click** on one of the selected clips.
4. Select **Timeline > Apply Copied Filters** from the context menu.

Behavior: If a filter with the same name already exists on the destination clips, Shotcut will **replace** it with the settings of the copied filter rather than creating a duplicate.

Copying Between Tracks and Source

- **Track to Clip:** You can copy filters applied to a **Track Header** and paste them onto an individual clip.
- **Clip to Track:** Conversely, you can copy a successful look from a single clip and paste it onto the Track Header to apply it to every clip on that track.
- **Source to Timeline:** Filters applied to a clip in the **Source** player can be copied and pasted onto any Timeline clip.

Copying Keyframes

When you copy and paste a filter, all associated **Keyframes** are also copied. However, keep in mind that keyframes are relative to the start of the clip. If you paste a filter with keyframes onto a clip of a different duration, you may need to adjust the timing in the **Keyframes** panel.

See also: [Filter Sets](#), [Timeline > Apply Copied Filters](#)

Filter sets first appeared in version 23.05.14.


A **Filter Set** is a collection of filters. A few examples are included:

- Glow Intensity
- Obscure With Blur
- Obscure With Mosaic
- Text Simple with Gradient Back


You can create your own from all enabled (checked) filters by clicking the **Save a filter set** button in the **Filters** panel.



This opens a window to let you give it a name. If you do not give it a name, the filters are copied to the clipboard.



To use a filter set, after clicking  to add a filter you can either search or click the **Sets** category to browse them.



To remove a filter set that you created (stock ones cannot be removed), go the **Sets** list and right-click the custom filter set you want to remove.

How to install a Filter Set

In the Forum [Resource](#) section, you will find many Filter Sets created and shared by members of the forum.


They are usually contained in a ZIP file and, once unzipped, they need to be moved to a specific folder on your computer.

Steps

- Download the ZIP file containing the filter set and un-zip it.

- In Shotcut, go to **Settings > App Data Directory > Show** to open the **filter-set** directory.
- In that directory you should see a ***filter-sets*** folder. If it is not there, you will have to create one (use that exact name: **filter-sets**).
- Open the **filter-sets** folder.
- Move the unzipped Filter Set file in that folder.
- **Close and re-launch** Shotcut.

To find your new filter set in Shotcut, go to **Filters panel > Sets**




Applies a **directional blur and blending treatment** to the top and/or bottom regions of a 360° equirectangular image.

The filter is designed to soften, hide, or visually integrate the polar areas (zenith and nadir) of 360° equirectangular video, which often contain strong projection distortion, stitching artifacts, or visible parts of the camera mounting (such as tripods or rigs).

Important Note:

This filter does not change the projection of the image; it only modifies the polar regions and is typically used in combination with other 360° projection or reframing filters.



This is a **360°-specific filter**. All controls operate in **angular space (degrees)**, not pixels.

General behavior (important)

- The filter operates on **angular bands** at the top and/or bottom of the frame.
- Effects are applied **progressively** between defined angular limits.
- The blur is **directional and animated relative to motion**, not a static matte.
- The effect must be evaluated **during playback**; a still frame can be misleading.

Note:

This filter is meant to be seen in motion. On a still frame, the result can look incorrect or overly aggressive.

Parameters

Interpolation

Defines how pixels are resampled when the image is warped and blurred.

- **Bilinear**
Smooth interpolation using neighboring pixels.
Produces softer transitions and is the preferred mode for natural footage.
- **Nearest-neighbor**
Uses the closest pixel without smoothing.
Produces hard edges and visible stepping; mainly useful for debugging or stylized results.

Top (enable checkbox)

Enables processing of the **top (zenith)** region of the 360° image.

If disabled, all Top-related parameters are ignored.

Start (0.000 - 90.000°)

Defines the **angle from the top pole where the effect begins**.

- Lower values affect a smaller cap near the pole.
- Higher values extend the effect further into the image.

End (0.000 - 90.000°)

Defines the **angle where the effect reaches full strength**.

- The region between Start and End is progressively affected.
- End values closer to 90° push the effect deeper toward the equator.

Fade (0.000 - 90.000°)

Controls how **gradually the effect fades** between Start and End.

- **0°**
Hard transition.
- Higher values
Softer, more progressive blending.

Blend

Controls how the processed region blends with the unmodified image.

In (0.000 - 90.000°)

Out (0.000 - 90.000°)

- **In** controls the blend entering the affected region.
- **Out** controls the blend leaving it.

These parameters soften the visual boundary and help avoid visible seams.

Blur

Controls the **directional blur envelope** applied to the cap.

Width at start (0.000 - 360.000°)

Height at start (0.000 - 90.000°)

Defines the blur extent at the **beginning** of the affected region.

- Width corresponds to horizontal angular spread.
- Height corresponds to vertical angular spread.

Width at end (0.000 - 360.000°)

Height at end (0.000 - 90.000°)

Defines the blur extent at the **end** of the affected region.

This allows the blur to **expand, contract, or change shape** across the cap.

Important behavior

- Blur is not uniform; it is **interpolated between start and end values**.
- Very small Height values can still produce strong perceptual blur due to angular compression near the poles.

Bottom (enable checkbox)

Mirrors all **Top** controls, but applies them to the **bottom (nadir)** region of the image.

All parameters behave identically, but relative to the bottom pole.

Keyframes

All parameters are keyframeable.

This enables:

- Dynamic hiding or revealing of poles

- Motion-reactive visual effects
- Stylized transitions tied to camera or subject movement

Visual characteristics

- Strong directional blur near poles
- Angular stretching and smearing
- Seamless blending when tuned correctly
- Effect intensity increases with motion

On static frames, the image may appear distorted or unpleasant; during playback, motion integrates the blur perceptually.

Recommended use cases

- Hiding tripods or camera rigs in 360° footage
- Masking stitching or pole artifacts
- Stylized motion effects in VR environments
- Making polar regions less visually distracting

Important!

Typical filter chains (Common 360° workflows)

- **Equirectangular** → **Stereographic**
→ *then* **Cap Top & Bottom**
(hide stretched poles after projection)
- **360° Reframe / Viewpoint change**
→ *then* **Cap Top & Bottom**
(soften polar distortion introduced by reframing)
- **Raw 360° footage**
→ **Cap Top & Bottom**
(hide tripod or stitching before further processing)

What *not* to expect

Using **Cap Top & Bottom alone** will never:

- Turn flat footage into a planet
- Change perspective
- Animate the scene

It is **not a generator**, only a **polar modifier**.

Limitations

- Only meaningful for equirectangular 360° video
- Difficult to judge on still frames

- Aggressive settings can look incorrect when paused
- Not intended for conventional (non-360°) footage, except for deliberate stylistic use.



Adds a black matte to the frame. Use this if you filmed using a 360 camera but only want to use part of the 360 image - for example if you and the film crew occupy the 90 degrees behind the camera.

Parameters

Horizontal

- **Start:** the width of the field-of-view in degrees of the un-matted area.
- **End:** the width of the field-of-view in degrees where the matte is at 100%.

Vertical

- **Start:** half the height of the field-of-view in degrees of the un-matted area.
- **End:** half the height of the field-of-view in degrees where the matte is at 100%.



Converts an equirectangular frame (panoramic) to a rectilinear frame (what you're used to seeing). Can be used to preview what will be shown in a 360 video viewer.

Parameters

- **Interpolation:** Nearest-neighbor or bilinear. Determines the sampling method. See also "Interpolation" in the [Glossary / Terminology](#)
- **Yaw, Pitch and Roll:** The direction of the image center in the panorama.
- **FOV:** The horizontal field of view, in degrees, of the resulting frame. Any value over 179 results in a fisheye projection.
- **Fisheye:** The amount of fisheye to mix in. 100 means that you get a 100% fisheye lens.

VUI


When this filter is active you can click and hold in the video area to drag the video around to adjust it:

- left-right for **Yaw**
- up-down for **Pitch**
- up-down for "Roll"

This does prevent being able to drag from the Source player. Either use a toolbar icon or keyboard shortcut to perform the intended operation, or select a different filter to turn off the VUI.

Projects a **360° equirectangular image** into a **stereographic projection**, producing a non-linear, globe-like or “tiny planet” style view.

The filter remaps spherical coordinates to a plane, transforming the full 360° environment into a visually compressed and highly distorted perspective.



This is a **projection filter**, not a blur or distortion effect. It changes how the spherical image is mapped to the frame.

360° context

In 360° video, the image represents a full sphere mapped onto a rectangle (equirectangular projection).

This filter converts that spherical representation into a **stereographic projection**, where angles are preserved but distances and areas are increasingly distorted away from the center.

Parameters

Interpolation

Controls how pixels are resampled during projection.

- **Bilinear**

Smooth interpolation using neighboring pixels.

Produces softer edges and is recommended for natural footage.

- **Nearest-neighbor**

Uses the closest pixel without smoothing.

Produces hard edges and visible stepping; mainly useful for stylized or technical results.

Yaw (–360.000 to 360.000°)

Rotates the view **horizontally** around the vertical axis.

- Changes the viewing direction left or right
- Equivalent to turning your head horizontally inside the 360° sphere

Pitch (–180.000 to 180.000°)

Rotates the view **vertically**.

- Positive and negative values tilt the view up or down
- Useful for centering the projection on the ground, horizon, or sky

Roll (–180.000 to 180.000°)

Rotates the view **around the viewing axis**.

- Spins the projected image clockwise or counter-clockwise
- Does not change what is being looked at, only its orientation

FOV (0.000 to 180.000°)

Controls the **field of view** of the stereographic projection.

- **Low values**
Narrow view, strong magnification, extreme distortion near edges
- **Medium values**
Balanced globe-like appearance
- **High values**
Wider view with less aggressive distortion

Important behavior

- Values near **180°** approach the limits of stereographic projection and can produce extreme stretching.

Amount (0 - 100%)

Controls how strongly the stereographic projection is applied.

- **0%**
With valid 360° equirectangular input, the image remains in its original projection.
- This filter has no true neutral state for non-360 footage; distortion occurs even at 0% Amount due to incompatible image geometry.
- **Low values**
Partial projection mixed with the original mapping.

- **100%**
Full stereographic projection.

Note

Amount blends between the original projection and the stereographic one.

Copy Parameters

Copies the current parameter values to the clipboard so they can be pasted into another instance of the same filter.

This is useful for:

- Applying identical projections to multiple clips
- Maintaining consistency across edits
- Reusing complex setups without manual re-entry

Keyframes

All parameters can be keyframed.

This enables:

- Animated camera moves inside a 360° scene
- Smooth transitions between viewpoints
- Dynamic projection effects over time

Visual characteristics

- Strong non-linear distortion
- Globe, “tiny planet,” or inverted sphere appearances
- Increasing distortion away from the center
- Motion-dependent perception; best evaluated during playback

Recommended use cases


- Creating “tiny planet” or spherical panorama visuals
- Stylized transitions from 360° footage to flat video
- Exploratory or artistic 360° presentations
- Animated reframing of 360° environments

Limitations

- Not intended for conventional (non-360°) footage
- Projection distortion is intentional and non-corrective
- Extreme FOV values can produce unreadable imagery
- Requires equirectangular 360° input for meaningful results

Applies a **selective wrap-and-blur treatment** to a rectangular region of a 360° equirectangular image.

The filter preserves a central angular window while **wrapping and blurring the areas outside it**, allowing discontinuities at the image edges to be softened or hidden.



This is a **360°-specific utility filter**. All parameters operate in **angular coordinates (degrees)**, not pixels.

360° context

In equirectangular 360° video, the **left and right edges of the image meet**, representing a full 360° horizontal wrap.

Hard cuts, reframing, or projection changes can make these edges visually discontinuous.

360: Equirectangular Wrap is designed to manage these transitions by wrapping and blurring areas outside a defined angular window.

General behavior (important)

- The filter defines a **central rectangular angular region** that remains readable.
- Everything **outside** this region is progressively blurred and wrapped.
- The effect is **spatial**, but perception improves during playback.
- It does **not** change the projection or viewpoint.

Note:

This filter defines what stays clear; everything else is treated as peripheral context.

Parameters

Horizontal

Defines the **horizontal angular window** to preserve.

Start (−180.000 to 180.000°)

Sets the **left boundary** of the preserved region.

End (−180.000 to 180.000°)

Sets the **right boundary** of the preserved region.

- The region between Start and End is treated as the main visible window.
- Areas outside this range are wrapped and blurred.

Vertical

Defines the **vertical angular window** to preserve.

Start (−90.000 to 90.000°)

Sets the **lower boundary** of the preserved region.

End (−90.000 to 90.000°)

Sets the **upper boundary** of the preserved region.

Together, the Horizontal and Vertical ranges define the **central rectangular area** you observed.

Blur

Controls how strongly the outside regions are softened.

Blur Start (0.000–2.000°)

Defines the blur amount at the **inner edge** of the preserved region.

End (0.000–2.000°)

Defines the blur amount at the **outer edge**.

Important behavior

- Blur is applied **only outside** the defined rectangle.
- Even small values have a strong perceptual effect due to angular compression.
- Blur is progressive between Start and End.

Note:

Blur values are angular, not pixel-based; small numbers go a long way.

Keyframes

All parameters can be keyframed.

This allows:

- Animated reframing of the preserved region
- Smooth transitions between focus areas
- Motion-guided emphasis inside 360° scenes

Parameter interaction

- **Horizontal + Vertical** define *what stays readable*.
- **Blur** defines *how softly the rest of the sphere is integrated*.
- Narrow windows increase perceived blur dominance.

If the preserved region is very small, the effect may appear as a blurred surround with a clear “window” in the center.

Visual characteristics

- Clear rectangular region in angular space
- Wrapped and blurred surroundings
- No hard seams at 360° boundaries when tuned correctly
- Best understood during playback

Does this filter need another filter?

No; it can operate on its own, but it is most useful in a 360° filter chain.

Typical combinations

- **Equirectangular Wrap → 360° Reframe / Viewpoint change**
Softens edge discontinuities after reframing.
- **Equirectangular Wrap → Equirectangular to Stereographic**
Helps integrate wrapped edges after projection changes.
- **Raw 360° footage → Equirectangular Wrap**
Reduces visual harshness near boundaries without altering projection.

This filter is **supportive**, not transformative.

Recommended use cases

- Softening 360° edge seams
- Limiting viewer attention to a central region
- Creating guided-focus 360° visuals
- Reducing discomfort caused by hard angular transitions

Limitations

- Not intended for non-360° footage
- Does not change projection or camera orientation
- Blur control range is intentionally narrow
- Effect can look confusing on still frames

Converts a video frame with two hemispherical images to a single equirectangular frame. The plugin assumes that both hemispheres are in the frame. If you have a camera like the Garmin Virb360 that produces two videos, one from each camera, you should start by converting them to a single movie by placing them side by side using, for example, [ffmpeg](#) (you can also add parameters to produce lossless, intra-only output here for easier editing):

```
ffmpeg \
  -i left.mp4 \
  -i right.mp4 \
  -filter_complex hstack \
  output.mp4
```

Parameters

- **Interpolation:** Nearest-neighbor or bilinear. Determines the sampling method. See also “Interpolation” in the [Glossary / Terminology](#)
- **Yaw, Pitch and Roll:** The two fisheye cameras of a 360 camera are rarely perfectly 180 degrees apart. These parameters adds a rotation to both lenses to correct for this.
- **Projection:** The fisheye projection type. Currently only equidistant fisheyes, like the Ricoh Theta and Garmin Virb360 are supported.
- **FOV and Radius:** The field of view of a single hemisphere in degrees, and the radius of the image circle, expressed as a fraction of the frame width.
- **X, Y and Up:** The image center and image orientation for the front- and back-facing camera. The **X** parameter is a fraction of the image width, and the **Y** parameter a fraction of the image height. **Up** is the “up” direction in the image, in degrees clockwise from a direction towards the frame top edge.
- **Nadir Radius and Nadir Start:** 360 cameras like the Theta have a problem with the nadir direction where, no matter what, you will have a little of the camera in the image. This parameter “stretches” the image near nadir to cover up the missing parts.



Converts a rectilinear (a normal-looking) image to an equirectangular image. Use this together with **360: Transform** to place “normal” footage in a 360 video.

Parameters

- **Interpolation:** Nearest-neighbor or bilinear. Determines the sampling method. See also “Interpolation” in the [Glossary / Terminology](#)
- **Horizontal:** The field of view’s width in degrees of the rectilinear image.
- **Vertical:** The field of view’s height in degrees of the rectilinear image.

Stabilizes 360 footage. The plugin works in two phases - analysis and stabilization. When analyzing footage, it detects frame-to-frame rotation, and when stabilizing it tries to correct high-frequency motion (shake).

How to Stabilize Video

1. Disable **Realtime (frame dropping)** in the **Settings** menu. It is important that the filter sees all frames.
2. Add the 360 footage
3. Apply a **Hemispherical to Equirectangular** filter to it so it is in equirectangular format.
4. Apply a **Transform 360** filter
5. Apply the **Stabilize 360** filter.
6. Select a file to store stabilization data in.
7. Enable **Analyze** mode.
8. Use the **Transform 360** filter to rotate the footage so that the point straight ahead is over the center cluster of track points.
9. Play the footage from start to finish. The FUD (Filter-Up Display) should show an increasing frame count.
10. When the footage has completed playing, switch off the **Analyze** mode.
11. You should now have stable 360 video.

The FUD (Filter-Up Display)

When you switch to analysis mode, the filter will overlay information on the video.

- In the top left is information about the file you use for storage and how many frames it has data for.
- The squares are the areas used for motion detection.
 - Red outline are the areas that were sampled.
 - Yellow outline is the search radius
 - Blue squares are fixed.
 - Green squares are where the sampled areas were found.
- When analyzing, the red, yellow and green squares are drawn first. Then the frame is transformed in such a way as to cancel the rotation from the previous frame, and the blue squares are drawn. If the detection and detected rotation are flawless, the blue and green squares should overlap completely and form cyan squares. Note: "Analysis: Apply Transform" must be enabled.
- The waveform drawn in the lower middle shows the detected motion. Red is yaw, green is pitch and blue is roll.

Parameters

- **Mode:** Toggle this checkbox to go from stabilization mode to analysis mode.
- **File:** Path to file that will be used to store the analysis data.
- **Start Offset:** The offset into the stabilization file that corresponds to the start of this clip. Press the **Undo** button to set it from Shotcut timeline. For example, if you have a 30 second clip, analyze it all, and then split it into three clips of 10 seconds each, then the start offsets should be 0s, 10s, and 20s.
Added in Shotcut 20.11
- **Interpolation:** Output quality.

Analysis

- **Apply Transform:** When checked, the filter will apply the frame-to-frame transform in order to show the quality of the analysis (see the above section about the FUD). When you are satisfied with the analysis quality you can turn this off, as it adds approximately 50% to the analysis time.
Added in Shotcut 20.11
- **Sample Radius:** The radius of the square that the stabilizer will sample.
- **Search Radius:** The maximum amount of motion the stabilizer will detect.
- **Offset:** The distance between the track points.
- **Use backwards-facing track points:** If set, six backwards-facing track points will also be used to detect pitch and yaw motion. Disable if, for example, you show up holding the camera there.

Yaw, Pitch & Roll

- **Amount:** The amount of stabilization to apply. 100% means that the stabilizer will make the camera as steady as it can. Smaller values reduce the amount of stabilization.
- **Smoothing:** The number of frames to use to smooth out the shakes. The higher the value, the slower the camera will follow any intended motion.
- **Time Bias:** Shift the frames used to smooth out the shakes relative to the stabilized frame. A value less than zero will give more weight to past frames, and the camera will seem to lag behind intended movement. A value greater than zero will give more weight to future frames, and the camera will appear to move ahead of the intended camera movement. A value of zero should make the camera follow the intended path.



Rotates a panoramic image.

Parameters


- **Interpolation:** Nearest-neighbor or bilinear. Determines the sampling method. See also “Interpolation” in the [Glossary / Terminology](#)
- **Yaw, Pitch and Roll:** The amount to rotate the image.

VUI

When this filter is active you can click and hold in the video area to drag the video around to adjust it:

- left-right for **Yaw**
- up-down for **Pitch**
- up-down for “Roll”

This does prevent being able to drag from the Source player. Either use a toolbar icon or keyboard shortcut to perform the intended operation, or select a different filter to turn off the VUI.



This filter adjusts the **alpha channel** in a number of ways.

Mode works in conjunction with **Amount**.

The modes primarily affect the edges of the alpha channel, which is very helpful for chroma keying. It is cascadable, so for example one can do a soft shrink first and then threshold, which gives a slightly different result than a hard shrink.

- **No Change**
- **Shave** tries to remove the “hairy” stuff, and also shrinks the selection a bit.
- **Shrink Hard**
The “hard” operations introduce no new values to the alpha channel, so if you have a “hard” key (only 0 and 255) it will stay that way.
- **Shrink Soft**
The “soft” operations will introduce interpolated values, making the edge softer.
- **Grow Hard**
- **Grow Soft**
- **Threshold**
- **Blur** simply blurs the alpha channel with a quasi Gaussian blur.
- **Reset** (added in version 23.09) changes the alpha channel to all zero for transparent. Combine this with **Invert* to make the alpha channel entirely opaque.

Amount controls how much the selected mode applies.

Invert means transparent becomes opaque, and opaque becomes transparent. However, these are just the terms for 0 or 100%. If the alpha level of a pixel is 50%, then invert does nothing. Or, 25% inverts to become 75%, for example.

See also

[frei0r/src/filter/alpha0ps at master · dyne/frei0r](#)

A large collection of free and portable video plugins - dyne/frei0r




The **Alpha Channel: View** filter is intended for temporary use to assist in visualizing the alpha (transparency) channel of your video. It's particularly useful for fine-tuning transparency effects, chroma keying, and masking operations.

It displays the transparency information of your clip, making it easier to see which areas are transparent and which are opaque.

See the [Glossary/Terminology](#) page to learn more about Alpha Channels.


The **Alpha Channel: View** filter was added in Shotcut version 16.01



How to Use the Alpha Channel: View Filter

1. **Add the Filter:** In the Filters panel, search for the “**Alpha Channel: View**” filter and add it to your video clip.
2. **Choose a Display option:** The filter will display the alpha channel of your clip differently depending on the **Display** option you choose.

To show the difference between the **Display** options, we will use this sample project: A fully opaque image on track V1 and a clip with transparencies on track V2.



Display options


Gray Alpha

Shows the alpha channel as a grayscale image, where **white** areas are fully opaque, **black** areas are fully transparent, and **shades of gray** represent varying levels of transparency.



Red & Gray Alpha

Shows the fully transparent pixels in gray and overlays a red-pink color over the other pixels in the image. The intensity of the red-pink color varies according to the level of transparency.





Checkered Background

Shows a checkered background under the transparent and translucent areas.



Black, Gray or White Background

Shows a black, gray or white background under the transparent and translucent areas.






Alpha Strobe Video Filter

The **Alpha Strobe** filter alternates the alpha channel (transparency info) between its original state and full transparent on a timer. Think of it as the blink effect.

This filter was added in Shotcut version 25.05.




Use **Interval** to control the timer; a higher % is slower.

The **Audio Dance Visualization** filter is a creative tool that allows to animate your video based on the audio waveform of a video or audio track. This filter is perfect for adding dynamic, eye-catching effects to your projects.

This filter was added in Shotcut version **19.04.30**

IMPORTANT: The clip or track that the filter is applied to must have audio. The filter will not respond to the audio from a different track or clip.

If your music is in an AUDIO track and want the visualization to appear on the video from video tracks, you should add the filter to the “Output” track in the top, left corner of the Timeline.




Settings

- **Initial Zoom:** The amount of zoom to apply to the image before any motion occurs.
- **Oscillation:** Oscillation can be useful to make the image move back and forth during long period of sound.
- **Zoom:** The amount that the audio affects the zoom of the image.
- **Up, Down, Left, Right:** The amount that the audio affects the vertical and/or horizontal offset of the image.
- **Clockwise, Counterclockwise:** The amount that the audio affects the rotation of the image.
- **Low Frequency:** The low end of the frequency range to be used to influence the image motion.
- **High Frequency:** The high end of the frequency range to be used to influence the image motion.
- **Threshold:** The minimum amplitude of sound that must occur within the frequency range to cause the image to move.

Displays a real-time **visual representation of audio amplitude** as animated bars or segmented meters over a video frame.

The visualization reacts to the audio signal associated with the clip and renders directly into the video output.



This filter is intended for **visual feedback**, not audio analysis or correction.

Important requirement (mandatory)

The audio source must be on a video track; audio-only tracks are not supported.

- The filter reads audio **only from the clip it is applied to**.
- A video clip with embedded audio works as expected.
- Audio-only clips on audio tracks **cannot drive this filter**.


If no usable audio is present on the video clip, the visualization will not animate.

Parameters

Type (Bar | Segment)

Selects the **visual style** of the level meters.

- **Bar**
Continuous bars that grow and shrink smoothly with the audio level.
- **Segment**
Bars divided into discrete segments, similar to LED or broadcast meters.




Graph Colors (1 - 10)

Defines the **color ramp** used for the levels.

- You must explicitly set the **number of colors** (1 to 10).
- Colors are interpolated across the meter.
- Click directly on the horizontal color bar to:
 - Set individual colors
 - Adjust intermediate color stops

Important behavior

- With only one color, the entire meter uses that color.
- Multiple colors create a gradient that follows level intensity.



Background Color

Sets the background behind the visualization.

- Default is **transparent**
- Can be changed using the color picker or by clicking the color control

Note:

The background is rendered as part of the video image. Timeline track compositing works normally; alpha transparency in the exported file is relevant **only** if transparency is required and the export format supports an alpha channel.

Thickness (0 - 100 px)

Controls the **thickness of each bar or segment**.

- Affects visual weight, not sensitivity
- Larger values produce bolder meters

Position (X, Y)

Defines the **on-screen position** of the visualization.

- Can be adjusted numerically
- Can be moved directly in the preview window

Size (Width, Height)

Defines the **overall bounding box** of the visualization.

- Can be resized numerically
- Can be resized interactively in the preview

Note:


Size controls the drawing area; Thickness controls the bar width inside it.

Mirror the Levels (checkbox)

Mirrors the meters around a central axis.

- When enabled, levels expand **symmetrically up and down** from the center.
- When disabled, levels grow in a single direction.

This is a **visual-only** option.



Reverse the Levels (checkbox)

Reverses the **horizontal direction** of the visualization.

- Left becomes right
- Right becomes left

Useful for layout or compositional symmetry.

Channels (1 - 10)

Controls the **number of audio channels visualized**.

- Each channel is displayed as a separate bar or segment column.
- This does **not** create or mix audio channels.

- If the source audio has fewer channels, remaining meters may duplicate or remain inactive depending on the source.

Important clarification:

This represents audio channels being visualized, not physical speakers.

Segments (1 - 100)

(Segment mode only)

Controls how many **discrete segments** each bar is divided into.

- Low values → coarse, blocky meters
- High values → fine-grained segmentation

Segment Gap (0 - 100)

(Segment mode only)

Controls the **spacing between segments**.

- **Low values** → thin gaps, visually thicker segments
- **High values** → wide gaps, thinner-looking segments

Note:

Increasing the gap reduces the visible segment thickness.

Keyframes

This filter does **not** support keyframes.

All animation is driven **entirely by the audio signal**.

Visual characteristics

- Real-time response to audio amplitude
- Vertical meters following music, speech, or sound effects
- Deterministic and repeatable for the same audio
- Rendered directly into the video output

Recommended use cases


- Music visualizers
- **Podcast** or voice-over videos
- Educational or technical demonstrations
- Audio-reactive overlays without scripting

Limitations

- Requires audio on a **video track**
- No control over frequency bands (not a spectrum analyzer)
- No peak hold or RMS/peak mode selection

Displays an **audio-reactive light shape** that pulses in intensity based on the audio signal.

The result is a soft, glowing, ellipsoid-like gradient that brightens and dims over time, rather than a meter or spectrum that follows specific frequencies or waveforms.




This filter is intended for **ambient, abstract audio visualization**, not for precise audio analysis.

Important requirement (mandatory)

The **audio must be part of a video** clip on a video track.

- The filter reads audio **only from a video clip that also provides video frames**.
- A normal video with embedded audio works.
- A static image or album cover **with embedded audio** (e.g. music + cover art) also works.
- Audio-only clips, even if placed on a video track, **do not drive this filter**.
- **But, when applied to the Output**, the filter uses the project's mixed audio and runs for the entire duration of the timeline.



Parameters

Waveform Color (1 - 10)

Defines the **color gradient** of the light visualization.

- You must explicitly choose the **number of colors** (1 to 10).
- Click the horizontal color bar to:
 - Set colors
 - Adjust intermediate color stops

Important behavior

- One color → uniform glow
- Multiple colors → smooth gradient transitions within the light shape

Position (X, Y)

Controls the **on-screen position** of the visualization.

- Adjustable numerically
- Can be moved directly in the preview window

Size (Width, Height)

Defines the **bounding area** in which the light visualization is drawn.

- Adjustable numerically
- Can be resized interactively in the preview

The light shape is always constrained to this rectangle.

Oscillation (0 - 10 Hz)

Controls the **base pulsing rate** of the light.

- **Low values**
Slow, smooth breathing-like motion
- **Higher values**
Faster flickering or pulsing

Important behavior

- Oscillation defines the *temporal rhythm* of the effect.
- Audio intensity modulates this motion rather than replacing it.

Note:

- Oscillation sets how fast the light wants to pulse; audio loudness controls how strong the pulse becomes.

Low Frequency (20 Hz - 19 900 Hz)

Defines the **lowest frequency** considered when measuring audio intensity.

- Raising this value ignores bass content
- Useful to avoid low-frequency dominance

High Frequency (120 Hz - 20 000 Hz)

Defines the **highest frequency** considered.

- Lowering this value limits the effect to bass and mids
- Higher values include treble energy

Threshold (–60 to 0 dB)

Sets the **minimum audio level** required to trigger visible response.

- **Lower values**
More sensitive; quiet audio produces visible light
- **Higher values**
Less sensitive; only louder passages trigger strong response

Note:

- This effect is time-based and audio-driven; it must be evaluated during playback, as **no motion is visible on a paused frame**.

Keyframes

This filter does **not** support keyframes.

All animation is driven by:

- The audio signal
- The Oscillation setting

Visual characteristics

- Soft, glowing, ellipsoid-like gradient
- Pulsing brightness rather than discrete motion
- No frequency-by-frequency display
- Best evaluated during playback

The visualization reacts primarily to **overall intensity (loudness)**, not to rhythm details or pitch structure.

Recommended use cases


- Ambient music visuals
- Background motion for static imagery
- Minimalist audio-reactive overlays
- Atmospheric or abstract video styles

Limitations

- Requires audio embedded in a video clip
- Not a waveform or spectrum analyzer
- No per-channel or per-band control
- UI elements are rendered into the video

Displays a **real-time frequency spectrum** of the audio signal as animated lines, bars, or segments over a video frame.

Unlike level meters, this filter visualizes **frequency content** (low to high frequencies) rather than overall amplitude.



The visualization is rendered into the video output and reacts in real time to the audio signal, so it **must be evaluated during playback** rather than on a paused frame.

Important requirement (mandatory)

The audio source must be **attached to a video clip**.

- The filter reads audio **only from the video track it is applied to**.
- Video clips with embedded audio work as expected.
- Audio-only clips on audio tracks **cannot drive this filter**.

If no usable audio is present on the video clip, the spectrum will not animate.


Parameters

Type (Line | Bar | Segment)

Selects the **visual representation** of the spectrum.


- **Line**

A continuous curve following frequency magnitude.




- **Bar**

Vertical bars representing frequency bands.



- **Segment**

Bars divided into discrete blocks, similar to LED-style analyzers.



Graph Colors (1 - 10)

Defines the **color ramp** used to draw the spectrum.

- You must explicitly choose the **number of colors** (1 to 10).
- Click the horizontal color bar to:
 - Set colors
 - Adjust intermediate color stops

Important behavior

- One color → uniform spectrum color
- Multiple colors → interpolated gradient along the spectrum



Background Color

Sets the background behind the spectrum visualization.

- Default is **transparent**
- Can be changed using the color picker or by clicking the color control in the preview

Note:

The background is rendered into the image. Alpha transparency on export matters only when transparency is needed and the chosen format supports it.

Thickness (0 - 100 px)

Controls the **stroke thickness** of lines, bars, or segments.

- Affects visual weight only
- Does not change frequency resolution or sensitivity

Position (X, Y)

Sets the **on-screen position** of the spectrum.

- Adjustable numerically
- Can be moved directly in the preview window

Size (Width, Height)

Defines the **drawing area** for the spectrum.


- Adjustable numerically
- Can be resized interactively in the preview

Note:

Size defines the available space; Thickness defines how bold the spectrum appears inside it.

Fill the area under the spectrum (checkbox)

- Applies to **Line** mode only
- Fills the area beneath the curve using the selected Graph Colors
- Has no effect in Bar or Segment modes



Mirror the Levels (checkbox)

Mirrors the spectrum around a central horizontal axis.

- When enabled, frequencies expand **upward and downward** from the center
- When disabled, the spectrum grows in a single direction

This is a visual option only.

Reverse the Spectrum (checkbox)

Reverses the **frequency direction**.

- Low frequencies appear on the right instead of the left
- High frequencies appear on the left instead of the right

Tension (0.0 - 1.0)

Controls **curve smoothing** in Line mode.

- **Low values**
Angular, sharp transitions between bands
- **High values**
Smooth, rounded curves

Important behavior

- Tension affects **Line mode only**.
- It has no visible effect in Bar or Segment modes.

Segments (1 - 100)

(Segment mode only)

Defines the **number of discrete blocks** per frequency band.

- Low values → coarse segmentation
- High values → fine segmentation

Segment Gap (0 - 100)

(Segment mode only)

Controls the **spacing between segments**.

- Low values → thin gaps, visually thicker segments
- High values → wide gaps, thinner-looking segments

Frequency analysis parameters

(Apply to Line, Bar, and Segment modes)

Bands (2 - 100)

Controls the **number of frequency bands** analyzed and displayed.

- Fewer bands → coarse frequency resolution
- More bands → finer spectral detail

Low Frequency (20 Hz - 19 000 Hz)

Defines the **lowest frequency** included in the spectrum.

- Raising this value removes bass content from the display

High Frequency (1020 Hz - 20 000 Hz)

Defines the **highest frequency** included in the spectrum.

- Lowering this value limits the display to mid or low frequencies

Threshold (-60 to 0 dB)

Sets the **minimum signal level** required to be displayed.

- **Lower values**
More sensitivity; quiet frequencies appear
- **Higher values**
Less sensitivity; low-level noise is suppressed

Keyframes

This filter does **not** support keyframes.

All animation is driven entirely by the **audio signal**.

Visual characteristics

- Real-time frequency response
- Vertical frequency-based animation
- Deterministic for a given audio source
- Best evaluated during playback

Recommended use cases


- Music visualizers
- Educational audio demonstrations
- Podcast or voice-over visualization
- Audio-reactive motion graphics

Limitations

- Requires audio on a **video clip**
- Not a calibrated measurement tool
- No logarithmic/linear scale selection
- UI elements are rendered into the video




The **Audio Waveform Visualization** filter allows the user to display a waveform of the audio on top of the video.



The clip or track that the filter is applied to must have audio. The filter can not show the audio from a different track or clip. If you have music in an audio track and want the visualization to appear on the video from video tracks you should add the filter to the **Output** track in the top, left corner of **Timeline**.



The **Blend Mode** filter lets you override the way a clip blends with the the bottom video track. The default blending mode follows the video track’s blending mode, which is set in **Properties** with a default of **Source Over** (simple alpha compositing).



Most of the options are the standard **Porter/Duff** blend modes:

- **Source Over** is the default, standard way layers stack: an upper layer is drawn over a lower layer.
- **Destination Over**: if for some reason you are layering tracks to be composed in a bottom-up instead of top-down fashion. For example, you have the clip with the transparent background under the background clip on the timeline.
- **Lighten** is useful if you have a clip with a **black** background and no transparency. It makes black transparent without needing a chroma key filter.
- **Darken** is useful if you have a clip with a **white** background and no transparency. It makes white transparent without needing a chroma key filter.

Quick Reference

Category	Blending Mode	Core Function	Primary Practical Use
Darkening	Multiply	Darkens the base layer by multiplying values. Ignores white pixels.	Compositing Dark Textures/Shadows. Best for

Category	Blending Mode	Core Function	Primary Practical Use
			removing white backgrounds from assets.
	Color Burn	Darkens more aggressively than Multiply. High contrast, rich shadows.	Deepening Shadows and creating a moody, high-contrast cinematic tone.
	Darken	Compares and keeps only the darker pixels from either layer.	Knocking out white backgrounds on graphics or text (similar to Multiply but simpler).
Lightening	Screen	Brightens the base layer by multiplying inverted values. Ignores black pixels.	Compositing Light Effects (fire, flares, smoke, etc.) over black backgrounds.
	Color Dodge	Dramatically brightens and increases saturation. Creates a luminous glow.	Intense Glare and Spotlight Effects. Pushes highlights towards white with energy.
	Lighten	Compares and keeps only the lighter pixels from either layer.	Knocking out black backgrounds on high-contrast visual effects.
Contrast	Overlay	Combines Screen (on light areas) and Multiply (on dark areas).	Non-Destructive Contrast Enhancement and adding seamless textures/color grades.
	Soft Light	Applies soft contrast and subtle lighting effects.	Subtle atmospheric effects and gentle warming/cooling of the scene.
	Hard Light	Applies strong contrast and dramatic lighting effects.	Dramatic Lighting or compositing intense, high-contrast black/white textures.
Component	Color	Applies the blend layer's Hue and Saturation, keeping the base layer's Luminosity.	Color Grading: Quickly applying a specific tone (e.g., teal/orange) while preserving detail.
	Hue	Applies only the blend layer's Hue (color), keeping base Saturation and Luminosity.	Changing a specific color (e.g., a shirt) without affecting its shading or brightness.
	Luminosity	Applies the blend layer's Luminosity (brightness), keeping the base layer's Hue and Saturation.	Sharpening and Contrast Control separate from color (to avoid color artifacts).
Creative	Difference	Subtracts the lighter color from the darker color,	Creative Glitch Effects and precise edge detection (often

Category	Blending Mode	Core Function	Primary Practical Use
		causing inversion.	used with offset layers).

More Information

Blending modes are categorized by how they calculate the final color by comparing the pixels of the top layer (blend) with the pixels of the bottom layer (base).

Isolation and Compositing Modes (Darken / Lighten)

These modes simplify compositing by isolating only the darkest or lightest parts of the blend layer.

Blending Mode	Practical Use	Example
Darken	Removing Highlights/Combining Shadows: Compares the pixels and keeps only the <i>darker</i> one. This is perfect for compositing elements where you only want the shadows or dark parts of the top layer to show.	You have a logo or text graphic with a white background and dark text. Place it over your video and set it to Darken —the white background disappears, and only the dark text remains, quickly knocking out a light background.
Lighten	Removing Shadows/Combining Highlights: Compares the pixels and keeps only the <i>lighter</i> one. This is the opposite of Darken.	Use this to overlay a series of different spark or lens flare effects (that typically have black backgrounds) onto your footage. The black background disappears, and only the bright flares/sparks are visible because they are lighter than the base video.

Color Dodge

Color Dodge is a member of the **Lightening Group**. It creates a much stronger, more luminous brightening effect than Screen by dramatically increasing the brightness and saturation of the base layer.

Practical Use	Explanation
Intense Glare and Light Flares	Use a soft, bright color (like a vibrant orange, yellow, or white) on the blend layer and set it to Color Dodge . This creates a powerful, focused light bloom or lens flare that looks highly energetic and realistic, especially over existing highlights.

Practical Use	Explanation
Cinematic Focus/Glow	Create a shape or brush stroke over a face or object using a subtle, bright color. Color Dodge will make the area glow brightly, drawing the viewer's eye and creating a dreamy or heroic spotlight effect.
Saturated Color Grading	Applying a solid, bright color layer (e.g., a fiery red-orange) at a low opacity and setting it to Color Dodge will inject intense saturation into the mid-tones and highlights, giving the entire scene a vibrant, almost painted look.

Color Burn

Color Burn belongs to the **Darkening Group**. It creates a stronger, deeper darkening effect than Multiply, giving the result an intense, rich mood, especially in the shadows.

Practical Use	Explanation
Deepening Shadows and Blacks	Use Linear Burn on a duplicate clip or a blurred copy of your footage. It significantly enriches the shadows and dark areas, instantly giving the video a deeper, moodier, and more dramatic contrast, especially useful in horror or dramatic genres.
Gritty Texture Compositing	When applying a texture (like dirt, grime, or deep film grain) that you want to aggressively darken the underlying image, Linear Burn is far more effective than Multiply. It burns the dark texture into the base layer with a higher contrast and less saturation than Color Burn.
High-Contrast Cinematic Tone	Use a solid, dark color (e.g., a deep teal or maroon) on the blend layer at a low opacity. Setting it to Linear Burn infuses that color into the shadows and dark mid-tones while intensely deepening the black points, creating a strong, stylized color grade.

Contrast and Lighting Modes (Overlay, Hard Light & Soft Light)

These modes modify the contrast and brightness of the base layer using the blend layer's color values.

Blending Mode	Practical Use	Example
Overlay	Non-Destructive Contrast Enhancement: It's perfect for adding contrast, texture, and color grading in a non-destructive way.	Duplicate your clip and set the top layer to Overlay . This instantly increases the overall contrast and saturation, giving the image a punchy, professional look. You can then lower the opacity to fine-tune the intensity.

Blending Mode	Practical Use	Example
Hard Light	Dramatic Lighting/Texture Overlays: Creates a much more intense effect than Soft Light. It's ideal for adding dramatic lighting effects or quickly compositing detailed textures.	Use a black-and-white texture (like scratched film or dust) over a clip. Set the blend layer to Hard Light —white areas will brighten the base, and black areas will darken it, adding instant grit.
Soft Light	Subtle Lighting/Atmosphere: Excellent for non-destructively brightening or darkening a scene, adding a soft glow, or applying a colored atmosphere (e.g., a warm, hazy glow) without blowing out highlights or crushing shadows.	Place a solid color layer (e.g., warm orange) or a gradient over your footage and set it to Soft Light to mimic sunset lighting.


Subtractive Modes (Difference / Exclusion)


These modes are used for creative, often psychedelic, or glitch effects by subtracting one color from the other.

Blending Mode	Practical Use	Example
Difference	Color Inversion and Alignment Check: It inverts the colors of the base layer based on the brightness of the blend layer. If the blend layer is pure black, nothing happens; if it's white, the base layer is fully inverted.	Use a copy of your base video as the blend layer and offset it slightly. Setting the blend mode to Difference creates a colorful, high-contrast, edge-detection effect, popular for abstract or glitch aesthetics.
Exclusion	Softer Inversion Effects: Similar to Difference but with lower contrast, often used to create a more dreamy or muted inverted look.	It can be used for artistic color shifts and subtle, high-contrast text overlays, where the color changes depending on the background video.

Versions before 25.10

In previous versions the blending options were a little different because it was using a **component** that could not be updated to support 10-bit video.





In this version **Over** is equivalent to **Source Over**.


Component Modes (Hue / Color / Luminosity)

These modes are used for advanced color grading and correction by manipulating the separate components of a color (Hue, Saturation, Luminosity).

Blending Mode	Practical Use	Example
Hue	Changing Colors without Affecting Brightness: Takes the hue (color) of the blend layer and applies it to the base layer, retaining the base layer's saturation and luminosity.	To change the color of a specific object (e.g., a shirt or car) without affecting its lighting or shadows, mask the object, apply a new color layer above it, and set the color layer to Hue .
Color	Applying a Specific Tone/Tint: Applies both the hue and saturation of the blend layer, while retaining the base layer's luminosity (brightness). This is the best mode for coloring black-and-white footage or adding a precise color grade.	Place a color adjustment layer over your footage (e.g., a cool blue) and set the blend mode to Color to apply a professional-looking cinematic blue tone across the whole image.
Luminosity	Sharpening and Contrast Control: Applies the luminosity (brightness) of the blend layer while retaining the base layer's hue and saturation. This mode is mainly used for manipulating contrast separate from color.	To increase sharpening without introducing color artifacts, create a high-pass filter effect on the blend layer and set the mode to Luminosity or Soft Light .

The **Blur Box** filter applies a **box blur to the entire frame**.

The blur strength is controlled independently in the horizontal and vertical directions, allowing directional blur effects.



This filter operates **spatially** and affects the whole image, not a selected region.

Parameters

Width (0.0–100.0)

Controls the amount of **horizontal blurring**.

- Lower values
Little or no horizontal blur
- Higher values
Stronger blur **along the X axis**

This parameter blurs pixels **left to right** across the frame.

Height (0.0–100.0)

Controls the amount of **vertical blurring**.

- Lower values
Little or no vertical blur
- Higher values
Stronger blur **along the Y axis**

This parameter blurs pixels **top to bottom** across the frame.

Keyframes

Both **Width** and **Height** parameters can be **keyframed**, allowing blur strength to change over time.

This enables:

- Progressive blur or unblur effects
 - Directional motion-style blurring
 - Animated emphasis or abstraction
-

Blur alpha

When enabled, the blur is also applied to the **alpha channel**.

- Disabled
Only RGB color data is blurred; transparency remains unchanged
- Enabled
Alpha values are blurred as well, softening transparency edges

This is useful when blurring images or overlays that contain semi-transparent areas.

Parameter interaction

- **Width** controls blur strength horizontally
- **Height** controls blur strength vertically

Using both together produces a uniform box blur, while using only one produces a **directional blur**.

Visual characteristics

Typical effects include:

- Uniform blur applied across the entire frame
- Directional blur when only one axis is used
- Softening of detail and texture
- Transparency smoothing when Blur alpha is enabled

Recommended use cases

- General-purpose blur
- Directional blur effects
- Stylized abstraction
- Blurring overlays or graphics with transparency
- Animated blur transitions using keyframes

Creative uses

Because Blur Box applies independent horizontal and vertical blurring across the entire frame, it can be used for several stylized effects beyond simple softening:

- **Directional motion blur**
Using only Width or only Height creates horizontal or vertical motion blur.
- **Defocus or dream effects**
Moderate values on both axes soften the image uniformly.
- **Animated blur transitions**
Keyframing Width and/or Height allows blur to ramp in or out for transitions or emphasis.
- **Glow and highlight enhancement**
When combined with blending modes on a duplicated layer/track, Blur Box can be used to create glow or bloom effects.
- **Smear or analog-style distortion**
Strong directional blur can simulate smear or degraded analog visuals.

These effects rely on keyframing and filter order rather than masking.

Minimalist cheat-sheet

Effect	Width	Height	Keyframes
Motion blur	High	0	Optional
Defocus	Medium	Medium	Optional
Speed ramp	0 → High → 00		Yes
Glow	Medium	Medium	No
Smear	High	0	Optional


Limitations

- No region or mask control
- No feathering or falloff control
- Blur strength only, no quality or radius options
- Blur alpha cannot be keyframed



Applies a **non-linear blur** where pixel influence decreases exponentially with distance.

Unlike **Gaussian blur**, which spreads blur evenly in all directions, exponential blur concentrates its effect near edges and high-contrast transitions, producing a softer yet more localized smoothing.



This filter is designed for **perceptual softening** rather than precise optical simulation.

Parameters

Amount (0.0 - 100.0%)

Controls the **strength of the exponential blur**.

- **0.0%**
No blur is applied.
- **Low values (5% - 20%)**
Subtle softening. Fine detail is reduced while edges remain relatively defined.
- **Medium values (20% - 60%)**
Noticeable blur. Texture and small features are smoothed, but the image does not spread as widely as with Gaussian blur.
- **High values (60% - 100%)**
Strong blur. The image becomes soft and diffuse, with reduced haloing compared to Gaussian blur at similar strength.

Important behavior

- Blur intensity increases rapidly at low values.
- The effect is **not radius-based**; Amount controls perceptual strength rather than distance in pixels.

Note:

Exponential blur fades quickly with distance, so blur feels concentrated rather than spread out.

Keyframes

The **Amount** parameter can be keyframed.

This allows:

- Gradual focus shifts
- Soft transitions
- Animated emphasis or de-emphasis of detail

Visual characteristics

- Smooth, localized softening
- Less edge haloing than Gaussian blur
- Reduced spread into neighboring areas
- More “matte” appearance at high values

Comparison with other blur filters

Gaussian Blur

- Even, symmetrical blur
- Radius-based
- Spreads blur uniformly
- Commonly used for defocus or general smoothing

Exponential Blur

- Non-linear falloff
- Concentrates blur near edges
- Less uniform spread
- More perceptual than physical

Box blur

- Harder edges
- Less natural transitions
- Performance-oriented

Summary of differences

Blur type	Spread behavior	Typical use
Gaussian	Uniform, wide	Optical defocus
Exponential	Rapid falloff, local	Perceptual softening
Box / Fast	Hard, blocky	Speed or stylization

Recommended use cases


- Softening harsh digital edges
- Subtle beauty or matte effects
- Background diffusion without strong haloing
- Transitional blur animations

Limitations

- Not physically accurate lens simulation
- Single control only
- Can flatten contrast at high values
- Not suitable where precise blur radius is required



A Gaussian blur smooths the image by averaging pixels using a bell-shaped (Gaussian) distribution, producing a softer and more natural blur than a box blur.



This filter operates **spatially** and affects the whole image uniformly.

Parameter

Amount (0.0–100.0)

Controls the **strength of the Gaussian blur**.

- Lower values
Subtle softening with minimal detail loss
- Higher values
Strong blur with significant reduction of fine detail

This parameter defines how widely pixel values are averaged around each pixel.

Keyframes

The **Amount** parameter can be **keyframed**, allowing blur strength to change over time.

This enables:

- Blur-in / blur-out transitions
- Progressive defocus effects
- Emphasis or de-emphasis over time

Blur alpha

When enabled, the blur is also applied to the **alpha channel**.

- Disabled
Only RGB color data is blurred; transparency edges remain sharp
- Enabled
Alpha values are blurred as well, softening transparency edges

This is useful when blurring images, titles, or overlays that contain semi-transparent regions.

Visual characteristics

Typical effects include:

- Smooth, natural softening
- Reduction of noise and fine detail
- No directional bias
- Softened transparency edges when Blur alpha is enabled

Gaussian blur produces fewer hard artifacts than simple box blur, but is generally more computationally expensive.

Creative uses

- **Defocus or depth-of-field simulation**
Uniform blur to suggest loss of focus.
- **Soft transitions**
Keyframed blur changes smoothly over time for gentle scene changes.
- **Glow or bloom effects**
When combined with blending modes on a duplicated track.
- **Background softening**
Reducing detail to emphasize foreground elements.


Limitations

- Full-frame effect only (no region or mask)
- No directional control
- Blur alpha cannot be keyframed
- High values may remove significant detail



Applies a **low-pass spatial filter** to the video clip, attenuating high-frequency detail while preserving low-frequency structures.

In practical terms, it removes fine texture and noise before it removes larger shapes and edges.



This filter is based on the same conceptual idea as a **low-pass filter in signal processing**: allow slow variations to pass, suppress rapid changes.

What “low pass” means

In images:

- **Low frequencies** correspond to:
 - Large shapes
 - Smooth gradients
 - Broad lighting changes
- **High frequencies** correspond to:
 - Fine detail
 - Texture
 - Noise
 - Sharp edges

A **low-pass filter** reduces high-frequency content while keeping low-frequency information intact.

Note:

Low-pass blur removes detail first, not structure.

Parameters

Amount (0.0 - 100.0%)

Controls how strongly high-frequency detail is suppressed.

- **0.0%**
No filtering; the image is unchanged.
- **Low values (5% - 20%)**
Fine texture and noise are reduced. Edges remain mostly intact.
- **Medium values (20% - 60%)**
Noticeable softening. Small details disappear, larger forms remain readable.
- **High values (60% - 100%)**
Strong smoothing. The image becomes flat and diffuse, with minimal fine structure.

Important behavior

- The effect increases rapidly at low values.
- Amount controls **frequency suppression**, not a blur radius in pixels.

Keyframes

The **Amount** parameter can be keyframed.

This enables:

- Progressive smoothing or reveal of detail
- Temporal separation of detail and structure
- Transitions between sharp and abstract states

Visual characteristics

- Smooth, uniform softening
- Reduced fine texture and noise
- Edges fade gradually rather than haloing
- Less “spread” than Gaussian blur at comparable strength

Comparison with other blur filters

Gaussian Blur

- Uniform spatial blur
- Radius-based
- Spreads pixels evenly in all directions
- Produces visible haloing at higher values

Low Pass Blur

- Frequency-based smoothing
- Suppresses detail before structure
- More controlled edge degradation
- Less haloing, more matte appearance

Exponential Blur

- Rapid falloff from edges
- Perceptual, localized softening
- Not frequency-oriented

Summary comparison

Blur type	Core behavior	Typical use
Gaussian	Even spatial spread	Defocus, general blur
Low Pass	High-frequency suppression	Noise/detail reduction
Exponential	Localized falloff	Perceptual softening

Recommended use cases

- Reducing fine digital noise
- Preprocessing before Posterize or Dither
- Softening textures while preserving shapes
- Creating abstract or painterly looks
- Preparing footage for heavy compression


Limitations

- Not a physical lens simulation
- Single control only
- Can flatten contrast at high values
- Not suitable when edge precision must be preserved



Creates a blurred padding area **outside a defined rectangular region**, using a scaled version of the original image.

The interior of the rectangle remains unmodified, while the surrounding area is filled with a **zoomed and blurred extension** of the same frame.



This filter is primarily intended to replace black bars or empty space by generating a visually consistent background derived from the clip itself.

Note:

This filter is commonly used to create blurred backgrounds for **vertically oriented videos presented in a landscape frame** by extending and softening the original image, avoiding lateral black bars without manual compositing.

Parameters

Position (X, Y)

Controls the **center position of the unblurred rectangle**, expressed as normalized coordinates.

- Can be adjusted numerically or directly from the preview overlay.
- Moving the position changes **which part of the image remains sharp**.
- The blurred padding always occupies the area **outside** this region.

Typical use:

- Centered position for pillarbox/letterbox filling
- Offset position to protect a subject that is not centered

Size (Width, Height)

Controls the **size of the unblurred rectangular area**.

- Can be adjusted numerically or via the preview overlay.
- The rectangle defines the **protected region** that remains untouched.
- Everything **outside** this rectangle is replaced by the blurred padding.

Important behavior

- The blurred area is derived from a **scaled (zoomed) version** of the original image.
- This prevents visible seams between the sharp and blurred regions.

Note:

Size defines what stays sharp, not what gets blurred.

Blur (0.00 - 10.00%)

Controls the **strength of the blur applied to the padding area only**.

- **0.00%**
No blur; the padding is a scaled copy of the image.
- **Low values (0.50 - 2.50%)**
Strong blur already. Image details in the padding become hard to identify.
- **Medium to high values (>2.50%)**
Extremely soft background. Additional increases have diminishing visual impact.

Important behavior

- Blur affects **only the padded area**, never the protected rectangle.
- The blur radius increases very rapidly; small values produce large perceptual changes.

Note:

Blur is intentionally coarse; it is designed to hide detail, not to be finely tuned.

Keyframes

- **Position** and **Size** can be keyframed.
- **Blur** cannot be keyframed.

Keyframing enables:

- Tracking a subject while maintaining a sharp foreground
- Dynamic reframing for vertical or square output
- Animated transitions between protected regions

Parameter interaction


- **Position + Size** define the protected (sharp) area.
- **Blur** affects only the generated padding outside that area.
- The padding is always a **zoomed version of the same frame**, ensuring color and motion consistency.

Visual characteristics

- Sharp rectangular region overlaid on a blurred background
- No visible edges or hard seams between regions
- Background blur follows the motion and color of the source clip
- Padding scales automatically with the frame

Recommended use cases

- Converting horizontal footage for vertical formats
- Filling black bars without cropping content
- Social media delivery (Shorts, Reels, Stories)
- Quick alternative to mask-based background generation




Limitations

- Blur strength has a narrow useful range
- Padding cannot use an external or custom background
- Shape is limited to a rectangle / square
- Blur behavior is not physically accurate

Adjusts the **overall luminance** of the image / video by scaling pixel intensity values.

This filter makes the image uniformly darker or brighter without targeting contrast, midtones, or color balance.



Brightness operates as a **global gain** on luminance rather than a tonal remapping.

Parameters

Level (0.0 - 200.0%)

Controls the **brightness multiplier** applied to the image.

- **100.0%**
No change.
- **Below 100%**
Darkens the image.
- **Above 100%**
Brightens the image.

Observed behavior

Although applied globally, the adjustment is **most noticeable in shadows and darker regions**.

Highlights reach clipping more quickly and therefore show less visible change at high values.

Note:

Brightness lifts or lowers the signal as a whole; dark areas respond more visibly than bright ones.

Keyframes

The **Level** parameter can be keyframed.

This allows:

- Gradual exposure changes
- Fade-ins or fade-outs via luminance
- Temporal brightness modulation

Visual characteristics

- Uniform luminance scaling
- No selective midtone control
- No contrast expansion or compression
- Color hue is preserved, but saturation may appear altered due to clipping

Comparison with related controls

Brightness vs **Contrast**

- **Brightness** shifts overall luminance
- **Contrast** changes the difference between dark and light regions

Brightness vs **Levels / Gamma**

- **Brightness** is a simple gain
- **Levels / Gamma** remap tonal ranges and midtones more precisely

Note:

Use Brightness for coarse exposure correction, not fine tonal shaping.

Recommended use cases

- Quick exposure compensation
- Matching shots with small luminance differences
- Simple fades to black or white
- Pre-adjustment before contrast or gamma corrections


Limitations

- Can clip highlights or crush shadows
- No protection for midtones
- Not color-aware
- Not suitable for precise exposure correction




A “choppy video filter” refers to an effect that intentionally makes a video appear uneven or stuttery, like it’s skipping frames. This look mimics a low frame rate or gives the impression of stop-motion animation. It’s often used creatively to add a sense of retro, glitchy, or dramatic style to videos.

You can see the difference between the original video and the same clip with the filter, but the frame rate reduced by half.





Preserves a selected color while desaturating all other colors in the image. The result is a mostly monochrome image in which one color (or a narrow range of colors) remains visible.



This effect is often referred to as **selective color** in still image editing and is used for emphasis or visual isolation rather than color correction.

Parameters

Color

Selects the **reference color** to preserve.

- Can be chosen:
 - From the color palette
 - Directly from the image using the picker in the preview
- The chosen color defines the **center of the preserved color range**.

Important behavior

- The filter operates in color distance, not by named hues.
- Picking a color from the image usually produces the most predictable result.

Distance (0.0 - 100.0%)

Controls the **tolerance around the selected color**.

- **0.0%**
Only pixels extremely close to the selected color are preserved.
Most of the image becomes grayscale.
- **Low values (5% - 20%)**
A narrow color range is kept. Subtle variations of the selected color may be lost.

- **Medium values (20% - 50%)**

A broader range of similar hues is preserved. This is often the most usable range.

- **High values (50% - 100%)**

Many colors remain unaffected. At high values, the effect becomes weak or nearly invisible.

Note:

Distance controls how strict the color selection is, not how strong the desaturation feels.

Keyframes

The **Distance** parameter can be keyframed.

This enables:

- Gradual reveal or isolation of a color
- Animated transitions between monochrome and color
- Emphasis shifts over time

Visual characteristics

- Non-selected colors are converted to grayscale
- Selected color range remains fully colored
- No change to geometry or luminance
- Sharp transitions may appear if Distance is too low

Comparison and context

Chroma Hold vs Color Grading

- **Chroma Hold** isolates color by similarity
- **Color grading** adjusts colors globally or tonally

Relation to “Selective Color” in image editors

- Conceptually similar: one color is preserved, others are muted
- Chroma Hold operates dynamically and can be animated over time

Important distinction:

This filter does not recolor pixels; it selectively desaturates them.

Recommended use cases

- Highlighting a subject or object by color
- Creating dramatic or stylized visuals
- Drawing attention in narrative or documentary footage
- Transitional effects between color and monochrome


Limitations

- No control over edge softness or feathering
- Can produce hard boundaries between color and grayscale
- Sensitive to lighting changes and color noise
- Not suitable for precise color isolation in complex scenes



Introduced in version [16.01](#)

This filter makes certain colors of your video transparent. Thus, you need to have something under the video to see its effect. You can also use it after a **Mask** filter and before **Mask: Apply** if you want to make other filters affect only certain colors, for example for secondary color correction.



Key color

This is the color to select. This is the center point of the selected color subspace. You can use the dropper to pick a color from the screen. When using the dropper, drag a small rectangle to average the color values within the rectangle.

Color space

This specifies in which of the two color spaces (**RGB**, **HCI**) the delta controls will work. HCI is separated luma/chroma (a cylindrical “Hue Chromacity Intensity”) space. It suffers from chroma subsampling, and will give less sharp results than RGB.

NOTE: Chroma subsampling is not a fault of this plugin. Frei0r plugins work in RGB 4:4:4. If you supply it with a truly 4:4:4 video, it will produce perfectly sharp alpha from both color spaces.

NOTE: HCI is slow, because it has to calculate the arctangent and hypotenuse for each pixel.

Red/Green/Blue delta

Hue/Chroma/Intensity delta

These three parameters determine the size of the color subspace along each axis. Bigger value means bigger tolerance on that axis. For example, setting a high intensity delta (in HCI), will allow the selection of a specific color in both light and

shadows, but will also discard most of the high-bandwidth luma signal, making the selection less spatially accurate.

Shape

This determines the shape of the color subspace. Options are: **Box**, **Ellipsoid** or **Diamond**. Box is the biggest of them (by volume) and diamond the smallest. Imagine an octahedron inscribed inside an ellipsoid, which is in turn inscribed in a box. The tips of the diamond touch the ellipsoid, and the box, at the center of the sides of the box.

Edge

This has five options: **Hard**, **Fat**, **Normal**, **Thin** and **Slope**. **Hard** means the alpha channel will be thresholded to two values only, fully opaque and fully transparent. This is mainly useful for keying. The **Slope** mode is also intended for keying. It is particularly useful when keying partly transparent or fuzzy (like hair) objects. The remaining options (**Fat**, **Normal** and **Thin**) create a gradual transition between transparent and opaque. Alpha will be a function of the difference of color from the selected color, with cutoff at the delta points. The fatter the choice, the more the selected areas are filled towards the rim. This is useful with alpha controlled color adjustment tools.

Slope

When edge mode is set to **Slope**, this parameter controls a gradual transition between opaque and transparent. 0% gives a hard-edged key, increasing the value of this parameter increases the range of colors around the selected color, that will be rendered partially transparent.

Operation

This determines the way in which the alpha channel will be written. These options combine the current selection with the pre-existing alpha of the source material. This way complex selections can be built. The **Minimum**, **Maximum**, **Add** and **Subtract** options allow cascading of other **Alpha Channel** or **Mask** filters; **Overwrite** does not.

Invert


When ON, the selected color will be transparent, as normally used with keying. When OFF (default) the selected color will be opaque, for example for alpha controlled adjustment of that color only.



Introduced in version **16.01**

Location: Filters > Video > Chroma Key: Simple

If you have an item in your playlist or timeline that has a uniform background, you can use the “Chroma Key: Simple” filter to remove it. This is commonly used to remove a green screen to superimpose a picture or video onto a background that was captured or acquired separately. That means the asset has to be placed on any track but V1, and a background has to be placed on a lower track. So for example some green screen footage is placed on V2, and a background is placed on V1.




There are two options:


The **Key colour** sets the background colour. The eye dropper icon can be used to activate the colour picker. By clicking now on the background, the colour will be selected as the key colour.



The **Distance** adjusts which shades or colours are treated like the key colour. This is especially important for similar colours that are present in the background as well as in the foreground. If the colours are very similar, only few percentages decide whether both are keyed out or not. This is also useful for retaining fine details like hairs, which might be keyed out as well.

Preset is for saving a set of key colour and distance.





In this example the colour of the text is very similar to the background. With a small percentage the background is keyed out, but the text remains. Increasing the value gradually reduces the opacity of the text.






In this example one strand of hair is partially removed beyond a certain percentage.



In this example the lighting is subpar, and the background consists of various shades of green. With finding a key colour in the middle of these shades and adjusting the distance, the background can be removed to large parts. The lighting is partly responsible for the green screen remains behind the hair. Using only this filter will not remove it completely without affecting other parts.




If “Chroma Key: Simple” does not lead to satisfying results, **“Chroma Key: Advanced”** and/or the key spill filters (**“Key Spill: Simple”** and **“Key Spill: Advanced”**) might help.



Applies **tonal color correction** by independently adjusting the color balance of **shadows, midtones, and highlights**.

This filter allows precise control over the emotional tone, visual coherence, and stylistic identity of footage.



Color Grading does **not** correct exposure errors by itself; it reshapes how colors are distributed across tonal ranges.

What color grading is (and why it matters)

Color grading is the process of **intentionally shaping color and tone** to support mood, narrative, and visual consistency.

In cinema, it is a fundamental storytelling tool:

- **Spaghetti westerns (e.g. Sergio Leone)**
Heavy warm highlights, cool or desaturated shadows, strong tonal separation
→ Emphasizes dust, heat, tension, and moral contrast
- **Blade Runner (1982)**
Strong color bias in shadows and highlights (teal, amber, neon hues)
→ Reinforces atmosphere, futurism, and emotional isolation
- **In the Mood for Love (2000)**
Selective midtone grading with warm chromatic bias and subdued highlights, combined with controlled shadow density
→ Supports a sense of closeness, repetition, and emotional tension without relying on overt contrast

In all cases, the look is not “natural”; it is **designed**.

Main controls

Each tonal range has the **same control structure** and behavior.

Shadows (Lift)

Affects **dark areas** of the image.

- Controls shadow color bias
- Influences perceived depth and mood
- Overuse can crush blacks or introduce color noise

Midtones (Gamma)

Affects **mid-level brightness and color**, where most image detail lives.

- Primary tool for skin tones
- Strongly affects overall look
- Small changes have large perceptual impact

Highlights (Gain)

Affects **bright areas** of the image.

- Controls highlight color cast
- Shapes perceived lighting temperature
- Can easily clip if pushed too far

Color wheel (common to all three)

Each tonal range uses an identical **color wheel**.

- The **black dot** in the middle represents neutral (no color bias)
- **Moving the dot** toward a color:
 - Pushes that color into the selected tonal range
 - Automatically adjusts the corresponding RGB values

Important behavior

- The wheel and R / G / B inputs are fully linked
- Changing one updates the others

Vertical bar (luminance strength)

The vertical bar next to each wheel controls the **intensity of the adjustment**.

- **Top (white)**
Maximum influence
- **Center (default)**
Balanced influence

- **Bottom (black)**
Decreases luminance for the selected tonal range, resulting in darker output.

This bar directly affects the **range and sensitivity** of the RGB input fields.

RGB input fields (R / G / B)

Provide **numerical control** over the same adjustment defined by the wheel.

- The available range depends on the vertical bar position
- Default center position typically limits values to smaller ranges
- Moving the vertical bar expands or contracts the usable range

Important behavior

- RGB values do not represent absolute color
- They represent **offsets applied to the selected tonal range**

Note:

RGB inputs refine the wheel adjustment; they do not replace it.

Keyframes

All three tonal controls are **keyframeable**.

This enables:

- Gradual mood shifts
- Day-to-night color transitions
- Emotional emphasis tied to narrative beats

Visual characteristics

- Independent color control by tonal range
- No geometry or sharpness changes
- Can introduce stylized or cinematic looks
- Errors are more visible on skin tones and neutral areas

Recommended use cases


- Establishing a cinematic look
- Matching shots from different cameras
- Enhancing mood and atmosphere
- Stylized genre looks (western, noir, sci-fi)

Limitations

- Not a replacement for exposure correction
- Easy to overdo; subtlety is critical
- Can reveal noise in shadows
- Requires calibrated viewing for precision work



Adjusts the **difference between dark and light areas** in the image. Increasing contrast pushes dark values darker and bright values brighter, while decreasing contrast compresses the tonal range toward midtones.



Contrast does **not** change overall brightness; it redistributes how luminance values are separated.

What contrast means

In image terms:

- **Low contrast**
Dark and light areas are closer together
→ Flat, soft, or washed-out appearance
- **High contrast**
Dark and light areas are further apart
→ Strong separation, punch, and visual clarity

Contrast defines how clearly forms, textures, and edges are perceived.

Note:

Brightness moves everything up or down; contrast spreads values apart or squeezes them together.

Parameters

Level (0.0 - 100.0%)

Controls the **strength of contrast expansion or compression**.

- **0.0%**
Minimal contrast. The image appears flat, with reduced distinction between dark and light areas.

- **Low values (10% - 30%)**
Soft contrast. Useful for foggy, low-key, or muted looks.
- **Medium values (40% - 70%)**
Natural contrast range for most footage.
- **High values (70% - 100%)**
Strong contrast. Blacks deepen, highlights brighten, and midtones become more pronounced.

Important behavior

- Increasing contrast can **clip highlights** and **crush shadows**.
- Decreasing contrast can reduce perceived sharpness.

Keyframes

The **Level** parameter can be keyframed.

This allows:

- Gradual build-up of visual intensity
- Transitions between flat and dramatic looks
- Contrast fades independent of brightness

Visual characteristics

- Affects luminance distribution only
- Preserves color hue, but saturation may appear stronger at high contrast
- Emphasizes edges and texture
- Impacts perceived depth and clarity

Relation to black, white, light, and dark

- **Blacks**
Higher contrast pushes blacks closer to pure black.
- **Whites**
Higher contrast pushes highlights closer to pure white.
- **Midtones**
Become more separated, increasing perceived detail.

Lower contrast compresses all values toward mid-gray, reducing distinction.

Note:

Contrast controls the distance between black and white, not their position.

Comparison with related controls

Contrast vs Brightness

- **Brightness** shifts the entire image lighter or darker.
- **Contrast** changes the separation between light and dark values.

Contrast vs **Levels / Gamma**

- **Contrast** is global and symmetrical.
- **Levels / Gamma** allow targeted control of shadows, midtones, or highlights.


Recommended use cases

- Restoring punch to flat footage
- Enhancing depth and separation
- Matching shots from different sources
- Establishing dramatic or stylized looks

Limitations

- Easy to overdo
- Can hide detail in shadows and highlights
- Not selective; affects the entire image uniformly
- Not a replacement for proper exposure correction

This filter lets you transform (change the size, position, and shape) to fit inside a shape defined by four points. This is not limited to a 4-sided polynomial (quadrilateral) because the corners can be moved beyond the lines between other points to create weird effects. A common use case for this is to fit one video to a rectangular area within another video. For example, replacing a sign or window, or placing something on a wall.



This screenshot is a super simple example, but typically you will use this on a timeline with more than one video track where the video with this filter is on a video track above another video below to composite them. (One could also use this in conjunction with **Mask: Simple Shape** and **Mask: Apply** video filters to composite the distorted video on top of itself with additional filters applied only to the corner pinned video by stacking them before **Mask: Apply**.)

This filter also features a video user interface (VUI) meaning the corners are shown with circles on the video preview area, and you can drag these circles to interactively define the four **Corner X Y** parameters. These corner controls will snap to the player grid if that feature is turned on.

Stretch X Y parameters will do a non-linear scaling horizontally (X) or vertically (Y) to help create perspective when the plane onto which you are projecting is not facing parallel with the camera.

Interpolator is the algorithm used to define how to do sub-pixel rendering and affects the quality and speed. **Nearest Neighbor** is fastest but the worst quality. **Lanczos** is probably the best quality but also slowest. The default **Bicubic Smooth** represents a good combination of medium quality and speed.

Alpha Operation controls how the **alpha channel** created by this filter combines with the alpha channel that is already in the video. When not set to **Opaque**, this filter's alpha channel makes everything outside of the lines connecting the corners transparent while the rest is opaque.

Feathering controls the softness of the edge between opaque and transparent.


For **Simple Keyframes**, this filter uses only the following parameters:

- **Corner 1**
- **Corner 2**
- **Corner 3**
- **Corner 4**
- **Stretch**
- **Feathering**

Note: added in version 20.04.12



Applies a **circular crop mask** to the clip, keeping a circular region visible while replacing everything outside the circle with either a solid color or transparency.



This filter is a **masking tool**, not a geometric crop in the traditional rectangular sense.

General behavior (important)

- The filter defines a **circle centered on the frame**.
- The **inside of the circle** shows the original clip.
- The **outside of the circle** is either:
 - Filled with a solid color, or
 - Made transparent (if enabled)

The circle expands or contracts based on the Radius setting.

Parameters

Radius (0.0 - 100.0%)

Controls the **size of the circular visible area**, expressed as a percentage of the frame.

- **0.0%**
The circle has zero radius.
 - With **Transparent disabled**: the entire frame is filled with the selected color.
 - With **Transparent enabled**: the clip is fully transparent.
- **Low values (10% - 40%)**
A small visible circular window appears at the center of the frame.

- **Medium values (40% - 80%)**
The visible area expands, cropping less of the image.
- **100.0%**
The circle fully covers the frame.
- The original clip is fully visible.
- Color and transparency settings have no effect.

Note:

Radius controls how much of the clip remains visible, not how much is removed.

Color

Defines the **fill color used outside the circle**.

- Can be selected from:
 - The color palette
 - The color picker in the preview

When the circle does not cover the full frame, this color replaces the cropped area.

Transparent (button)

When enabled:

- The area **outside the circle becomes transparent**
- **Underlying video tracks show through**

When disabled:

- The outside area is filled with the selected Color

Important behavior

- Transparency affects **timeline compositing**, not the clip itself.
- Export transparency is preserved only if the chosen export format supports an alpha channel.

Keyframes

All controls are keyframeable.

This enables:

- Animated circular reveals or hide effects
- Spotlight or iris-style transitions
- Dynamic masking over time

Visual characteristics

- Hard-edged circular boundary
- No feathering or soft edge
- Centered circle only (no position control)
- Clean, predictable masking behavior

Recommended use cases

- Circular reveals or transitions
- Picture-in-picture with circular framing
- Spotlight or focus effects
- Masking content for overlays or UI elements

Limitations

- Circle is always centered; cannot be repositioned
- No edge softness or feather control
- No aspect-ratio control (always circular)
- Transparency on export depends on format support

Crop: Rectangle is a filter used to crop a specified region of a video clip or image file to the current **Video Mode**. **Keyframes** are available with this filter.


Filter can be applied to Source, selected clip, Video Track or Output.

Parameters

- **Preset:** Save current settings for later use. To use later, just click the drop-down menu for selection.
- **Position & Size:** Values can either be input by keyboard or adjusted by the adjustable border corners and center.
- **Corner radius:** Apply a corner radius to the cropped selection. Can be adjusted by the slider bar, entering a value, or just click & scroll on the percent value.
- **Padding color:** Default is set to Black with no opacity (Alpha channel = 255). Any padding color can be selected and alteration of the Alpha channel can all be adjusted. You can also pick a color using the eye dropper tool.
Tip: To adjust opacity, lower Alpha channel from 255 to 0.


Example #1:

Video Mode 1920x1080, Image 1920x1080, Padding Alpha channel set to 255.




Example #2:

Video Mode 1920x1080, Image 1080x1350, Padding Alpha channel set to 0.




Apply to Source




The **Apply to Source** button in the **Crop: Rectangle** filter is the primary way to simply "crop a video".

Exporting without using **Apply to Source**



The clip is exported at its original size, with black bars around the cropped area.

Exporting using **Apply to Source**



Only the cropped area of the clip is exported. The clip will have the same dimensions as the ones set in the Crop filter.

Notes


- The **Apply to Source** button was added in Shotcut version 23.11.29.
- The **Apply to Source** button is only enabled when the aspect ratio of the source media matches your project's Video Mode because it is impossible to include the black padding that Shotcut adds to a source. To make sure the project and the media dimensions are the same, set the **Video Mode** to **Automatic** before importing your media file.
- Since this feature is intended to crop a single clip, you do not need to use the timeline.

How To Export a Cropped Area

1. On a new Shotcut project, go to **Settings > Video Mode** and choose **Automatic**.
2. Go to **File > Open File...** and import your media file.
3. Add the **Crop: Rectangle** video filter.
4. Adjust the size of the cropped area.
5. Click the **Apply to Source** button in the **Filters** panel.
6. Choose **Yes** in the dialog asking to change the Video Mode.
At this point, your clip will be cropped, and the **Crop: Rectangle** filter will be replaced by the **Crop: Source** filter.
Also, the **Export** settings will automatically be adjusted to match the new dimensions of the project.
7. Export


Residual Black Bars

It happens sometimes that the cropping is not perfect and leaves small black bars on one or two sides of the clip.




If that happens, edit the parameters of the **Crop: Source** filter to remove those bars.

Depending on which side(s) the bar(s) are located, adding 1 or 2 pixels to one or more of the parameters highlighted below should be enough to remove them.





Crop: Source removes rows or columns of pixels from the edges of the source video or image before the image is scaled and padded to match the current **Video Mode**. When cropping, selection is centered to the Video Mode because padding is used to maintain the aspect ratio and prevent the image from appearing stretched or squished. Keyframes are not available with this filter.



This filter can only be applied to **Source** or a selected clip; it cannot be applied to a track or **Output** because these are not sources.

Preset

Save the current parameter values for later use. To use later, just click the drop-down menu for selection. Preset folder: `crop`

Parameters

- **Center** automatically computes the crop values to remove artificial black bars that have been added to the source to make it match your Video Mode. For example, if you put an image with a 3:2 aspect ratio into a 16:9 project, it will be padded with black on the left and right so the whole image is displayed without the distortion of stretching. Turning on **Center** removes the black bars by removing some of the image from the top and bottom.
- **Center bias** only works when **Center** is checked. Instead of removing the same number of pixels from opposite edges, this is a convenient way to adjust the center or balance of the cropping to give control over the framing or region of interest/focus.
- **Top** removes rows of pixels from the top edge.
- **Bottom** removes rows of pixels from the bottom edge.
- **Left** removed columns of pixels from the left edge.
- **Right** removes columns of pixels from the right edge.

Notes

The values are in pixel units in terms of the current Video Mode, not the source resolution that may change depending upon whether a proxy is being used. Therefore, to convert the values from source to video resolution use the equations:

- $\text{Top/Bottom} = \text{source top} / \text{source height} * \text{mode height}$

- Top/Bottom = source-top / source-height * mode-height
- Left/Right = source-left / source-width * mode-width


Now, since these values must be rounded to enter them without a decimal, you may need to add 1 to one or more edges to get rid of black edges. You may find the Video Zoom helpful to look for black edges.



Applies a **softening** that spreads light within textures rather than simply blurring edges. The effect reduces micro-contrast and detail while largely preserving overall shapes, producing a softer, more organic image. Diffusion is a **spatial effect** and operates per frame. It does not analyze motion or time.

NOTE:

This filter is only available when GPU processing is enabled. It does not appear in CPU-only processing modes.



Parameters

Radius (0.00 - 20.00)

Controls the spatial extent over which diffusion spreads.

- **Low values ($\approx 0.00 - 3.00$)**
Very subtle diffusion. Fine texture is slightly softened with minimal spread.
- **Mid values ($\approx 3.00 - 10.00$)**
Noticeable diffusion. Light and detail begin to bleed into neighboring pixels.
- **High values ($\approx 10.00 - 20.00$)**
Strong diffusion. Texture and fine detail are heavily softened.

Behavior note:

- Radius defines **how far** the diffusion reaches, not how strong it appears.
- Increasing Radius increases the area of influence, not blur intensity.

Blurriness (0 - 100%)

Controls how strongly the diffused result blends with the original image.

- **Low values**
Diffusion is faint; original texture remains mostly intact.
- **Higher values**
Diffusion becomes more pronounced, reducing fine detail and local contrast.

Important distinction:

Blurriness does not increase the diffusion radius. It controls the **visibility and strength** of the diffusion effect.

Parameter interaction

- **Radius** determines the scale of diffusion.
- **Blurriness** determines how dominant the diffused image becomes.
- Small Radius with higher Blurriness softens texture without obvious blur.
- Large Radius with high Blurriness produces a pronounced soft-focus look.

Visual characteristics

- Reduced micro-contrast
- Softer internal texture rather than sharp edge blur
- Slight “glow-through” within detailed areas
- Overall shapes remain readable longer than with standard blur filters

Recommended use cases

- Softening harsh digital texture
- Subtle beauty or diffusion pass
- Reducing noise-like detail without heavy blur
- Creating a mild optical softness reminiscent of diffusion filters used in photography
- Stylized or atmospheric visuals


Limitations

- GPU processing required
- Not keyframeable
- Can reduce perceived sharpness if overused
- Not a physical simulation of optical diffusion filters



Applies a **time-varying geometric distortion** to the image, warping the frame along a moving wave pattern.

The filter displaces pixels spatially over time, creating wobble, ripple, or oscillation effects rather than changing color or tone.



This filter is intended for **stylization and motion distortion**, not for lens correction or stabilization.

Parameters

Amount (0.0 - 100.0%)

Controls the **strength of the distortion**.

- **0.0%**
No distortion. The image is unchanged.
- **Low values (5% - 20%)**
Subtle wobble or gentle image drift.
- **Medium values (20% - 60%)**
Clearly visible distortion. The image bends and shifts over time.
- **High values (60% - 100%)**
Strong warping. The image becomes unstable and heavily deformed.

Important behavior

- **Amount is a master control.**
- If Amount is set to **0%**, **Frequency and Velocity have no effect**, regardless of their values.

Note:

Amount defines whether distortion exists at all.

Frequency (0.0 - 100.0%)

Controls the **spatial repetition of the distortion pattern**.

- **Low values (0% - 5%)**
Large, smooth bends affecting broad areas of the frame.
- **Medium values (5% - 30%)**
Multiple waves appear across the image.
- **High values (30% - 100%)**
Dense, zig-zag or ripple-like patterns.

Important behavior

- Frequency affects **how many waves** exist, not how fast they move.

Note:

Frequency controls *how often* the image bends across space, not time.

Velocity (0.0 - 100.0%)

Controls the **speed and direction of the distortion's motion over time**.

- **0.0%**
Static distortion pattern.
- **Low values (5% - 30%)**
Slow drifting motion.
- **Medium values (30% - 70%)**
Continuous wobbling.
- **High values (70% - 100%)**
Rapid oscillation and strong temporal instability.

Important behavior

- The default diagonal motion (top-left to bottom-right) is intrinsic to the effect.
- Velocity does not change distortion strength, only its movement.

Note:

Velocity controls how fast the distortion travels, not how strong it is.

Keyframes

All parameters can be keyframed.

This enables:

- Gradual buildup or release of distortion
- Pulsing or rhythmic warping
- Synchronization with music or motion

Parameter interaction

- **Amount** enables and scales the effect.
- **Frequency** defines the number of spatial waves.
- **Velocity** animates those waves over time.

If Amount is zero, the filter is effectively disabled, regardless of the other parameters.

Visual characteristics

- Time-based geometric warping
- Diagonal wave motion by default
- Zig-zag or ripple patterns at higher Frequency
- No color, brightness, or alpha changes

This is a **spatial + temporal** effect and must be evaluated during playback.

Recommended use cases

- Stylized distortion or glitch-like effects
- **Simulating heat, water, or instability**
- Transitional effects between scenes
- Music videos or experimental visuals

Limitations


- Not suitable for corrective distortion
- Direction of motion is fixed
- Can cause motion sickness at medium-high values
- Compression may exaggerate artifacts



Applies ordered dithering to the image to approximate smooth tonal transitions using a limited number of discrete values.

Instead of creating smooth gradients, the filter distributes pixels according to a fixed pattern so that tone changes are perceived visually rather than represented directly.

Dithering predates digital video. It was used in early printing, photography, and computer graphics to simulate continuous tones on media that could only represent a small number of intensity levels. Conceptually, dithering is closely related to [halftoning](#): both replace smooth shading with structured patterns that rely on human visual averaging.



Dither is a perceptual filter, not a blur; it suggests missing tones through spatial patterns rather than creating them.

Parameters

Levels (0.0 - 100.0%)

Controls the **degree of tonal quantization** before dithering is applied.

- **0%**
Very few tonal levels. **Strong pattern visibility** and high contrast between light and dark regions.

- **Low values (10% - 30%)**
Heavy dithering with clearly visible structure. Gradients are replaced by pronounced dot or line patterns.
- **Medium values (30% - 70%)**
Balanced dithering. Patterns remain visible but tones read more smoothly.
- **100%**
Maximum available levels for this filter.
Dithering remains visible; the effect is never fully disabled.

Important behavior

- Like [Posterize](#), 100% does **not** restore continuous tones.
- The filter always replaces smooth gradients with patterned approximations.

Matrix (pattern selection)

Selects the **spatial pattern** used to distribute pixels.
The matrix determines how tonal error is arranged visually.

2×2 Magic Square

- Very small repeating pattern
- Highly visible structure
- Produces coarse, noisy-looking dithering
- Suitable for extreme retro or low-resolution looks

4×4 Magic Square

- Larger version of the 2×2 pattern
- More tonal steps and smoother gradients
- Still clearly patterned

4×4 Lines

- Uses linear, stripe-like arrangements
- Creates directional texture
- Can emphasize horizontal or vertical structure

6×6 90 Degree Halftone

- Dot-based pattern aligned at right angles
- Closely resembles traditional print halftone screens
- Produces a graphic, printed appearance

6×6 Ordered

- Regular ordered grid
- Balanced between smoothness and structure
- Less visually aggressive than smaller matrices

8×8 Ordered

- Large ordered grid
- Smoother perceived gradients
- Pattern is less obvious at normal viewing distances

Order-3 Clustered

- Groups pixels into small clusters
- Softer, less regular appearance
- More organic than ordered grids

Order-4 Ordered

- Medium-density ordered matrix
- Good compromise between pattern visibility and smoothness

Order-8 Ordered

- High-density ordered matrix
- Smoothest tonal transitions among the options
- Least visually distracting pattern

Keyframes

Levels can be keyframed.

This allows:

- Gradual transitions between smooth gradients and strong dithering
- Animated stylistic changes over time
- Controlled visual degradation effects

Parameter interaction

- **Levels** controls *how many tones exist*.
- **Matrix** controls *how missing tones are distributed spatially*.

Lower Levels with small matrices produce aggressive, high-contrast patterns.
Higher Levels with large matrices produce subtler, smoother dithering.

Visual characteristics

- Structured, repeating patterns
- No true blur or noise
- Detail is preserved as pattern density rather than smooth shading
- Strong dependence on viewing distance

This is a **spatial** effect; it does not change over time unless keyframed.

Recommended use cases

- Reducing banding in gradients without using blur
- Simulating print, newspaper, or early computer graphics
- Retro or technical visualization aesthetics
- Preparing footage for limited-color or stylized output

Common combinations


- **Posterize + Dither**
Produces controlled, intentional banding with structured tone approximation.
- **Dither + Old Film: Grain (subtle)**
Breaks up overly regular patterns and adds organic texture.
- **Dither + Gaussian Blur (very low radius)**
Softens pattern harshness while preserving dithering logic.

Limitations

- Patterns are fixed and may alias during scaling or compression
- No adaptive or error-diffusion dithering (e.g. Floyd - Steinberg)
- Can introduce visible structure in flat areas
- Not suitable for naturalistic footage without careful tuning



The Drop Shadow filter adds a background to an element optionally with soft edges and offset a little from the foreground. The default values creates the appearance of a shadow behind the subject:



This filter uses the input alpha channel. This means that it will not work as expected when used after a text filter on a video clip because that alpha channel is completely opaque. Rather, you must use a text clip on an upper track. However, this effect is not only for text. You can be overlaying a PNG or JSON Animation that has a transparent background, for example.

This filter was added in version 24.06.


Parameters

- **Color** is the color of the shadow or glow. The default is a 70% translucent black. Click the color block to open a color dialog to change it, or click the eye-dropper icon to pick a color from the screen.
- **Blur** is the softness of the edges of shadow. The default is a small amount.
- **X** and **Y** are the horizontal and vertical offset from the foreground.

Tip: you can also create a sort of glow effect by setting **X** and **Y** to 0 and increasing **Blur**.



Applies a **non-uniform, position-dependent scaling** to the image. Instead of scaling the entire frame evenly, this filter stretches or compresses different regions by varying amounts, producing bending or “elastic” deformation effects.



Elastic Scale is a **geometric distortion filter**, not a crop or zoom tool.

What “elastic” means here

In a normal scale:

- Every pixel moves proportionally
- Shapes remain straight

In **elastic scaling**:

- Some areas scale more than others
- Straight lines can bend
- The image appears stretched, squeezed, or curved around a center

Note:

Elastic Scale redistributes space across the image instead of resizing it uniformly.

Parameters

Center (0.0 - 100.0)

Defines the **horizontal center point** of the elastic effect.

- **50.0**
Effect is centered in the frame.
- **Lower values**
Effect is biased toward the left.
- **Higher values**
Effect is biased toward the right.

This control determines **where the deformation pivots**, not how strong it is.

Linear width (0.0 - 100.0)

Defines the **width of the linear (uniformly scaled) region** around the center.

- **0.0**
No linear region exists.
- **Higher values**
A wider central zone is scaled evenly, with elastic deformation occurring outside it.

Important behavior

- When **Linear width = 0.0**, the **Linear scale factor has no effect**.

Note:

Linear width defines how much of the image behaves “*normally*.”

Linear scale factor (0.0 - 100.0)

Controls the **uniform scaling applied inside the linear region**.

- **50.0**
Neutral scaling for the linear region.
- **Lower or higher values**
Compress or expand the central region evenly.

Important behavior

- Has **no effect** unless Linear width is greater than 0.0.

Non-linear scale factor (0.0 - 100.0)

Controls the **strength and direction of elastic deformation** outside the linear region.

- **50.0**
Neutral state; no elastic distortion.

- **Lower values (<50)**
Image bends outward, producing a bulging or fisheye-like effect on one side.
- **Higher values (>50)**
Image bends inward, compressing space and creating a pinched or curved look.

This control affects **how space is redistributed**, not just how large areas appear.

Note:

Non-linear scale factor controls *curvature*, not size.

Keyframes

All parameters are keyframeable.

This enables:

- Animated warping effects
- Progressive bending or straightening
- Motion-driven distortion synchronized with action or music

Parameter interaction (important)

- **Center** defines *where* the deformation is anchored.
- **Linear width** defines *how much* of the image is unaffected by elastic bending.
- **Linear scale factor** affects only the linear region.
- **Non-linear scale factor** controls the elastic distortion outside that region.

If:

- Linear width = 0.0
 - Linear scale factor is ignored
 - Non-linear scale factor alone drives the effect

Visual characteristics

- Horizontal bending or squeezing
- Curved vertical lines
- Asymmetric distortion if Center is offset
- No change to color or luminance

The effect is purely **geometric**.

Recommended use cases


- Stylized distortion or surreal effects
- Simulating lens-like warping without true optics
- Transitional effects between shots

- Emphasizing or de-emphasizing regions without cropping

Limitations

- Distortion is horizontal only
- Can easily produce unnatural geometry
- Not suitable for corrective scaling
- No edge feathering or protection zones


Gradually reveals a clip at its beginning by transitioning from black or transparency to the full image over a specified duration.



This filter is commonly used to create smooth visual entrances and to avoid abrupt cuts at the start of a clip.


Access methods

- From the **Filters** list as **Fade In Video**
- Directly from the timeline by hovering near the top-left corner of a clip's start and dragging the rounded handle inward the clip (to the right).
 - This action automatically adds the filter to the clip with the chosen duration.
 - Double-click the fade control's circle handle to either add a one second fade or remove the fade.



General behavior

- The fade is applied from the clip's start only.
- If the clip contains audio, audio is faded in as well, and a corresponding **Fade In Audio** filter is automatically added.
- The fade affects the clip locally and does not alter adjacent clips.



Parameters

Duration (time value)

Sets how long the fade-in lasts.

- Format: **HH:MM:SS.SSS**
 - Example: 00:00:04.200 = 4.2 seconds
 - Depends upon **Settings > Time Format**. If using a timecode format, it is HH:MM:SS:FF where FF = number of frames since the last second and the maximum before returning to 0 depends on the project frame rate.
- Defines the interval from the clip's first frame to full visibility.
- Short durations create quick, subtle fades.
- Longer durations produce slow, deliberate entrances.
- Changing Duration updates the fade curve proportionally.

The adjacent **Set as default** button stores the current duration as the default for newly added fades.

Adjust opacity instead of fade with black (checkbox)

Changes **how** the fade is performed.

- **Disabled (default)**
The clip fades **from black** to full image.
- **Enabled**
The clip fades by **adjusting opacity**:
 - The clip becomes gradually opaque
 - Underlying tracks (if any) show through during the fade

Important behavior note (how the filter is added matters)

The default behavior of **Fade In Video** depends on how the filter is created:

Added from the **Filters** list

- The filter defaults to fade from black
- **Adjust opacity instead of fade with black** is disabled

Added directly from the timeline clip handle

- The filter defaults to opacity fade
- **Adjust opacity instead of fade with black** is enabled if the clip is NOT on the bottom video track (V1). Otherwise, on V1, it fades with black.

In both cases, the same filter is used; only the initial checkbox state differs.

Audio behavior

When the clip includes audio:

- A matching **Fade In Audio** filter is added automatically
- Audio fades in over the same **Duration**
- Audio and video remain synchronized

This happens whether the filter is added manually or via the timeline handle.

Visual characteristics

- Smooth, linear transition by default
- No geometry or color changes
- Predictable timing aligned to the clip start
- Behavior is identical whether added via filters or timeline interaction


Recommended use cases

- Softening the start of a clip
- Introducing scenes without hard cuts
- Matching visual and audio entrances
- Compositing overlays using opacity fades

Limitations

- Applies only at the **start** of a clip
- Duration is time-based, not frame-count based
- No curve or easing control
- Not keyframeable

Gradually removes a clip at its end by transitioning from the full image to black or to transparency over a specified duration.




This filter is used to create smooth visual exits and to avoid abrupt cuts at the end of a clip.

Access methods

- From the **Filters** list as *Fade Out Video*
- Directly from the **timeline** by hovering near the **top-right corner of a clip's end** and dragging inward
 - This action automatically adds the filter to the clip with the chosen duration
 - Double-click the fade control's circle handle to either add a one second fade or remove the fade.

General behavior (important)

- The fade is applied **to the end of the clip only**.
- If the clip contains audio, **audio is faded out as well**, and a corresponding **Fade Out Audio** filter is automatically added.
- The fade affects only the selected clip and does not modify following clips.



Parameters

Duration (time value)

Sets how long the fade-out lasts.

- Format: **HH:MM:SS.SSS**
- Example: 00:00:03.000 = 3 seconds
- Defines the interval from full visibility to complete fade at the clip's end.

Behavior notes

- Short durations create quick, subtle exits.
- Longer durations produce slow, deliberate fades.
- The adjacent **Save as default** button stores the current duration as the default for newly added fade-outs.

Adjust opacity instead of fade with black (checkbox)

Controls **how** the fade-out is performed.

- **Disabled (default when added from Filters list)**
The clip fades **to black**.
- **Enabled (default when added from the timeline handle)**
The clip fades by **reducing opacity**:
 - The clip becomes gradually transparent
 - Underlying tracks show through during the fade

Note:


Use opacity fades when compositing over other video; use black fades for standalone endings.

How the filter is added (important note)

The default fade mode depends on **how the filter is created**:

- **Added from the Filters list**
 - Defaults to **fade to black**
 - “**Adjust opacity instead of fade with black**” checkbox is **disabled**
- **Added from the timeline clip handle**
 - Defaults to **opacity fade**
 - “**Adjust opacity instead of fade with black**” checkbox is **enabled** if the clip is NOT on the bottom video track (V1). Otherwise, on V1, it fades with black.

The mode can always be changed manually after creation.



Audio behavior

When the clip includes audio:

- A **Fade Out Audio** filter is added automatically
- Audio fades out over the same Duration
- Audio and video remain synchronized

Visual characteristics

- Smooth transition from image to black or transparency
- No geometry or color changes
- Timing is locked to the clip's end
- Identical behavior whether added via the filter list or timeline handle

Recommended use cases

- Ending scenes cleanly
- Avoiding abrupt visual and audio cutoffs
- Transitioning between sequences
- Compositing overlays using opacity fades

Limitations

- Applies only at the **end** of a clip
- Duration is time-based, not frame-based
- No easing or curve control
- Not keyframeable



Overview

Convert fisheye video to rectilinear, and vice versa. It is based on the angular [mapping functions](#) actually used in fisheye lens design, to get the best possible results. It can also be used to correct the slight distortion of some wide-angle converters, or to bend the image beyond recognition for special effects and light shows.

Presets

Camera	Resolution	Type
4K CLONE	1080p	Wide (Action)
4K CLONE	1080p	Wide (Focus)
4K CLONE	1080p	Wide (Linear)
4K CLONE	720p	Medium (Focus)
4K CLONE	720p	Medium (Linear)
HERO3	1080	Wide (Focus)
HERO4	1080	Medium (Action)
HERO4	1080	Medium (Focus)
HERO4	1080	Medium (Linear)
HERO4	1080	SuperView (Action)
HERO4	1080	SuperView (Focus)
HERO4	1080	SuperView (Linear)
HERO4	1080	Wide (Action)
HERO4	1080	Wide (Focus)
HERO4	1080	Wide (Linear)
HERO4	1440 4:3	Wide (Action)
HERO4	1440 4:3	Wide (Focus)
HERO4	1440 4:3	Wide (Linear)
HERO4	1440	Wide (Action)
HERO4	1440	Wide (Focus)
HERO4	1440	Wide (Linear)
HERO5	1080	SuperView (Linear)
HERO5	1080	Wide (Linear)

Save current settings as a preset

Delete selected preset

Stored presets location: ~Shocut\presets\frei0r.defish0r

Settings → [App Data Directory](#) → Show

Selection: Add or remove fisheye effect.

Focal ratio

The amount of lens distortion

Default setting is 0.500

Range: 0.000 - 1.000

Quality (Resample)

Option
Nearest neighbor
Bilinear
Bicubic smooth
Bicubic sharp
Spline 4x4
Spline 6x6
Lanczos 16x16

Lens

Select a lens distortion that best matches your camera.

Option	Fisheye mapping function
Equidistant	Linear-scaled
Ortographic	Orthogonal
Equiarea	Equal-area
Stereogrphaic	Panoramic

- Non-Linear Scale: The image will be stretched/squished to fix camera scaling between 4:3 and 16:9. Like used in GoPro's superview.
 - Use negative values for up-scaled videos
 - Use positive values for down-scaled videos

Scale

Preset Scale methods.

Lock pixels at specific locations.

Presets:

- Scale to Fill
- Keep Center Scale
- Scale to Fit
- Manual Scale

Y - Separate Y Scale -0.49 to 05.0

Crop - Remove distorted edges

Aspect

Pixel aspect ratio

Presets:

- Square Pixel (Default)
- PAL DV 1.067
- NTSC DV 0.8889
- HDV 1.333
- Manual Aspect


Tips

- this filter is very CPU intensive on high resolution videos so using [Settings > Preview Scaling](#) will have a massive impact allowing realtime playback even for 4K sources.
- See also [Lens Correction Video Filter](#) which supports keyframes unlike this filter.



Flips the video / image **vertically**, turning it upside down.

This filter performs a simple geometric transformation without altering color, timing, or aspect ratio.



Flip is a **binary operation**: it is either applied or not; there are no adjustable parameters.

General behavior

- The entire frame is inverted top-to-bottom
- Left-right orientation is unchanged
- Audio is unaffected
- The transformation is applied uniformly to the whole clip

This is a **pure geometric transform**, not a distortion or rotation.

Controls

This filter has **no parameters**.

Applying or removing the filter toggles the vertical inversion.

Visual characteristics

- Top becomes bottom, bottom becomes top
- Straight lines remain straight
- No scaling, cropping, or resampling artifacts beyond the flip itself

Recommended use cases

- **Correcting inverted footage**
Some cameras or capture devices record video upside down depending on mounting orientation.

- **Ceiling- or rig-mounted cameras**
Common in surveillance, overhead rigs, microscopes, or experimental setups.
- **Projection and mirror workflows**
Correcting orientation when footage passes through mirrors, prisms, or optical systems.
- **Creative or stylistic effects**
Deliberate inversion for disorientation, dream sequences, or abstract visuals.
- **Preprocessing for further transforms**
Used before rotation, masking, or compositing to simplify later steps.

Comparison with related filters

Flip vs **Rotate**

- **Flip** inverts the image vertically in one operation.
- **Rotate** changes orientation by arbitrary angles and may introduce resampling or black borders.


Flip is lossless in terms of geometry; rotation is not.

Limitations

- Vertical only (no horizontal flip)
- No partial or masked application
- No animation or keyframing
- Always applies to the entire frame



The **Glitch** video filter introduces randomized digital distortion to simulate, for example, corrupted video signals or transmission errors. The effect works by displacing horizontal bands of the image and altering color channels during short glitch events. It is useful for stylized transitions, motion graphics, or adding visual noise to footage.



Parameters

Frequency

Controls how often glitch events occur over time.

- Low values produce occasional, subtle glitches.
- High values increase the number of glitch events, creating a more chaotic and continuous effect.

This parameter affects **when** glitches happen, not how strong they are.

Block Height

Determines the vertical size of the horizontal strips that are distorted during a glitch.

- Small values create thin, fine-grained tearing.
- Large values produce thick, blocky horizontal bands.

This parameter defines the “resolution” of the glitch distortion.

Shift Intensity

Adjusts the horizontal displacement applied to the selected pixel bands and color channels.

- Higher values increase the sideways offset of the glitch bands.
- Also increases RGB channel separation, producing stronger color fringing.

- Does not affect glitch frequency.

This parameter controls **how far** the image shifts during a glitch.

Color Intensity

Controls the amount of color distortion applied during glitch events.


- Low values keep colors mostly intact.
- High values introduce strong channel misalignment, tint shifts, and chromatic artifacts.

This parameter influences the visual “noise” and color corruption of each glitch.



Adds a soft halo around bright areas by blurring and reintroducing luminance into surrounding pixels.

The effect increases perceived brightness and softness without altering geometry. It is commonly used to suggest light bloom, diffusion, or optical softness.



The Glow filter operates as a **spatial effect**. It does not analyze motion or time; any temporal change comes only from keyframing.

Parameters

Blur (0 - 100%)

Controls the radius and strength of the glow by increasing the amount of blur applied to bright regions.

- **Low values (0 - 20%)**
Subtle glow. Bright edges soften slightly with minimal spread. Often perceived as gentle diffusion rather than an obvious effect.
- **Mid values (20 - 60%)**
Noticeable halo around highlights. Light areas bleed into adjacent pixels, increasing perceived brightness and softness.
- **High values (60 - 100%)**
Strong bloom. Highlights expand significantly and can wash out nearby detail. Fine edges may lose definition.

Important behavior notes:

- Blur does **not** brighten dark areas on its own. The effect is driven primarily by existing brightness.
- Increasing Blur enlarges the glow area; it does not act like a sharpening or contrast control.
- The scale is non-linear in perception: higher values increase spread faster than they increase apparent intensity.

Keyframes

The **Blur** parameter can be keyframed.

This enables:

- Gradual introduction or removal of glow
- Pulsing or breathing light effects
- Simulating changes in exposure, light intensity, or dream-like transitions over time

Parameter interaction

This filter has a single control.

Perceived strength depends heavily on the clip's luminance range and contrast:

- High-contrast or high-exposure clips show stronger glow at the same Blur value.
- Flat or dark footage may require higher values to produce visible results.

Visual characteristics

- Soft halos around bright objects or highlights
- Reduced edge sharpness near light sources
- Increased perceived brightness without true exposure change
- At high values, highlights may appear smeared or over-bloomed

Recommended use cases


- Softening harsh digital highlights
- Creating light bloom or diffusion effects
- Stylized looks (dream sequences, memories, fantasy)
- Enhancing light sources such as lamps, reflections, or specular highlights
- Subtle beauty or diffusion pass when used at low values

Limitations

- No threshold or intensity control: all bright areas contribute equally.
- Can reduce fine detail and edge clarity at higher values.
- Not a physical simulation of camera or lens bloom; it is a visual approximation.
- Excessive use may cause washed-out highlights or loss of contrast.



The **GPS Graphic** filter allows you to overlay telemetry data—such as speed, altitude, coordinates, and maps—directly onto your video. This is an essential tool for action videography, including cycling, drone flight, motorsports, and hiking, where visualizing the “where” and “how fast” adds context to the footage. For example:



Requirements

To use this filter, you must have a GPS data source. Shotcut supports:

- **External Files:** `.gpx`, `.tcx`, `.fit`, `.nmea`, and `.kml`.
- **Embedded Telemetry:** Some cameras (like GoPros, certain DJI drones, and Insta360) embed GPS data directly into the video file metadata.
- Shotcut can extract the telemetry in GoPro camera videos to GPX using **Properties > menu-button > Export GPX**. (This requires the `gopro2gpx` executable if you are on a Linux or BSD system running a build we did not make.)

Getting Started: Choose a Preset

This filter comes with a long list of ready presets! It almost feels unusable without them, and they provide a great teaching aid. Please start with one of these to get started.

- 1D graph: Altitude
- 1D graph: Altitude, follow @90%
- 1D graph: HR
- 1D graph: HR, follow @90%
- 1D graph: Speed
- 1D graph: Speed, follow @90%
- 2D map: follow @90%
- 2D map: follow @95%
- 2D map: follow @98%
- 2D map: full map progress line
- Colors: default (5)
- Colors: green middle (5)
- Colors: inversed default (5)
- Map colors: gradient of grade
- Map colors: gradient of heart rate
- Map colors: gradient of speed
- Map style: 25% opacity thin future path
- Map style: hide future path
- Map style: show GPS points as dots @99%
- Simple line progressbar
- Speedometer: 0-30 km/h
- Speedometer: 0-60 mi/h**
- Speedometer: min-max km/h
- Text only: Altitude
- Text only: GPS coordinates
- Text only: Heart rate
- Text only: Percentage of trimmed track
- Text only: Speed

Parameters

Open file

GX019752.gpx

Preset

Speedometer: 0-60 mi/h

+

-

GPS options

GPS offset

+

1927

:

2

:

3

:

37

⏮

⏪

⏩

⏭

GPS smoothing

5 points

↺

Video speed

1.00 x

↺

↻

Graph data

Data source

Speed

Graph type

Speedometer

Trim time

0.0 %

↺

↻

100.0 %

↻

Crop horizontal

0.0

↺

↻

100.0

%

↻

Crop vertical

0.0

↺

↻

60.0

va

↻

Graph design

Color style

Two colors

Color

↺

↻

5 colors

↻

Now dot

✓

Color

↺

↻

Now text

✓

Rotation

0.0

°

↻

Thickness

5

px

↻

Draw legend

✓

Unit

mi/h

↻

Position

113

,

608

↻

Size

576

x

324

↻

Background options

Image path

!

↻

Scale

1.000

↻

The GPS Graphic filter is divided into several sections to manage the data source, visual style, and synchronization.

Parameter Description

File	Select the GPS data file. If your video has embedded data, click “Extract from Video.”
-------------	--

Parameter	Description
Graphic	Choose the type of visualization (e.g., Speedometer, Map, Altitude Chart, Compass).
Units	Toggle between Metric (km/h, meters) and Imperial (mph, feet).
Smoothing	Reduces “jitter” in the data caused by GPS signal noise.
Sync	Adjust the time offset if the GPS recording did not start at the exact same moment as the video.

Don't forget! Some of the parameters have a tool tip that you can hover the mouse over to get helpful information and tips.

Types of Graphics

Shotcut provides a wide variety of preset graphic types that can be customized for color, size, and position:

1. Speed and Distance

- **Speedometer (Dial or Digital):** Displays the current speed.
- **Odometer:** Shows total distance traveled.
- **VSI (Vertical Speed Indicator):** Useful for aviation or drone footage.

2. Maps and Navigation

- **Map:** Displays a top-down view of your path. You can choose different map providers (OpenStreetMap, Stamen) or use a “Path Only” view.
- **Compass:** Shows the heading/direction of travel.
- **Coordinates:** Displays the raw Latitude and Longitude.

3. Elevation and Force

- **Altitude:** Displays the current height above sea level.
- **Elevation Chart:** Shows a profile of the terrain covered.
- **G-Force:** Visualizes acceleration and cornering forces.

Workflow: Adding and Syncing

1. **Add Filter:** Select your clip on the Timeline and add the **GPS Graphic** filter.
2. **Load Data:** Click the folder icon to load your `.gpx` file or use the “Extract” button if using a GoPro.
3. **Choose a Preset:** Select a graphic type from the dropdown.
4. **Position the Graphic:** Use the **VUI (Visual User Interface)**—the rectangle in the preview player—to drag the graphic to a corner and resize it.
5. **Synchronize:** If the speed doesn't match the action in the video (e.g., the speed increases *after* the car starts moving), use the **Sync** slider or enter a time offset in seconds.

Customizing the Look

You can modify the appearance of the graphic to match your video's aesthetic:

- **Colors:** Change the background, text, and gauge colors.
- **Fonts:** Select any system font for the data readout.
- **Opacity:** Make the background of the gauges transparent to show more of the underlying video.

Important Notes

- **Privacy:** Remember that GPS data contains precise location information. Be mindful of this when sharing videos that show your home or private locations.
- **Hardware Acceleration:** Rendering complex GPS graphics (especially high-resolution maps) can be CPU-intensive.

See also: [GPS Text Video Filter](#)



Introduction

The **GPS Text** filter allows you to add useful GPS related data (such as speed, altitude or distance) on top of and in sync with your video. It is available since version 21.08.



Note: you need to record your GPS track with a device (phone, smartwatch, bike computer or dedicated GPS tracker) AND export your track to a .gpx or .tcx file. If you record using an action camera that has GPS (like GoPro), you need to extract the data to a separate file ([there](#) are [online](#) and [offline](#) tools for this).

Update: Shotcut 22.09 now supports .gpx extraction itself:

Properties panel → (hamburger menu) → Export GPX.

Tip: hover your mouse over filter elements for some extra explanations


Usage

1. Add the **GPS Text** filter on a clip in the timeline.

2. Click the Open file button in the filter panel to load the **.gpx** or **.tcx** file.




3. Use the **GPS offset** row to sync the GPS track time to the video time (check out some hints below; if you're *extremely lucky* they are already in sync and you can skip this step).




4. Edit the filter **text panel** with any keywords from the available list. You can also format the text to your liking just like **Text: Simple** filter.

Tip: hover your mouse over the list for extra keyword arguments.



5. Check out the results live in the video player.



6. If needed, change the more advanced options in the filter panel. There is also basic keyframing support for in/out and text position/size.

7. Export

Syncing workflow details and tips

TLDR: watch out for timezones; buttons have tooltips; remember to film a clock from time to time for easy sync later (but do remember its timezone); Also: scrollwheel to add/subtract a few seconds for quick corrections.


At the very bottom of the filter panel (you'll probably need to scroll to see it) there are 2 important date-time notes: the Video's and the GPS's start times (Note: they will be first updated after you load a GPS file). The **GPS time** is always in the UTC timezone, but the **Video start time** is read from the file's **Creation Time** metadata field, which is currently very not-standardized across devices, and so, it might be in your local timezone or not, it might be set to the start or end of the video or it might just be a very random time if your device loses power and doesn't sync it at startup. The important thing to remember here is that no matter **when** you shot your video, *this is what Shotcut sees as your file's start time* (converted to UTC, so the same as GPS time) therefore this is what will be used for offseting the GPS track. You can use the `#file_datetime_now#` keyword to see it live in the output.

Sync buttons

Syncing the GPS track to the video file will be the most annoying part of using this filter. To the right of the sync text areas there are a number of 4 buttons that will probably cover 95% of all use cases thus avoiding any need for doing time math. This will be frustrating at first but you will notice that there is a pattern to follow and it will be the same everytime for each device.

Note: if your device doesn't support auto time sync over the internet you will most likely still need to match a few seconds (maybe even minutes) to the real time in addition to one of the helper buttons. I recommend you take a photo of a (digital) clock and check the photo's properties for the time.

I'll go over the buttons from left to right, first being the clock one, and last the "Pause" one.



Sync button 1: Timezone removal (GoPros)

Probably the most used button as I expect most shots that would work with a GPS overlay would come from an action camera and GoPro is doing this as well.

Why? If you can't set the timezone in the camera clock settings, it will probably be wrongly set to UTC (well it's wrong unless you live in UTC area) so the file will match the GPS track a few hours ahead or behind.

Devices? GoPros, some DSLRs/Mirrorless, most non professional grade video cameras will have this issue. This is an easy one to find as it will remain constant for all files.

Sync button 2: Video duration removal (Android)

Why? Some devices (*cough* android phones *cough*) decided it's a good idea to set the video's creation time after you finish recording. So if you record a 15 minute video starting at 12:00, it will store the time as 12:15.

Devices? Android smartphones are the only ones that do this as far as I know. It's a big pain as you can't copy paste the filter to multiple clips, each clip has to be individually updated as it will most likely have a different duration.

Sync button 3: Sync beginnings (tests & dashcam)

If you just want to see the filter work just press this button and it will immediately sync the 2 times no matter how far apart they are.

When? This is good for testing random files or if you don't want to bother with times and just remember to start recording video and GPS at the same time.

Devices? My dashcam resets the time every time it loses power for more than a

few hours. I gave up trying to set it correctly so I just start the GPS track and immediately hit start record.


Sync button 4: Sync to now

When? If you want a very precise sync that doesn't need time keeping at all, this is the way: just start filming before the GPS device and make sure to record the moment the GPS device gets the first fix and starts recording. Move the playhead to that exact moment in the video, press the button and you're done.

You can also make a sign or cover your camera with your palm as a mark point. If you clap to sync audio to video this workflow will be very familiar.

Sync extra: record a precise clock

Your phone automatically syncs pretty precisely all the time. You can just record it's clock (make sure you can see the seconds) at any moment and manually match the `#gps_datetime_now#` time to what you see in the video. This is my preferred method as it's simply the easiest and you don't need to preserve file metadata (don't forget about your clock's timezone though).




More info

[Here](#) you can read the entire list of arguments and extra keywords. Most numeric fields accept obvious arguments (for example the speed or distance can be converted to imperial by typing “mi”, “mile”, “mi/h” between the `# ... #` (example: `#gps_speed mi#`), time accepts standard `%H:%M`, temperature “F” and “K”, etc).

The Gradient Map Video Filter was added in Shotcut 25.01.03


The Gradient Map Video Filter maps the colors of an image to a gradient according to their intensity. The controls for the filter include a configuration for the gradient. The colors on the left of the gradient are applied to dark colors. The colors on the right of the gradient are applied to bright colors.

The filter includes presets for Thermal and Sepia gradients.



Divides the clip into a **repeated tiled layout**, displaying multiple scaled copies of the same video in a regular grid.

The filter duplicates the source image spatially; it does not crop different regions of the frame. Every cell shows the **entire clip**, scaled to fit.



Grid is a **spatial replication effect**, not a layout or masking tool.

Parameters

Rows (0.0 - 100.0%)

Controls the number of vertical divisions in the grid.

- **0.0%**
No vertical subdivision. The clip is shown once, filling the frame.
- **Low values (≈5 - 15%)**
A small number of rows. On a 1080p project, the default value (10%) typically results in **3 rows**.
- **Higher values**
More rows are added. Each tile becomes shorter vertically as the grid density increases.

Important behavior notes:

- The percentage value does **not** correspond directly to a row count.
- The mapping is non-linear and resolution-dependent; exact row numbers cannot be inferred from the slider value.
- Increasing Rows reduces the vertical resolution of each tile.

Columns (0.0 - 100.0%)

Controls the number of horizontal divisions in the grid.

- **0.0%**
No horizontal subdivision. The clip is shown once, filling the frame.
- **Low values ($\approx 5 - 15\%$)**
A small number of columns. At the default value (10%), this typically results in **3 columns** on a 1080p project.
- **Higher values**
More columns are added. Each tile becomes narrower as the grid density increases.

Important behavior notes:

- Like Rows, the percentage does **not** map linearly to a column count.
- The effective number of columns depends on project resolution.
- Increasing Columns reduces the horizontal resolution of each tile.

Keyframes

Both **Rows** and **Columns** can be keyframed.

This enables:

- Animating from a single full-frame view into a tiled grid
- Dynamic changes in grid density over time
- Zoom-like transitions that break the frame into multiple repeating views

Parameter interaction

Rows and Columns are independent but cumulative:

- Increasing both increases the total number of tiles multiplicatively (rows $\times$ columns).
- Increasing only one creates a strip layout (vertical or horizontal repetition).
- High values on both sliders can rapidly reduce per-tile clarity due to downscaling.

Visual characteristics

- Repeated copies of the same video filling the frame
- Uniform grid with equal-sized tiles
- All tiles remain time-synchronized
- No borders or spacing between tiles
- Image detail decreases as grid density increases

Recommended use cases

- Stylized visual effects or motion graphics
- Abstract or experimental repetition patterns

- Visual emphasis on rhythm or motion
- Transitions from a single image to a multi-view layout
- Background textures using animated video content

Limitations


- Cannot display different parts of the image in different cells
- No direct numeric control over exact row or column count
- No control over spacing, borders, or offsets
- High grid densities can make motion hard to read due to heavy downscaling



Simulates **halftone printing**, a technique historically used in newspapers, magazines, and early color printing to reproduce images using dots instead of continuous tones.

Because printing presses could not vary ink intensity, tonal variation was achieved by changing **dot size and spacing**. Different ink colors were printed in separate passes, each with its own dot pattern.

This filter is a **visual approximation** of that process. It converts the image into colored dot patterns whose size and density represent brightness and color intensity.



Halftone is a **spatial effect**. It operates per frame and does not analyze motion.

Warning:

This filter has a known bug. Refer to the corresponding [bug report](#) for description, screenshots, and reproduction steps.

Parameters

Radius (0.0 - 100.0%)

Controls the size and spacing of the halftone dots.

- **Low values**

Small, dense dots. The image retains more detail and appears closer to continuous tone.

- **Mid values**

Clearly visible dot structure. Individual dots become distinguishable, and fine detail starts to break into patterns.

- **High values**
Large dots with wide spacing. Image detail is heavily abstracted and coarse.

Note:

- Radius affects **all color channels simultaneously**.
- Increasing Radius reduces spatial detail; this is not a blur but a pattern substitution.
- Perceived strength depends on output resolution; dots appear larger at lower resolutions.

Cyan (0.0 - 100.0%)

Controls the contribution of the cyan halftone layer.

- **0.0%**
Cyan dots are absent.
- **Higher values**
Cyan dots become more visible and dominant in the pattern.

Behavior note:

- This control adjusts dot **presence**, not hue rotation.
- It does not remap colors; it changes how strongly cyan dots participate in the halftone pattern.

Magenta (0.0 - 100.0%)

Controls the contribution of the magenta halftone layer.

- **0.0%**
Magenta dots are absent.
- **Higher values**
Magenta dots become more visible and prominent.

Yellow (0.0 - 100.0%)

Controls the contribution of the yellow halftone layer.

- **0.0%**
Yellow dots are absent.
- **Higher values**
Yellow dots become more visible and influential.

Important note on color controls:

- The sliders do **not** behave like color correction controls.

- They weight how strongly each CMY dot pattern is applied, not the overall image color balance.
- Black (K) is not a separate channel; dark tones emerge from dot overlap.

Keyframes

All parameters can be keyframed.

This enables:

- Animated transitions between continuous-tone and halftone looks
- Gradual increases in dot size or color separation
- Stylized temporal effects, such as print textures appearing or dissolving over time

Parameter interaction

- **Radius** defines the geometric structure of the pattern.
- **Cyan, Magenta, Yellow** define how color information is distributed within that structure.
- High Radius combined with high CMY values produces strong color separation and coarse patterns.
- Lower CMY values can be used to simplify the palette while retaining dot structure.

Visual characteristics

- Image rendered as colored dot patterns
- Reduced fine detail, replaced by geometric texture
- Visible color separation typical of print media
- Strong association with newspaper photos, comic art, and vintage printing

Recommended use cases

- Newspaper or magazine print emulation
- Stylized retro or comic-book visuals
- Abstracting photographic detail into graphic patterns
- Visual transitions referencing print or media reproduction
- Educational demonstrations of halftone printing concepts

Limitations

- Not a physically accurate CMYK simulation
- No control over dot angle, screen rotation, or black (K) channel
- Fine text and thin lines may become unreadable
- Pattern scale is resolution-dependent



The HSL Primaries Video Filter was added in Shotcut 25.01.03


The HSL Primaries Video Filter operates in the HSL color model. It can adjust the Hue, Saturation or Lightness of the 6 primary colors in the image. The hue can be shifted ± 180 degrees. The saturation and lightness can be scaled by a configurable factor.

The Overlap parameter controls how much an adjustment can fade into the adjacent color channels. This can be helpful to avoid banding when the image has a gradient that fades from one color to another.

This example shows an image with a blue car and no adjustments to the filter:



If a shift is applied to the blue and cyan color channels, the car appears purple:



If, instead, the lightness is scaled down, the car appears a darker blue:



See also:

[Hue/Lightness/Saturation Video Filter](#)

[HSL Range Video Filter](#)

The HSL Range Video Filter was added in Shotcut 25.01.03

The HSL Primaries Video Filter operates in the HSL color model. It can adjust the Hue, Saturation and Lightness of the colors in a specified color range. The hue can be shifted ± 180 degrees. The saturation and lightness can be scaled by a configurable factor.

The blend parameter controls how much an adjustment can fade into the adjacent color colors outside the range. This can be helpful to avoid banding when the image has a gradient that fades from one color to another.

This example shows an image with a blue car and no adjustments to the filter:



In this example, the color range was configured to include all the blue/cyan colors that make up the car. Then, the hue was shifted and the saturation increased to make the car appear bright red:



In the image above, you can see a color bar the displays the colors that are included in the range to be adjusted:



In this example, the color range was set to everything except blue and the saturation was scaled to 0% to make the background black and white:



See also:

[Hue/Lightness/Saturation Video Filter](#)


[HSL Primaries Video Filter](#)



Inverts the colors of the image to produce a **negative-style appearance**, similar to photographic film negatives.

This is a **color inversion effect**, not a mathematical channel inversion intended for data or signal processing.

The result resembles analog negative film, where light and dark areas are reversed and colors shift to their complementary counterparts in a perceptual, non-linear way.



Invert Colors is a **spatial effect** applied per frame.

Parameters

This filter has **no user controls**.

The inversion is applied uniformly to the entire image.

Visual behavior and interpretation

Although the effect appears similar to a simple RGB inversion, it should **not** be interpreted as a strict linear or value-based inversion for technical purposes.

Important behavior notes:

- The visual result is intended to look like a **photographic negative**, not a scientific color transform.
- Bright areas become dark; dark areas become bright.
- Colors shift toward their visual complements, but the transformation is not designed for round-trip accuracy (invert → invert ≠ original).
- Midtones may not invert symmetrically due to perceptual color handling.

This makes the filter suitable for **visual styling**, not for analytical workflows.

Historical context

In traditional photography, **negative film** records light inversely:

- Bright areas of a scene become dense (dark) on the film.
- Dark areas remain lighter.
- Colors are recorded as complementary dyes.

This allowed photographic prints to be produced by projecting the negative onto light-sensitive paper, which inverted the tones again to produce a positive image.

The Invert Colors filter references this visual logic, not the chemical process itself.

Visual characteristics

- Negative-like appearance
- Reversed luminance
- Strong color shifts
- High contrast in unexpected regions
- Familiar “film negative” aesthetic

Recommended use cases

- Stylized or experimental visuals
- Flashback, dream, or surreal sequences
- Transitional effects
- Graphic design or abstract motion backgrounds
- Visual emphasis or contrast inversion for artistic purposes

Limitations

- Not suitable for technical color correction or analysis
- Cannot be adjusted or fine-tuned
- Not a physically accurate film simulation
- May produce harsh or unreadable results on footage with extreme contrast



When working with green or blue screens, you often see a “halo” or colored reflection on your subject even after removing the background. These are known as “spills.” Shotcut provides two filters to fix this: **Key Spill: Simple** and **Key Spill: Advanced**.

Key Spill: Advanced – Precision Color Correction

The **Key Spill: Advanced** filter offers high precision for complex shots. It allows you to target specific color ranges using Hue, Chroma, and Intensity parameters to ensure your subject’s natural colors remain untouched.

When to Use

Use this if the “Simple” version is ruining other colors in your video, such as skin tones, hair, or clothing that might be similar to the background color.

How to Apply

1. **Apply Chroma Key First:** Ensure your background is already keyed out using a **Chroma Key** filter.
2. **Add the Filter:** Click the + button in the Filters panel, search for **Key Spill: Advanced**, and place it **below** the Chroma Key filter in the list.
3. **Fine-Tune the Settings:**
 - **Intensity:** Adjust this first to remove the primary spill.
 - **Hue/Chroma Ranges:** Fine-tune these sliders to ensure the filter only affects the edges of your subject and doesn’t “bleed” into other parts of the image.

Pro Tip

If you still see edges after tuning, try adjusting the **Distance** setting in your original **Chroma Key** filter first, then return to the **Key Spill: Advanced** filter for the final cleaning.

When working with green or blue screens, you often see a “halo” or colored reflection on your subject even after removing the background. These are known as “spills.” Shotcut provides two filters to fix this: **Key Spill: Simple** and **Key Spill: Advanced**.

Key Spill: Simple – Quick & Easy Spill Removal

The **Key Spill: Simple** filter is designed for speed. It works by automatically desaturating the key color (green or blue) across the entire frame to neutralize reflections on your subject.

When to Use

Use this for quick fixes when your subject doesn’t have many other colors that might be affected by desaturation.


How to Apply

1. **Apply Chroma Key First:** You must have a **Chroma Key: Simple** or **Chroma Key: Advanced** filter active on your clip first.
2. **Add the Filter:** Click the + button in the Filters panel, search for **Key Spill: Simple**, and add it **below** your Chroma Key filter.
3. **Select Color:** Choose the target color (Green or Blue) to see immediate results.

Pro Tip

If you find that this filter makes your subject look too gray or “washed out,” try adding a **Hue/Lightness/Saturation** filter after it to bring back some of the lost color vibrancy.

Applies a **radial geometric distortion** to the image, allowing manual compensation for lens-like warping or deliberate spatial deformation. Despite its name, this filter does **not** model real lens optics and does not use camera metadata. All adjustments are visual and manual.



With all parameters set to **50.0%**, the filter is neutral and produces no visible change.

Parameters

X Center (0.0 - 100.0%)

Sets the **horizontal pivot** around which the distortion is applied.

- **50.0%**
Distortion is centered horizontally.
- **Lower values**
The pivot shifts left, making distortion stronger on the right side.
- **Higher values**
The pivot shifts right, making distortion stronger on the left side.

Note:

This parameter does not move the image. It defines the origin of the distortion.

Y Center (0.0 - 100.0%)

Sets the **vertical pivot** around which the distortion is applied.

- **50.0%**
Distortion is centered vertically.
- **Lower values**
The pivot shifts upward.
- **Higher values**
The pivot shifts downward.

Correction at Center (0.0 - 100.0%)

Controls **how early distortion begins as you move away from the center**, not the center itself.

- **50.0%**
Neutral reference point.
- **Below 50.0%**
Distortion ramps up **closer to the center**, causing the region just outside the center to stretch outward.
- **Above 50.0%**
Distortion is delayed toward the edges, compressing the region just outside the center.

Important clarification:

The exact geometric center does not move.
This parameter affects the surrounding region, not the center point.

Correction at Edges (0.0 - 100.0%)

Controls **how strongly distortion increases toward the frame edges**, with inverted behavior relative to its name.

- **50.0%**
Neutral reference point.
- **Below 50.0%**
Edge distortion is reduced; edges appear more stretched or pulled.
- **Above 50.0%**
Edge distortion is amplified; compression becomes more pronounced near the borders.

Important clarification:

Despite the name, higher values increase distortion at the edges rather than “correcting” it.

Note on parameter naming

Some parameter names in this filter do not reflect their observable behavior.

Adjustments should be made visually rather than by relying on the parameter names alone.

Keyframes

All parameters can be keyframed.

This allows:

- Progressive distortion or correction over time
- Animated warping effects
- Shifting distortion pivots during a shot

Parameter interaction

- **X Center** and **Y Center** define where the radial distortion originates.
- **Correction at Center** determines where distortion begins along the radius.
- **Correction at Edges** determines how strong distortion becomes near the frame borders.
- Both correction parameters shape a **single continuous distortion curve** and do not operate on isolated regions.

Note:

Adjustments should be evaluated visually using straight lines near the edges as reference.

Visual characteristics

- Radial stretching or compression
- Straight lines may bend, especially near borders
- Edge regions may lose apparent sharpness due to resampling
- Extreme values can expose empty areas at the frame edges

Recommended use cases

- Manual compensation for lens-like distortion
- Aligning architectural or horizon lines by eye
- Creative warping or stylized spatial effects
- Animated distortion using keyframes

Limitations


- There is no interpolation a.k.a. sub-pixel rendering. That means it can cause jagged or block image artifacts.
 - No automatic lens profiles or metadata support
 - Parameter names are misleading relative to observed behavior
 - Not a physical lens simulation
 - May require cropping or scaling after correction
-

See also: [***Fisheye Video Filter***](#) which does support interpolation.

The **Levels** filter adjusts the **tonal range** of an image by remapping how input brightness values are distributed between black, midtones, and white. It is one of the most important tools for **correcting exposure, contrast, and color balance**.

Levels works by redefining:

- What is considered **black**
- What is considered **white**
- How midtones are distributed between them



This filter can be applied globally or per color channel.

Parameters

Channel

Selects which channel the Levels adjustment applies to:

- **Value** (default)
Adjusts overall luminance without changing color balance
- **Red, Green, Blue**
Adjust individual color channels for color correction or creative grading

Histogram

Displays a live **histogram overlay** to assist with adjustment.

Available positions:

- Top left

- Top right
- Bottom left
- Bottom right
- **None**

Notes:

- The histogram is a **visual aid only**
- It **will be rendered into the video output** unless set to **None**
- The histogram size is fixed and constrained to the viewer corner
- **Remember to set Histogram to None before exporting**, unless you need to include it



Input Black (0-255) - keyframeable

Defines the **black point** of the input signal.

- Values below this level are mapped to pure black
- Raising Input Black increases contrast and removes shadow detail

Input White (0-255) - keyframeable

Defines the **white point** of the input signal.

- Values above this level are mapped to pure white
- Lowering Input White increases contrast and removes highlight detail

Gamma (0.01-4.00) - keyframeable

Adjusts the brightness of **midtones** without changing the black or white points.

- **Values < 1.0**
Darken midtones

- **Values > 1.0**
Brighten midtones

Gamma does not change the endpoints, only how values are distributed between them.

Output Black (0-255)

Defines the **output black level**.

- Raises or lowers the darkest output value
- Useful for limiting contrast or matching broadcast ranges

This parameter **cannot be keyframed**.

Output White (0-255)

Defines the **output white level**.

- Compresses or expands the output dynamic range
- Useful for broadcast-safe or stylistic output control

This parameter **cannot be keyframed**.

Parameter interaction

- **Input Black / White** define which parts of the input signal are clipped
- **Gamma** reshapes midtones without clipping
- **Output Black / White** remap the final output range

Channel-specific adjustments affect color balance; Value affects luminance only.

Why Levels is important

Levels is a foundational correction tool because it allows you to:

- Recover contrast from flat or washed-out footage
- Correct underexposed or overexposed material
- Balance color casts using per-channel adjustments
- Prepare footage for further color grading
- Match shots from different cameras or lighting conditions

Incorrect levels can make footage look dull, crushed, clipped, or improperly colored.

Visual characteristics

Typical effects include:

- Increased or reduced contrast

- Brighter or darker midtones
- Corrected black and white points
- Color shifts when adjusting individual channels

Recommended use cases

- Exposure correction
- Contrast normalization
- Color balance correction
- Shot matching
- Pre-grading cleanup

Levels is often one of the **first filters** applied in a color correction workflow.


Limitations

- Histogram overlay must be disabled before export
- Output Black and Output White are not keyframeable
- Aggressive settings can clip shadow or highlight detail
- Not a full color grading solution by itself


Adjust the intensity of the LUT (3D) filter

You can reduce the intensity of a LUT (or any other filter) using filters in this order


1. **Mask: Simple Shape** with Width and Height = 100%
2. **Opacity** (here is where you adjust intensity, including with keyframes)
3. **LUT (3D)** (or other filters)
4. **Mask: Apply**



The video filters beginning with “**Mask:**” let you select a portion of the image or video in which to apply other video filters. There are a few different ones, and this one lets you choose the portion of the image by a color and a distance from the chosen color.



See the “Tip” above. To really make this work you add additional video filters that affect the selected color range followed by the **Mask: Apply** video filter, which combines every thing in between. So, the order of the filters is very important. For example, for secondary color correction, you might put a **Color Grading, Levels**, or **Hue/Lightness/Saturation** video filter in between them:




- **Key color** lets you choose the color to be affected. Click the dropper icon next to the color box to pick a color from the video. The dropper also lets you drag select a small rectangle of the image as well to get an average color for the rectangle to avoid picking a single stray pixel.
- **Distance** helps you choose the range near **Key color** like a threshold.

See also [Mask: Apply](#)



The **Mask: Draw (Glaxnimate)** filter allows you to create complex, hand-drawn, and animated masks using vector tools. It integrates Shotcut with [Glaxnimate](#), an open-source vector animation program that is integrated with Shotcut. This filter is ideal for custom shapes, rotoscoping, and advanced transitions that go beyond the simple shapes provided by the *Mask: Simple Shape* filter. It is basically the same as [Mask: From File Video Filter](#) plus special integration with Glaxnimate.



Parameters

Parameter	Description
Filename	Displays the file path of the current vector drawing (.rawr, .svg, .aep, .lot, .lotte, .json, .riv).
New	Creates a new vector file and opens it in Glaxnimate.
Edit	Opens the currently linked file in Glaxnimate for modifications.
Invert	Reverses the mask; areas that were transparent become opaque and vice versa.
Operation	Determines how this mask interacts with the alpha channel (Overwrite, Add, Subtract, Maximum, Minimum).

The Workflow

Using this filter involves a “round-trip” process between Shotcut and Glaxnimate:

1. **Select a Clip:** Select the clip on the Timeline where you want the mask applied.
 2. **Add Filter:** Go to the **Filters** panel, click **+**, and search for **Mask: Draw (Glaxnimate)**.
 3. **Create a New Drawing:** Click the **New** button. You will be prompted to save a file unless you are using a project folder made using **New Project > Start**.
 4. **Draw in Glaxnimate:** Glaxnimate will open automatically. Use the **Path** tool (Bezier) or shape tools to draw your mask.
- **Fill:** Ensure your shape has a “Fill” applied; otherwise, the mask will only be a thin outline.

- **Alpha:** The alpha values of the colored areas of your drawing will become the masks; the empty background becomes transparent.
 - **Canvas:** Video on Shotcut's timeline appears as the background canvas within Glaxnimate so you can draw upon and around the things in the video. Seeking is synchronized automatically between the two applications.
5. **Save and Return:** In Glaxnimate, go to **File > Save**.
 6. **Update in Shotcut:** Return to Shotcut. The preview player will automatically update to reflect the drawing you just saved.

Rotoscoping and Animation

One of the most powerful features of this filter is the ability to animate the mask over time.

- In Glaxnimate, enable the **Record** mode (the red circle) to create keyframes.
- As you move the playhead in Glaxnimate and adjust the vector points, the mask will morph accordingly.
- Shotcut will sync the timing of the Glaxnimate animation with the clip on your timeline.

Mask Operations

The **Operation** setting is crucial when using multiple masks on the same clip:

- **Minimum:** For each pixel, chooses the lesser the alpha value, more translucent or transparent, between the existing alpha channel and the new mask.
- **Maximum:** For each pixel, chooses the lesser the alpha value, more translucent or transparent, between the existing alpha channel and the new mask.
- **Overwrite:** This mask replaces all previous alpha channel data.
- **Add:** Adds the shape of this mask to any existing masks.
- **Subtract:** Uses this shape to “cut a hole” out of previous masks.

Important Notes

- **Bundled Software:** You do not need to install Glaxnimate separately; it is included with the Shotcut if you are using one of our builds.
- **Preview Resolution:** If you find the drawing doesn't align perfectly, ensure your Shotcut **Preview Scaling** is set to “None” for the most accurate alignment during the drawing process.
- **Performance:** Complex vector paths with many points may slow down preview playback. Use **Proxy** mode to improve performance while editing.


See also: [Mask: From File Video Filter](#)



Applies a **luminance- or alpha-based mask** to a clip using built-in static mask presets or an external image or video file.

The mask determines which areas of the clip are affected, based on brightness or transparency values.

All built-in mask presets are **static**. Animation requires keyframing, clip motion, or a video file used as the mask.



NOTE:

This filter operates is not available in GPU processing mode.

Mask source behavior

The mask source can be:

- A built-in **static mask preset**
- **Any** external media image, video, drawing, or animation media file supported by Shotcut

Mask interpretation rules:

- **Brightness mode**
Dark areas of the mask make the current clip visible.
Bright areas make the underlying clip visible.
Gray values produce a gradual transition between the two.

- **Alpha mode**
Transparent areas of the mask make the current clip visible.
Opaque areas make the underlying clip visible.

When a video file is used as the mask, its values are evaluated per frame.

Parameters

File (preset / file selector)

Selects the mask source.

The dropdown includes:

Built-in geometric masks

- **Bar Horizontal / Bar Vertical**
A single rectangular band across the frame.
- **Barn Door Horizontal / Vertical / Diagonal SW–NE / Diagonal NW–SE**
Four-panel shapes opening from the center, similar to camera barn doors.
- **Diagonal Top Left / Diagonal Top Right**
Half-frame diagonal masks.
- **Barn V Up**
V-shaped opening expanding upward from the bottom.
- **Iris Circle**
Circular opening centered in the frame.
- **Double Iris**
Two mirrored circular openings.
- **Iris Box**
Rectangular opening expanding from the center.
- **Box Bottom Right / Box Bottom Left / Box Right Center**
Rectangular masks anchored to specific edges or corners.
- **Clock Top**
Radial, clock-style sector opening from the top.

Pattern-style masks

- **Matrix Waterfall Horizontal / Vertical**
Repeating rectangular segments arranged in rows or columns.
- **Matrix Snake Horizontal / Vertical**
Sequential block patterns arranged in a snake-like order.

- **Matrix Snake Parallel Horizontal / Vertical**


Multiple parallel snake patterns.

These masks are **not animated**. Any **motion must be created via keyframes** or by using a video mask.

Custom...

Opens a file browser to select an external image or video file.

- Supports common image formats and video files.
- When a custom file is selected, its **filename is displayed below the dropdown list**.



Invert (checkbox)

Inverts the mask interpretation.

- Revealed areas become blocked
- Blocked areas become revealed

Applies after Channel selection.

Reverse (checkbox)

Reverses the **directional orientation** of the mask.

- Affects directional presets (bars, doors, diagonals)
- Example: a right-side mask becomes a left-side mask

This does **not** invert luminance or alpha values.

Channel (radio buttons)

Defines how the mask source is interpreted.

- **Brightness**

Uses the luminance of the mask source.

Bright areas reveal the effect; dark areas block it.

- **Alpha**

Uses the alpha channel of the mask source.

Opaque areas reveal the effect; transparent areas block it.

Important notes:

- Alpha mode requires the mask file to contain an alpha channel.
- In Alpha mode, RGB color information in the mask is ignored.
- Most video files do not contain an alpha channel unless explicitly encoded.

Threshold (checkbox + slider, 0.00–100.00%, keyframeable)

Converts the mask into a **binary (hard) mask**.

- **Unchecked**

Continuous grayscale or alpha values are used.

- **Checked**

Values are clipped to fully on or fully off.

Slider behavior:

- **Lower values**

More of the mask becomes active.

- **Higher values**

Only the brightest or most opaque areas remain active.

Softness (0.00–100.00%)

Controls edge feathering **after thresholding**.

- Only active when **Threshold** is enabled.
- Low values produce hard edges.
- Higher values produce smoother transitions.

Softness does not blur the mask image itself.

Keyframes

Only the **Threshold** slider can be keyframed.

This enables:

- Animated reveals and wipes using static masks
- Time-based control when using video masks

Parameter interaction

- **File** defines the mask shape or pattern.
- **Channel** defines whether luminance or alpha drives the mask.
- **Invert** swaps masked and unmasked regions.
- **Reverse** flips directional orientation.
- **Threshold** forces hard masking.
- **Softness** smooths threshold edges.

All operations are applied per frame.

Visual characteristics

- Clean geometric masking with presets
- Arbitrary shapes with custom images
- Animated masking only when using video files or keyframes
- Mask edge quality depends on source resolution and Softness

Recommended use cases

- Custom wipes and transitions
- Layer compositing and selective effects
- Masks drawn in external image editors
- Video-driven animated masks
- Precise, repeatable masking workflows

Limitations

- CPU-only filter
- Built-in masks are static
- No built-in mask animation controls
- Mask quality depends on source resolution
- Video masks can increase processing load

Built-in mask presets: quick reference

Linear / band masks

- **Bar Horizontal** — Horizontal rectangular band across the frame
- **Bar Vertical** — Vertical rectangular band across the frame

Barn door masks

- **Barn Door Horizontal** — Four panels opening left and right from the center

- **Barn Door Vertical** — Four panels opening up and down from the center
 - **Barn Door Diagonal SW–NE** — Diagonal barn-door opening from southwest to northeast
 - **Barn Door Diagonal NW–SE** — Diagonal barn-door opening from northwest to southeast
 - **Barn V Up** — V-shaped opening expanding upward from the bottom
-

Diagonal masks

- **Diagonal Top Left** — Diagonal half-frame mask anchored to the top-left
 - **Diagonal Top Right** — Diagonal half-frame mask anchored to the top-right
-

Iris / aperture masks

- **Iris Circle** — Circular opening centered in the frame
 - **Double Iris** — Two mirrored circular openings
 - **Iris Box** — Rectangular opening expanding from the center
-

Box / edge-anchored masks

- **Box Bottom Left** — Rectangular mask anchored to the bottom-left
 - **Box Bottom Right** — Rectangular mask anchored to the bottom-right
 - **Box Right Center** — Rectangular mask centered vertically on the right edge
-

Radial masks

- **Clock Top** — Radial, clock-style sector opening from the top
-

Matrix / pattern masks


- **Matrix Waterfall Horizontal** — Repeating horizontal rectangular segments
 - **Matrix Waterfall Vertical** — Repeating vertical rectangular segments
 - **Matrix Snake Horizontal** — Sequential horizontal block pattern
 - **Matrix Snake Vertical** — Sequential vertical block pattern
 - **Matrix Snake Parallel Horizontal** — Multiple parallel horizontal snake patterns
 - **Matrix Snake Parallel Vertical** — Multiple parallel vertical snake patterns
-

Notes

- All built-in presets are **static**.
 - Directional behavior can be modified using **Reverse**.
 - Edge hardness is controlled using **Threshold** and **Softness**.
 - Animation requires keyframes or a video file used as the mask.
-

See also: [Mask: Draw \(Glaxnimate\) Video Filter](#)

The **Mask: Simple Shape** filter is one of the most versatile tools in your Shotcut toolkit. Think of it as a way to “cut out” a specific part of your video to apply changes only to that area. While other filters might select parts of a video based on color (like a Green Screen), this filter selects based on geometry, using shapes like squares or circles to define your workspace.



When you apply this filter, you are essentially telling Shotcut: *“Only focus on what is inside this window.”*

How You Control the Shape

Once you select the filter, you can customize your “window” using these settings:

- **Shape:** Pick the best fit for your subject—Rectangle, Ellipse, Triangle, or Diamond.
- **Position & Size:** Use the Horizontal/Vertical and Width/Height sliders to place the mask exactly where you need it.
- **Rotation:** Tilt the shape to align with angled objects, like a computer screen or a tilted horizon.
- **Softness:** This is the secret to professional edits. It blurs the edges of your shape so the effect blends naturally into the rest of the video rather than looking like a harsh cutout.

Practical Ways to Use It

- **Privacy:** Place a soft-edged circle over a face or a license plate to hide sensitive information.

- **Selective Editing:** Want to make the sky bluer without affecting the grass? Use a rectangle mask on the top half of your frame.
 - **Spotlighting:** Brighten a person in a dark room by placing an ellipse over them and increasing the exposure.
 - **Split Screens:** Use a rectangle mask to “crop” a clip so you can show two different videos side-by-side.
-

The “Sandwich Logic” Workflow

In Shotcut, simply creating a mask isn’t enough; you have to tell the software how to handle the effects inside that shape. The most reliable method is the **Sandwich Logic**. You build your filter stack in three layers, just like a sandwich.

1. The Top Bun: Mask: Simple Shape

This is always your first filter. It defines the “where.”

- **What to expect:** When you add this, the area outside your shape will usually turn black. Don’t worry—this means it’s working! You have successfully “cut out” your shape from the rest of the footage.

2. The Filling: Your Creative Effects

Any filter you place directly below the mask is the “filling.” These are the changes that will appear **only** inside your shape.

- **Popular Fillings:** You might add a **Blur** to hide a face, a **LUT** for color grading, or **Brightness** to create a spotlight.
- **Pro Tip:** Add an **Opacity** filter here. It acts as a master “strength” slider for all your other fillings.

3. The Bottom Bun: Mask: Apply

This is the most important step that many editors forget. You **must** add the **Mask: Apply** filter at the very end of your list.

- **What it does:** This “seals” the sandwich. It tells Shotcut to stop applying effects to the shape and to bring back the rest of your original video. The black background will disappear, and your effects will now be perfectly contained within your shape.
-

Your Filter Stack Cheat Sheet


To make sure your mask works every time, your Filters Panel should follow this exact order:

Order	Filter Name	The Goal
1	Mask: Simple Shape	Define the area (The Window)

Order	Filter Name	The Goal
2	Any Effect (e.g., Blur)	The visual change (The Filling)
3	Any Effect (e.g., Color)	Additional changes
4	Mask: Apply	Seal the mask (The Bottom Bun)

Why follow this? Without the “Bottom Bun” (**Mask: Apply**), your video will often be stuck with a black box around your effect. This workflow prevents that error and gives you professional, selective control over your edits.


A Mirror Video Filter, is a digital effect that flips a video horizontally. If you've ever wondered **why the letters on your t-shirt appear backwards in a selfie, this is why**. Some phones let you choose whether to keep the mirrored image or correct it in the camera settings.



Divides the image into rectangular blocks and replaces each block with a single averaged color.

The result is a pixelated, low-resolution appearance where fine detail is intentionally removed.

The term *mosaic* reflects a long-standing image construction principle: building pictures from discrete units rather than continuous detail. Ancient mosaics used small, uniformly shaped pieces (*tesserae*) to form larger images, a logic later echoed in early color photography where images were filtered through dense grids of tiny colored elements. Modern digital mosaic effects apply the same idea computationally by grouping pixels into visible blocks.



Over time, the look became a deliberate tool for obscuring detail or stylizing footage, particularly in broadcast and documentary contexts.

Parameters

Width (0.0% - 20.0%)

Controls the **horizontal block size** as a percentage of the image width.

- **0.0**
No horizontal pixelation.
- **Low values (1.0 - 5.0)**
Fine pixelation with partial detail retention.
- **Medium values (6.0 - 12.0)**
Clear block structure; faces and text become difficult to identify.
- **High values (13.0 - 20.0)**
Very coarse blocks; only large shapes and colors remain.

This parameter defines how many pixels are grouped left-to-right into a single block.

A value of 10% means each mosaic block spans roughly one-tenth of the frame width.

Height (0.0% - 20.0%)

Controls the **vertical block size** as a percentage of the image height.

- **0.0**
No vertical pixelation.
- **Low values (1.0 - 5.0)**
Subtle vertical grouping.
- **Medium values (6.0 - 12.0)**
Strong vertical block structure.
- **High values (13.0 - 20.0)**
Extremely coarse vertical grouping.

This parameter defines how many pixels are grouped top-to-bottom into a single block.

A value of 10% means each mosaic block spans roughly one-tenth of the frame height.

Keyframes

Both **Width** and **Height** can be keyframed.

This enables:

- Tracking moving subjects when combined with a mask
- Gradual reveal or conceal of detail
- Dynamic censorship or stylized transitions

Parameter interaction

- Width and Height operate independently.
- Equal values produce square blocks.
- Unequal values produce rectangular blocks, which may better match faces, license plates, or signage.

The effective loss of detail is dominated by the **larger** of the two values.

Visual characteristics

- Hard-edged rectangular blocks
- Each block uses the average color of its underlying pixels

- Complete removal of fine detail within blocks
- No blur or transparency; edges remain discrete

This is a purely **spatial** effect with no temporal behavior.


Recommended use cases

- Obscuring faces, license plates, or sensitive information
- Anonymizing individuals in documentary or journalistic footage
- Intentional low-resolution or abstract visual styles

Face or object hiding (with Mask: Simple Shape)

When combined with **Mask: Simple Shape**:


1. Apply **Mosaic** to the clip
2. Add **Mask: Simple Shape** above it
3. Add **Mask: Apply** below it (to see the effect)
4. Restrict the mosaic to a defined region
5. Keyframe the mask to follow the subject if needed



This approach allows localized obfuscation while preserving detail elsewhere. It is available as **Filter Set > Obscure With Mosaic** that comes with Shotcut.

Limitations

- No automatic subject or face detection
- Block size is abstract and not expressed in pixels
- No control over block shape beyond Width and Height
- Coarse blocks may still reveal motion through color changes



This filter is required to use motion tracking in other filters. It only does the preview and analysis and is not intended to include effects but rather drive other video filters...

See the **Load Keyframes from Motion Tracker** button in the following video filters:

- **Corner Pin**
- **Crop: Rectangle**
- **GPS Text**
- **Mask: Simple Shape**
- **Size, Position & Rotate**
- **Size & Position (GPU)**
- **Spot Remover**
- **Text: Rich**
- **Text: Simple**
- **Timer**


Notes

- This filter first appeared in version 23.05.14.
- The **Motion Tracker** can be on a different clip than the clip with the above filters.
- The dialog only shows **Motion Trackers** that have successfully completed analysis.
- Click **Reset** in the dialog to remove the keyframes added by motion tracker and restore the state of these parameters to before motion tracking.
- Clicking **Apply** in the dialog more than once automatically resets the affected parameters before adding new keyframes; so, you do not need to reset manually between changes.

- You can edit the added keyframes, but if you reapply the motion tracker your changes are discarded.
- To change the initial position of something after applying tracking you should **Reset**, make changes, and then **Apply** again.
- Disable the **Motion Tracker** when your changes and its analysis are done (or turn off its preview). I do not recommend removing this filter when you are done loading keyframes because if you need to make an adjustment, this filter must still be available to list it in the dialog and get its tracking data.



The **Nervous** filter creates a **jittery, unstable motion effect** by randomly replacing the current frame with frames from nearby points in time. This produces a nervous, erratic visual appearance often associated with analog glitches, damaged recordings, or psychological tension effects.



This filter operates **temporally** and does not modify color or geometry.

Parameter

Amount (2–30 frames)

Controls the **temporal range**, in frames, from which replacement frames may be selected.

- Lower values
 - Subtle jitter using nearby frames
 - Motion appears slightly unstable
- Higher values
 - Stronger, more chaotic jitter
 - Frames may jump noticeably forward or backward in time

This parameter defines **how far from the current frame the filter is allowed to sample**.

How it works

For each output frame, the filter may:

- Keep the current frame, or
- Substitute it with a frame randomly chosen from within the specified frame range

The randomness creates:

- Temporal discontinuity

- Repeated or skipped frames
- A “nervous” visual rhythm

Visual characteristics

Typical effects include:

- Jittery or trembling motion
- Irregular frame repetition
- Perceived instability even in static shots
- Increased intensity as the frame range increases

The effect is independent of image content.

Example: static camera with moving subjects

When applied to footage shot with a **static camera** where **people or objects are moving**, the Nervous filter causes:

- The background to appear stable
- Moving subjects to jitter, stutter, or briefly jump backward and forward
- Moving elements appear fragmented or unstable, while the camera and background remain static.

This effect occurs because the filter substitutes frames at different moments in time, disrupting motion continuity while leaving static elements largely unchanged.

Recommended use cases

- Stylized glitch or distortion effects
- Psychological or tension-driven visuals
- Simulating damaged or unstable recordings
- Music videos or experimental edits
- Only moving elements are affected and appear to jitter; a moving background is affected as well.


Limitations

- Does not preserve motion continuity
- Not suitable for natural or smooth motion
- Effect strength depends on frame rate
- Parameter cannot be keyframed

Rearranges the image by **splitting it into regions and swapping their positions**, producing a deliberate misalignment effect reminiscent of analog signal loss or desynchronization.

The filter does not distort geometry or color; it **repositions rectangular regions** of the frame.

No Sync is a **spatial effect**. It operates per frame and does not analyze motion or time.



With both parameters at **50.00%**, the image is divided into four equal quadrants and swapped diagonally.

Parameters

Vertical (0.00 - 100.00%)

Controls the **vertical split position** of the image.

- **50.00%**
The frame is split evenly horizontally.
The bottom quadrant swaps with the top.
- **Below 50.00%**
The lower portion of the image becomes smaller.
The split line moves downward, giving more vertical space to the top region.
- **Above 50.00%**
The upper portion becomes smaller.
The split line moves upward.
- **0.00% or 100.00%**
The image is not split.

Horizontal (0.00 - 100.00%)

Controls the **horizontal split position** of the image.

- **50.00%**
The frame is split evenly vertically, contributing to the four-quadrant swap.
- **Below 50.00%**
The right portion of the image becomes larger.
The split line moves left, giving more horizontal space to the right region.
- **Above 50.00%**
The left portion becomes larger.
The split line moves right.
- **0.00% or 100.00%**
The image is not split.

Keyframes

Both **Vertical** and **Horizontal** parameters can be keyframed.

This allows:

- Animated region swapping
- Progressive desynchronization effects
- Glitch-like transitions between layouts

Parameter interaction

- **Vertical** and **Horizontal** define the split positions along their respective axes.
- At **50.00% / 50.00%**, the image is divided into four equal regions that swap diagonally.
- Moving either parameter away from 50.00% changes the relative size of the regions before they are swapped.
- Extreme values collapse the effect into a simple two-part swap along one axis.

Visual characteristics

- Hard rectangular splits
- Abrupt spatial jumps
- No blending or feathering between regions
- Clear, mechanical rearrangement of image areas
- Strongly reminiscent of video sync or signal errors

Recommended use cases

- Glitch and analog-error effects
- Stylized transitions

- Visual emphasis or disruption
- Abstract or experimental video design
- Simulating desynchronization without distortion or blur


Limitations

- No control over edge softness
- No rotation or scaling of regions
- Effect is purely positional
- Can be visually jarring on detailed footage



Noise: Fast filter adds synthetic **colored noise** to the image using a lightweight, performance-oriented algorithm.

The filter introduces random variation directly into color channels, producing visible RGB noise rather than luminance-only grain.



This filter prioritizes speed over physical accuracy. It is designed for real-time playback and simple texture breakup, not for film-accurate noise simulation.

Parameters

Amount (0.0 - 100.0%)

Controls the **intensity of the noise contribution**.

- **0%**
No noise is added.
- **Low values (1% - 10%)**
Subtle texture breakup. Useful for reducing banding or digital smoothness.
- **Medium values (10% - 40%)**
Clearly visible colored noise. Flat areas begin to shimmer with chromatic variation.
- **High values (40% - 100%)**
Strong, obvious noise. Image detail becomes secondary to noise structure.

Important behavior

- Noise is applied **per color channel**, not as pure luminance noise.
- The noise has no film-grain weighting; shadows, midtones, and highlights are affected similarly.

Note:

Amount controls how much random color variation is injected, not the size or scale of the noise.

Keyframes

This filter does **not** support keyframes.

Noise intensity and behavior remain constant over time.

Type of noise (technical note)

- **Colored (RGB) noise**, not monochrome
- No correlation between channels
- No film-style luminance weighting
- No temporal evolution control

Editor's note:

Visually, this noise resembles GIMP's **CIE LCh Noise** filter at low chroma correlation, although the underlying algorithms are different. The similarity is only visual, not algorithmic.

This is not:

- Film grain
- Sensor noise simulation
- Dither (no pattern, no tone approximation)

Visual characteristics

- Fine, high-frequency noise
- Random color speckling
- Even distribution across the frame
- No structured pattern or clustering

Noise is especially visible in:

- Flat color regions
- Gradients
- Shadows

Recommended use cases

- Breaking up gradient banding after heavy compression
- Adding texture to flat or synthetic imagery
- Masking **posterization** artifacts
- Creating digital or degraded aesthetics

Common combinations


- **Noise: Fast + Posterize**
Reduces harsh band edges by introducing stochastic variation.
- **Noise: Fast + Dither (low Levels)**
Adds randomness to ordered dithering patterns, reducing visible regularity.
- **Noise: Fast + Gaussian Blur (very subtle)**
Softens chromatic speckling while retaining texture.
- **Noise: Fast + Color Grading**
Can be used deliberately to introduce instability or analog-digital hybrid looks.

Limitations

- Not film-accurate grain
- No control over noise color balance or channel weighting
- Can amplify compression artifacts at high Amount values
- May shimmer noticeably after export or platform re-encoding

Adds synthetic visual noise to the image to introduce randomness, texture, or signal degradation.

Unlike **film grain** simulations, this filter produces a **digital-style noise** that can be animated over time.



Noise: Keyframe is a **spatial effect** evaluated per frame.

Parameters

Amount (0.0 - 100.0%)

Controls the intensity of the noise applied to the image.

- **Low values (≈0.0 - 20.0%)**
Subtle noise. Fine speckling is visible primarily in flat or midtone areas.
- **Mid values (≈20.0 - 60.0%)**
Clearly visible noise. Texture becomes apparent across most of the image.
- **High values (≈60.0 - 100.0%)**
Strong noise. A structured pattern may become visible, with noise appearing to align along rows and columns.

Notes:

- The noise is **not monochrome**; particles contain subtle color variation.
- At high values, the noise can appear **grid-like or patterned**, rather than fully random.
- Noise is applied uniformly across the frame; it does not respond to luminance or edges.

Keyframes

The **Amount** parameter can be keyframed.

This enables:

- Gradual buildup or reduction of noise
- Pulsing or rhythmic noise effects
- Time-based transitions between clean and degraded images

Visual characteristics

- Fine to coarse digital noise
- Slightly colored noise particles
- Uniform distribution across the frame
- At high values, visible structure rather than purely random grain
- Can reduce perceived image clarity and contrast

Recommended use cases

- Glitch or digital degradation effects
- Simulating signal interference or instability
- Adding texture to flat or synthetic imagery
- Transitional effects using animated noise
- Abstract or experimental visuals

Tips and usage notes

- Use **low Amount values** for subtle texture; higher values quickly become dominant.
- Evaluate the effect at full resolution, as noise structure is resolution-dependent.
- Keyframing small fluctuations in Amount often looks more natural than large jumps.
- This filter is not a replacement for film grain or photographic noise simulations.


Limitations

- Noise pattern may appear structured at high values
- No control over noise color, scale, or distribution
- No luminance-based or adaptive behavior
- Can introduce visible banding or patterning when overused



The **Old Film: Dust** filter simulates **dust particles and fine debris** commonly visible in analog film projection.

In physical film formats, dust, lint, and hair could settle on the film surface or inside the projector gate, briefly appearing as spots or short streaks as the film passed through the light path.



This filter recreates those transient artifacts for **creative and historical effect**, not restoration.

Parameters

Size (1 - 100)

Controls the **apparent size of dust particles**.

- Lower values
Small, fine dust specks and hair-like artifacts
- Higher values
Larger, more noticeable debris

This parameter affects the visual scale of each particle.

Note:

High values can quickly become unrealistic.

Amount (1 - 400)

Controls the **number of dust particles** visible over time.

- Lower values
Sparse, occasional dust artifacts

- Higher values
Dense dust activity, with many particles appearing

This parameter controls the effective footprint of dust particles.

Higher values can cause particles to overlap and cluster, increasing the perceived density of dust.

Usage notes

- The effect is **temporal**: dust particles appear, move, and disappear over time.
- Evaluating the filter on a **single frame is misleading**.
- For accurate preview, **play the video** to observe particle motion and frequency.

Subtle settings usually produce the most realistic results.

Visual characteristics

Typical effects include:

- Small specks briefly appearing in the frame
- Short-lived artifacts that move or flicker
- Random distribution over time
- Greater realism at lower Size and Amount values

The artifacts are **not fixed to the image content** and change from frame to frame.

Dust vs Scratches

Although both filters simulate film damage, they represent **different physical causes**:

- **Old Film: Dust**
Simulates loose particles (dust, lint, hair) that intermittently block or scatter light during projection.
Artifacts are small, transient, and randomly distributed.
- **Old Film: Scratches**
Simulates permanent damage to the film surface itself.
Artifacts are elongated, usually vertical, and persistent across frames.

Using both together can produce a more complete aged-film effect.

Recommended use cases

- Vintage or archival-style visuals
- Simulating projected film rather than film stock alone

- Subtle degradation effects
- Combining with other *Old Film* filters (Scratches, [Projector](#), [Grain](#),...)


Limitations

- Parameters are **not keyframeable**
- Procedural effect, not derived from image content
- High values may obscure details
- Intended for stylized use



This filter emulates the visible texture caused by photographic film stock rather than digital sensor noise. It adds film-style grain to the image by introducing luminance noise.

In photochemical film, grain originates from **silver halide crystals** embedded in the emulsion. These crystals vary in size and distribution, and when exposed and developed, they create a fine, irregular texture. Faster (high-ISO) film stocks used larger crystals, producing more visible grain. The effect was intrinsic to the medium and varied over time, exposure, and development.



In modern workflows, grain is often reintroduced intentionally to reduce the sterile appearance of digital footage, **unify mixed sources**, or support historical or archival aesthetics.

Parameters

Noise (1% - 200%)

Controls the **intensity of the grain pattern**.

- **Low values (1% - 20%)**
Subtle grain, often barely perceptible on flat areas. Suitable for breaking up digital smoothness without drawing attention.
- **Medium values (30% - 100%)**
Clearly visible grain across midtones and highlights. Typical for film-like texture.
- **High values (100% - 200%)**
Strong, exaggerated grain. The upper range exists to:
 - Simulate very coarse or pushed film stock
 - Remain visible after downscaling, compression, or export
 - Allow creative or deliberately degraded looks

Why 200% exists

Grain is frequently reduced by:

- Video compression
- Scaling
- Noise reduction downstream

Allowing values above 100% ensures the effect survives these processes. It is not a “double noise” mode, but a practical headroom for post-processing loss.

Brightness (1 - 400)

Controls how strongly the grain affects **luminance levels**.

- **Low values (1 - 100)**

Grain is subtle and primarily affects midtones and highlights.

- **Medium values (100 - 200)**

Grain becomes more visible across a wider tonal range.

- **High values (200 - 400)**

Grain is pushed strongly into brighter regions, increasing contrast between grain and image.

Important behavior

- **Pure black pixels remain pure black**, even at maximum values.
- Grain is applied **additively**, not by lifting the black level.
- This matches film behavior: unexposed or fully opaque areas do not reveal grain.

Note:

Brightness does not brighten the image.

It controls how visible the grain is relative to luminance.

Keyframes

This filter does **not** support keyframes.

The grain pattern and its intensity remain constant over time.

Parameter interaction

- **Noise** controls *how much grain exists*.
- **Brightness** controls *where the grain is visible* in the tonal range.

High Noise with low Brightness produces dense but restrained grain.

Lower Noise with high Brightness emphasizes grain primarily in brighter areas.

Visual characteristics

- Monochrome, film-style grain
- Even distribution across the frame
- No chromatic noise
- Static behavior per frame (no evolving pattern)

Grain is most noticeable in:

- Flat midtones
- Bright areas
- Soft gradients

Recommended use cases

- Adding film texture to clean digital footage
- Matching modern shots to archival or historical material
- Reducing banding and overly smooth gradients
- Supporting period aesthetics when combined with other **Old Film** filters:
 - **Old Film: Projector** (instability, flicker)
 - **Old Film: Dust** (physical debris)
 - **Old Film: Scratches** (emulsion damage)


Used together, these filters approximate different physical artifacts of film capture and projection rather than a single effect.

Limitations

- Grain pattern is not animated; it does not evolve over time
- No control over grain size or color
- Can be partially suppressed by export compression at low Noise values

The **Old Film: Projector** filter simulates **mechanical and optical artifacts** introduced during analog film projection, particularly common in **home and educational projectors** (8 mm, Super 8, 16 mm).

Unlike film **grain** or **scratches**, these artifacts do **not originate** from the film stock itself, but from the **projection process**: unstable film transport, uneven illumination, mechanical vibration, and inconsistent film development. The result is an image that **jumps, flickers, and fluctuates in brightness**, often making projection highly unstable compared to modern digital playback.



This filter recreates those characteristics for **creative and nostalgic effect**.

Parameters

Vertical amount (0 - 200)

Controls the **strength of vertical image displacement**.

- Lower values
Subtle vertical instability
- Higher values
Strong vertical jumping, similar to misaligned film transport

This simulates frame misregistration caused by imperfect sprocket movement.

Vertical frequency (0% - 100%)

Controls **how often vertical movement occurs**.

- Lower values
Occasional jumps
- Higher values
Frequent or continuous vertical motion

Brightness up (0 - 100)

Controls the **maximum upward brightness fluctuation**.

This simulates brief over-exposure caused by uneven projection light or shutter timing.

Brightness down (0 - 100)

Controls the **maximum downward brightness fluctuation**.

This simulates dimming caused by shutter occlusion, lamp instability, or film density variations.

Brightness frequency (0% - 100%)

Controls **how often brightness fluctuations occur**.

Higher values produce noticeable **flicker**, a hallmark of analog projection.

Uneven develop up (0 - 100)

Controls **bright variations** caused by uneven film development.

Simulates areas where film chemistry produced lighter regions.

Uneven develop down (0 - 100)

Controls **dark variations** caused by uneven film development.

Simulates density inconsistencies across the film strip.

Uneven develop duration (0 - 1000)

Controls how long uneven development artifacts persist before changing.

- Lower values
Rapid, unstable changes
- Higher values
Slow, drifting brightness irregularities

Note:

Uneven develop duration controls how long the defect stays, **not** how fast it moves.

Parameter interaction

- **Vertical amount + frequency** simulate mechanical instability
- **Brightness controls + frequency** simulate projector flicker
- **Uneven develop controls + duration** simulate chemical inconsistencies in film stock

Combined, these effects recreate the **imperfect, unstable viewing experience** of projected analog film.

Visual characteristics

Typical effects include:

- Vertical image jitter or jumping
- Irregular brightness flicker
- Uneven exposure across frames
- Persistent brightness patches that drift over time

The image may appear **distracting or unstable by modern standards**, intentionally reflecting historical projection limitations.

Historical context

In home and school environments, film projection was often:

- Mechanically imprecise
- Poorly aligned
- Affected by worn projectors and lamps
- Sensitive to film wear and chemical inconsistencies

This filter reproduces that experience rather than the idealized film image.

Historical preset suggestion

8 mm Classroom Projector

Simulates mechanical and optical instability typical of home and classroom film projectors.

Suggested values (starting point)

- Vertical amount: 35 - 60
- Vertical frequency: 15 - 25%
- Brightness up: 10 - 20
- Brightness down: 15 - 25

- Brightness frequency: 20 - 35%
- Uneven develop up: 10 - 20
- Uneven develop down: 15 - 30
- Uneven develop duration: 300 - 600

Note:

These presets produce mild vertical jitter and irregular flicker typical of worn 8 mm educational film projection on poorly maintained projectors.

Recommended use cases

- Authentic vintage or archival simulations
- Educational or historical reenactments
- Memory or nostalgia sequences
- Stylized degradation effects
- Combining with other *Old Film* filters (Scratches, [Dust](#), Grain)


Limitations

- Parameters are **not keyframeable**
- Effect is uniform across the frame
- Intended for stylization, not restoration
- Extreme values may reduce readability



The **Old Film: Scratches** filter simulates vertical scratches caused by physical wear in analog film.

In traditional motion-picture formats (such as 8 mm, 16 mm, or 35 mm film), images were recorded on a physical strip that could be scratched, abraded, or damaged through handling, projection, or aging.



This filter recreates those artifacts by overlaying animated scratch marks for a vintage film appearance. It is intended for **creative effect**, not restoration.

Parameters

Width (1–100)

Controls the **thickness of the scratches**.

- Lower values
Thin, hairline scratches
- Higher values
Wider, more visible scratches

Amount (1–100)

Controls the **density of scratches** visible in the frame.

- Lower values
Few, isolated scratches
- Higher values
Many scratches appearing simultaneously

At high values, scratches may overlap and merge, producing **broader visual artifacts** rather than distinct lines.

Darkness (1–100)

Controls the **intensity of dark scratches**.

- Lower values
Dark scratches are faint
- Higher values
Dark scratches become dominant

This simulates scratches that reduce or block light in the film material.

Lightness (0–100)

Controls the **intensity of bright scratches**.

- 0
No light scratches are visible
- Higher values
Bright scratches become dominant

This simulates scratches that reflect or scatter light during projection.

Parameter interaction

- **Width** controls scratch thickness
- **Amount** controls scratch density
- **Darkness** and **Lightness** operate on the **same scratch pattern**

Darkness and Lightness are **mutually influential**: increasing one reduces the visual influence of the other. Scratches are rendered as either predominantly dark or predominantly light, reflecting how physical scratches interact with light rather than appearing as both simultaneously.

At high **Amount** and/or high **Width** values, overlapping scratches may interact with color channels and produce **chroma artifacts** (magenta & green). This can resemble **analog videotape degradation** (VHS/Betamax) rather than film scratches and is considered a stylized effect.

Visual characteristics

Typical effects include:

- Vertical scratch lines of varying thickness
- Animated motion over time
- Either dark or bright scratch dominance

- Possible color artifacts at high scratch density

The scratches are an **overlay effect** and do not alter the underlying image content.

Recommended use cases


- Vintage or retro film simulations
- Archival-style visuals
- Historical reenactments
- Stylized degradation effects
- Combination with other *Old Film* filters (grain, dust, Technicolor)

Limitations

- Procedural effect, not film restoration
- Does not react to image content
- Uniform application across the frame
- High Amount values may produce non-film-like artifacts



The **Old Film: Technocolor** filter applies a stylized color transformation inspired by historical Technicolor™ film processes. It modifies color relationships rather than performing technical color correction or restoration.



The filter works by **shifting colors along two opposing color axes**, altering the balance between complementary hues to produce a saturated, vintage film look.

Parameters

Green → Red (–300 to 300)

Controls the balance between the **green and red color channels**.

- Negative values shift the image toward **green**
- Positive values shift the image toward **red**

This adjustment affects:

- Skin tones
- Vegetation
- Warm vs. cool color balance in midtones

Higher absolute values exaggerate the separation between green and red hues, contributing to a stronger stylized effect.

Yellow → Blue (–300 to 300)

Controls the balance between **yellow and blue hues**, corresponding broadly to a warm–cool axis.

- Negative values shift the image toward **yellow**
- Positive values shift the image toward **blue**

This adjustment affects:

- Overall warmth or coolness

- Highlights and skies
- Perceived color temperature

As with the Green → Red control, larger values produce more pronounced color shifts.

How the controls work together

The two sliders operate on **independent color axes**:

- **Green** ↔ **Red** adjusts chromatic balance along one axis
- **Yellow** ↔ **Blue** adjusts chromatic balance along a second, perpendicular axis

Together, they allow selective exaggeration or suppression of specific color relationships rather than a uniform saturation change.

Visual characteristics

Typical results include:

- Increased color separation
- Stronger, more vivid hues
- Reduced color neutrality
- A distinctly stylized, cinematic appearance

The effect is not intended to preserve accurate color reproduction.


Recommended use cases

- Creative grading to evoke a vintage or classic film aesthetic
- Stylized color looks where realism is not required
- Situations where increased color contrast between hues is desired

Limitations

- Not a color correction or white balance tool
- May significantly alter skin tones and natural color relationships
- High values can produce unrealistic or extreme color shifts

The **Opacity** filter controls the **alpha (transparency) level** of a video or image clip. It modifies how strongly the clip contributes to the final composite.



This filter affects only **transparency**, not color values.

Parameter

Level (0–200%)

Controls the **opacity multiplier** applied to the clip's alpha channel.

- **0%** Fully transparent (invisible)
- **100%** Original opacity (no change)
- **Above 100% (up to 200%)** Amplifies the alpha channel
Partially transparent pixels become more opaque

Internally, this parameter applies a **gain to the alpha channel**, not a clamp limited to 100%.

Keyframes:

The *Level* parameter can be keyframed to animate opacity changes over time. This enables fades, dissolves, and animated transparency effects.

Why the range goes to 200%

Opacity values above 100% support **compositing workflows** where media already contains **partial transparency**.

This includes:

- Images or animations with alpha channels
- Clips affected by prior filters or blending modes
- Overlays, titles, or graphics with softened edges

Increasing opacity above 100% allows:

- Reinforcing transparency that was reduced earlier in the filter chain
- Restoring visual dominance without reordering filters
- Fine control over layered compositions

Final output opacity is still clamped at fully opaque.

PNG and animated WebP transparency behavior

Formats such as **PNG** and **animated WebP** store **RGB color data even for fully transparent pixels**.

Although these pixels are invisible when alpha is 0 (at 100% on the slider), their color values still exist in the file and can influence compositing when:

- Opacity is increased
- Blending modes are applied
- Alpha is amplified above 100%

As a result, increasing opacity can make **semi-transparent areas and their underlying color data more visible**, which explains why opacity amplification is meaningful for these formats.

Transparency behavior by format

Format	Alpha support	RGB data under full transparency
PNG	Yes	Yes
WebP / Animated WebP	Yes	Yes
GIF	1-bit transparency	No (palette-based)
Video with alpha (e.g. ProRes 4444)	Yes	Yes

Visual characteristics

- Below 100%: increasing transparency
- At 100%: unchanged appearance
- Above 100%:
 - Semi-transparent areas become more solid
 - Overlays and edges appear stronger
 - No change to color values, only alpha

Recommended use cases

- Fades and dissolve control
- Strengthening overlays, titles, or graphics
- Correcting reduced opacity after blending or compositing
- Adjusting transparency without changing clip order

Limitations


- Does not affect color, brightness, or saturation
- Fully opaque pixels remain unchanged; opacity amplification affects only pixels with partial transparency.
- Effect depends on **existing** alpha information
- Hidden RGB data may become visible when opacity is amplified

The **Outline** video filter uses the image or video's alpha channel (transparency) to draw a colored line also known as a stroke. Therefore, it works best with something that has a transparent background and hard edges (not large areas of translucency). Since it is common to put outline on text and Text: Rich does not have its outline feature, this is very useful for rich text.

WARNING: This filter uses the CPU very heavily! Therefore, if you are using a text filter that has its own outline option such as **Text: Simple** use that instead. Use it for short periods if possible. If the thing you are outlining is a still image that you will show for a long time such as a so-called watermark or superimposed logo; use this to apply the outline, export it as a new file, and then use that.

This filter was added in Shotcut version 25.07.

NOTE: Since this is based on the alpha channel of the *filter's input* it will not work as expected when used *after a text filter* on a video or image clip. That is because a filter's input is the output of a previous filter, and the output of the text filter does not create an alpha channel with transparency. Rather, you must use a text clip (**New Generator > Text**) on an upper track.




- **Color** is the obviously the color of the line to draw. Click the color block to open the color dialog to change it, or click the dropper icon to pick a color from your screen.
- **Thickness** is the width of the line. The greater the thickness the more it uses your CPU and slows preview or export. The resolution of your project's Video Mode affects how much thickness you need. More resolution generally requires more thickness but also requires more pixels to process slowing everything. Proxy and Preview Scaling are your friends here, but expect long export times if you use this much.
- **Color** and **Thickness** support Keyframes to change over time for cool effects!
- **Preset** is where you can save these parameters to use them again for other clips or projects.



Reduces the number of distinct tonal values in the image, replacing smooth gradients with discrete color or brightness steps.

The result is visible banding and simplified shading, where continuous transitions are intentionally removed.

Posterization originates from both artistic and technical constraints. Historically, it appeared in printmaking and early photographic reproduction, where limited ink densities or exposure steps forced images into discrete tone bands. In digital imaging, posterization became a deliberate effect used to simplify imagery, emphasize shapes, or create graphic, illustration-like results.



Posterize reduces tonal resolution but does not perform edge detection or alter spatial detail.

Parameters

Levels (0 - 100%)

Controls the **number of tonal steps** used to represent the image, expressed as a percentage.

- **0%**
Extremely low number of tones. Large flat areas of color or brightness dominate the image. Fine shading is almost entirely removed.
- **Low values (5% - 30%)**
Strong posterization. Gradients break into a few clearly defined bands. Details are simplified into graphic regions.
- **Medium values (30% - 70%)**
Moderate posterization. Banding is visible but some tonal variation remains.
- **100%**
Maximum available levels for this filter.

Banding is still present and the image does not return to a fully continuous tone representation.

Important behavior

- 100% does **not** mean “no effect.”
- The filter always quantizes tones; higher values only increase the number of steps, not eliminate posterization entirely.

Note:

Levels controls how many steps exist, not how strong the effect feels. Even at maximum, tones remain discretized.

Keyframes

The **Levels** parameter can be keyframed.

This allows:

- Gradual transitions from realistic footage to a stylized or graphic look
- Temporal emphasis during edits or transitions
- Animated abstraction effects

Visual characteristics

- Visible tonal banding in gradients
- Flat color or brightness regions
- Emphasis on edges and large shapes
- Reduced fine shading detail

The effect is purely **spatial** and applies uniformly across the frame.

Recommended use cases

- Creating graphic or illustrative looks
- Stylizing footage toward a cartoon-like or poster-print appearance
- Preparing footage for further processing (e.g. edge detection, thresholding)
- Emphasizing form and contrast over realism

Example combinations

- **Posterize + Threshold**

Produces a stark, graphic look by first reducing tonal steps and then collapsing them into black and white. Useful for silhouette-style or high-contrast illustration effects.

- **Posterize + Sharpen (subtle)**

Emphasizes boundaries between tonal bands without attempting true edge

detection. This can make shapes read more clearly in posterized footage.

- **Posterize + Color Grading / Saturation**

Reinforces flat color regions, creating pop-art or screen-print-like visuals.

- **Posterize + Gaussian Blur (very low radius)**

Softens harsh band transitions between posterized tone regions while preserving the reduced tonal palette.

- **Posterize + Mosaic (very subtle)**


Reinforces abstraction by simplifying both tone and spatial detail.

Limitations

- No control over which tones are preserved or discarded
- Always introduces banding, even at maximum Levels
- Can exaggerate compression artifacts in smooth gradients
- Not suitable when natural shading must be preserved



HQDN3D (High Quality DeNoise 3D) reduces noise by smoothing pixel variations across **space and time**. The “3D” refers to processing in two spatial dimensions (X, Y) and one temporal dimension (frames).



This filter reduces noise by **averaging similar values**, not by reducing precision or analyzing edges.

Parameters

Spatial (0.0–100.0)

Controls the amount of **spatial noise reduction** applied **within each frame**.

This parameter uses a **normalized, unitless scale** representing relative denoising strength.

- Lower values
Minimal smoothing; preserves fine detail and edges
- Higher values
Stronger smoothing; increased noise reduction with potential loss of texture

Spatial denoising primarily targets:

- Grain
- Sensor noise
- Compression artifacts within a single frame

Temporal (0.0–100.0)

Controls the amount of **temporal noise reduction** applied **across adjacent frames**.

This parameter also uses a **normalized, unitless scale**.

- Lower values
Little or no frame averaging
- Higher values
Stronger noise reduction over time

Temporal denoising primarily targets:

- Flickering noise
- Temporal grain
- Low-light shimmer

High values may introduce motion artifacts such as smearing or ghosting.

Keyframes

Both **Spatial** and **Temporal** parameters can be **keyframed**, allowing noise reduction strength to vary over time.

This enables adaptive workflows such as:

- Stronger denoising in dark or static sections
- Reduced denoising during fast motion
- Gradual transitions between noise profiles

Parameter interaction

The two controls operate on **different dimensions**:

- **Spatial** smooths variations within a single frame
- **Temporal** smooths variations between frames

Moderate values for both parameters typically produce better results than extreme values on either control alone.

Visual characteristics

Typical effects include:

- Reduced grain and flicker
- Smoother shadows and flat areas
- Loss of fine detail at high Spatial values
- Motion artifacts at high Temporal values

Recommended use cases

- Low-light or high-ISO footage
- Compressed or archival video
- Static or slow-moving scenes
- Footage where spatial-only denoising is insufficient

Limitations

- Not edge-aware
- May soften fine detail
- Temporal denoising can cause ghosting on motion
- High values are unsuitable for fast action or handheld footage



This filter is available since version 20.06.

About wavelet denoising

Wavelet denoisers are excellent when decent-bitrate footage comes from a decent-quality camera, and the only real problem with the footage is high ISO noise or high thermal noise or photon shot noise. Wavelets are great at targeting inconsistency at the pixel level, as opposed to algorithms that average patches or regions, and end up smearing pixels in the process. A nice side effect of this precision is that gradients (especially skin tones and out-of-focus areas) are rendered extremely smoothly, which is a benefit that cannot be overstated. My primary use for this filter is to clean up low-light video. It usually increases the level of realism in the process.

That said, wavelets are not the best tool for making VHS tapes or 19th-generation MPEG-2 videos look great. Wavelets will usually perceive blocky borders as intentional detail and preserve them rather than eliminate them, hence the need for a decent-quality start point. Use NLMeans or HQDN3D for restoration purposes.

Demonstration using the attached images

Below are attached two high-ISO images to give everyone a common reference point. *Film.jpg* was shot at ISO 8000, and *Dashboard.jpg* is a frame grab from an ISO 6400 video, both from Micro Four Thirds cameras.

Using a 2160p UHD video mode, add the following filters to *Film.jpg* and compare them one at a time to the original:

- “Reduce Noise: Wavelet” at Heavy preset
- “Reduce Noise: HQDN3D” at 70% spatial
- “Reduce Noise: Smart Blur” ... any

Then using the same UHD video mode, compare *Dashboard.jpg* with these filters:

- “Reduce Noise: Wavelet” at Medium preset
- “Reduce Noise: HQDN3D” at 25% spatial
- “Reduce Noise: Smart Blur” ... any

Naturally, this is best observed with a full-screen external monitor.

Observations to help choose between HQDN3D and Wavelet

- On the Film image, Wavelet converted the noisy background into liquid gold. Very smooth gradients at the left and right edges. When HQDN3D tried the same thing at 70% strength, it got very close, but there were still some banding

and macroblocking artefacts inside the gradients. This becomes more evident during video playback when the bands wiggle like worms.

- On the Dashboard image, HQDN3D had to be dialed back to 25% because more strength would cause the skin tones to look like a plastic doll. Being scaled back this far (compared to 70% in the Film image) left many more bands, macroblocks, and color patches on the dashboard compared to Wavelet. Skin tones frequently prevent HQDN3D from being used at a higher strength.
- On both images, HQDN3D shifts the image down and right by a noticeable amount (several pixels) compared to the original, which makes me very uneasy. When applied to video, the shift is not constant either, so the video shakes. I don't know if this is the usual temporal ghosting issue or if this is an implementation issue. Maybe [@Paul_B_Mahol](#) has some ideas. It also happens with command-line FFmpeg. Meanwhile, Wavelet "just works" every single time.
- The main downside of Wavelet is that it takes 2.75 times longer to process than HQDN3D (on my hardware at least). For longer clips, I sometimes process them separately, then bring them back into the main Shotcut project as DNxHR intermediates.
- I could not find any combination of settings that made Smart Blur even remotely match the quality of HQDN3D and Wavelet. Smart Blur produced blocky gradients and blurry edges by comparison. As for export speed, it was only 30% faster than HQDN3D on a 16-core server, and identical in speed on a 4-core laptop. I struggle to find a use case for Smart Blur when much higher-quality options are available.

Why bother with another denoiser filter when Shotcut already has two:

HQDN3D and Wavelet use different techniques to serve different needs for different audiences. For video with significant motion, I've found that HQDN3D and Wavelet produce results that are visually similar. My eye can't track the finer details during motion, so I'll sometimes take the speed benefits of HQDN3D. But for a slow dramatic shot of a fancy restaurant interior in low light where the audience wants to see every detail of fine gold without artefacts... Wavelet scores a clear win. So this duo of filters satisfies the "Fast" crowd (HQDN3D) and the "Quality" crowd (Wavelet). Since my audience wants "Quality" 95% of the time, I made this interface to get it. Whether it's the right choice for you depends on your use case.

Happy low-light filming!






This filter reduces noise in video by **quantizing pixel values**, i.e. rounding them to discrete steps. Random, low-amplitude variations are removed because they fall below the quantization step and cannot be represented.

This method:

- Operates **spatially** (per frame)
- Does **not** analyze motion or neighboring frames
- Does **not** blur edges or average pixels over time

Noise is reduced by **removing precision**, not by detection or smoothing.



Parameters

Quantization (0–64)

Controls the **quantization step size**, i.e. how coarse the rounding of pixel values is.

Conceptually, pixel values are rounded to multiples derived from this parameter.

- Lower values apply finer rounding
 - Removes only very small variations
 - Preserves gradients and texture
- Higher values apply coarser rounding
 - Removes more noise and fine detail
 - Reduces tonal precision
 - May introduce banding or posterization at high settings

This parameter defines **which level of detail is discarded**.

Strength (0–100%)

Controls how much of the quantized result is **blended with the original image**.

- 0%: original image (no effect)
- 100%: fully quantized image

Intermediate values interpolate between the two, allowing noise reduction while limiting visible artifacts.

Mathematically, the output is a linear interpolation between the original and quantized images.

Parameter interaction

The two controls serve different purposes:

- **Quantization** determines the rounding threshold and loss of precision
- **Strength** controls how strongly that rounding is applied to the final image

Using a higher Quantization value with a lower Strength often produces cleaner results than using low Quantization at high Strength.

Visual characteristics

Typical effects include:

- Reduction of fine grain and compression noise
- Better edge preservation compared to blur-based noise reduction
- Possible banding in smooth gradients if overused

Recommended use cases

- Archival, CCTV, or low-quality sources
- Footage affected by compression artifacts (like JPEG images)
- Situations where temporal noise reduction produces ghosting or smearing

Limitations


- Not motion-adaptive
- Removes detail indiscriminately below the quantization threshold
- High settings may degrade smooth gradients or skin tones

Summary

Quantization noise reduction removes noise by **restricting the number of representable pixel values**, with Quantization defining the precision limit and Strength controlling how much of the quantized image replaces the original.



The **Smart Blur** filter reduces noise by applying a **selective blur** that attempts to smooth small variations while preserving edges. Unlike simple blur filters, it limits smoothing based on pixel differences to avoid blurring across strong edges.



This filter operates **spatially within a single frame** and does not use temporal information.

Parameters

Blur radius (0.1–5.0)

Controls the **size of the neighborhood** used for blurring.

- Lower values
 - Affect only immediate neighboring pixels
 - Preserve fine detail
- Higher values
 - Use a wider area for averaging
 - Increase smoothing and risk of detail loss

This parameter defines **how far the blur can spread**.

Blur strength (0.0–1.0)

Controls the **intensity of the blur** applied within the selected radius.

- 0.0
 - No visible effect
- Higher values
 - Stronger smoothing within the blur radius

This parameter determines **how strongly pixels are blended** once they are considered eligible for blurring.

Threshold (0–30)

Controls the **edge-preservation threshold**.

- Lower values
Blur is applied only to very similar pixels
Strong edge preservation
- Higher values
Blur is allowed across larger differences
Increased smoothing with reduced edge protection

This parameter defines **what is considered noise versus detail**.

Parameter interaction

The three controls work together:

- **Blur radius** defines the spatial extent
- **Blur strength** defines the blur intensity
- **Threshold** defines which differences are preserved as edges

Effective noise reduction typically requires balancing all three rather than increasing one parameter alone.

Visual characteristics

Typical effects include:

- Reduction of fine grain and small artifacts
- Preservation of strong edges at moderate thresholds
- Gradual loss of texture and edge definition at high settings

Recommended use cases

- Light spatial noise or grain
- Footage where temporal denoising is undesirable
- Static or lightly textured scenes
- Situations requiring edge-aware smoothing


Limitations

- Operates only within individual frames
- Cannot reduce temporal noise or flicker
- Parameters **cannot be keyframed**

- High values may blur fine detail and edges

The **Reflect** filter mirrors the image by **splitting the frame along a selected edge axis** and flipping one half to create a reflection.

When used alone, the reflection axis is fixed at the **center of the image**, and one half of the frame is mirrored onto the other.



This filter operates **spatially** and does not modify color or transparency.

Parameter

Direction

Selects which side of the image is used as the **source for the mirror**, determining which half of the frame is duplicated onto the other.

Available options:

- **Right**
The **right half** of the image is reflected onto the **left half**
- **Left**
The **left half** of the image is reflected onto the **right half**
- **Bottom**
The **bottom half** of the image is reflected onto the **top half**
- **Top**
The **top half** of the image is reflected onto the **bottom half**

The reflection is generated by flipping the image content across the chosen edge.

Visual characteristics

- Produces a mirrored version of the image
- Reflection is symmetrical relative to the selected edge

- No scaling, fading, or blending is applied by this filter alone
- The reflected area is a direct flip of the source content

Recommended use cases

- Creating mirror or symmetry effects
- Filling empty frame areas with reflected content
- Stylized layouts and graphic compositions
- Background extension effects

Workflow note: combining with Size, Position & Rotate

For greater control over **where the reflection appears**, use the **Size, Position & Rotate** filter **before** the Reflect filter in the filter list.

When **Size, Position & Rotate** is placed above Reflect:

- The image can be resized or repositioned first
- The reflection is applied **after** those transformations
- This allows precise control over:
 - Reflection placement
 - Visible reflected area
 - Composition balance

Filter order matters: Reflect always mirrors the result of all filters applied **above it**.


Limitations

- No adjustable reflection strength or opacity
- No gradient or fade control
- Reflection direction is discrete (fixed choices)
- Effect depends on prior transformations in the filter chain

Reframe or reframing is a way to crop a video or editing project in order to export it to a different aspect ratio. For example, you made a video project in 16:9 aspect ratio landscape/horizontal orientation but now want to make a 9:16 portrait/vertical orientation video for social media. Other forms of cropping are used primarily *within* a composition to remove bad edges (Crop: Source), zoom in to fill the frame (Size, Position & Rotate), add colored borders (Crop: Rectangle), or clip an overlay (Crop: Rectangle).


Reframe was added in version 24.09, but it is not available if **Settings > Processing Mode** is **Linear 10-bit GPU/CPU**.

One can only add the **Reframe** video filter to **Timeline > Output**. However, for convenience there is also a **Reframe** button in **Export > Video**



(**Resample** is typically used to change the resolution without changing aspect ratio. But one can use it and change aspect ratio in order to *intentionally* add black bars. For example, you compose a very wide film aspect ratio of 2.4:1 but want to deliver it as 16:9—letterboxing.)


Clicking **Reframe** adds the filter to **Output** or switches the user interface to it:



You can drag this rectangle control around and resize it to even dimensions only (for codec and pixel format compatibility), but you cannot push any part of the

rectangle outside of the frame.

It supports keyframes and motion tracking. However, if you change the size, only the size at the beginning is used. Here, I tracked the white horse in the first clip only and then re-centered it.




There are presets for 9:16 (defaults), 16:9, 4:3, and Square.

You can save the Reframe filter on your project and simply disable it when you want to export without it. You can have multiple Reframe filters and the first enabled one is used on export.

Separates the red, green, and blue color channels by offsetting them horizontally and/or vertically.

Neutral areas such as white, gray, and highlights split into distinct red, green, and blue bands, making the effect immediately visible.

The filter is a **spatial color offset**, not a depth or stereoscopic effect, even though it can resemble a glasses-free “3D” look.



With both parameters set to **50.0%**, the filter is neutral and produces no visible change.

Parameters

Horizontal (0.0 - 100.0%)

Controls horizontal separation of the RGB channels.

- **50.0%**
Neutral position. All color channels are aligned.
- **Below 50.0%**
Red, green, and blue channels separate horizontally in one direction.
White or light areas split into **three vertical color bands**.
- **Above 50.0%**
Channel separation occurs in the opposite horizontal direction.

Behavior note:

- The farther the slider moves away from 50.0%, the greater the distance between the three color channels.
- Image geometry does not move; only color channels are displaced.

Vertical (0.0 - 100.0%)

Controls vertical separation of the RGB channels.

- **50.0%**
Neutral position. All color channels are aligned.
- **Below 50.0%**
Red, green, and blue channels separate vertically.
White or light areas split into **three horizontal color bands**.
- **Above 50.0%**
Channel separation occurs in the opposite vertical direction.

Keyframes

Both parameters can be keyframed.

This enables:

- Gradual appearance or disappearance of channel separation
- Animated color drifting
- Glitch-style pulses or oscillations

Parameter interaction

- Horizontal and Vertical offsets are independent and cumulative.
- Using both produces diagonal RGB separation.
- Increasing distance between channels increases color visibility without changing luminance placement.

Visual characteristics

- White and light areas split clearly into red, green, and blue components
- Colored fringes appear along high-contrast edges
- Apparent loss of sharpness due to channel misalignment
- Strong visual effect even at moderate values

Recommended use cases


- Stylized chromatic aberration
- Glitch or digital distortion effects
- Text effects
- Visual emphasis on movement or impact
- Abstract or experimental visuals
- Transitions and visual accents

Limitations

- Does not create real depth or stereoscopic 3D
- Can severely reduce readability of text and fine detail

- No control over individual channel distance or order
- Effect strength is resolution-dependent

The **Saturation** filter adjusts the **overall intensity of colors** in the image. It uniformly scales color saturation across all hues without favoring or protecting specific color ranges.



This filter affects color intensity only and does not change brightness or contrast.

For more controlled color enhancement, the **Vibrance** filter may produce more natural results.

Parameter

Level (0%–300%)

Controls the **global saturation level** of the image.

- **0%**
All color information is removed, producing a grayscale image
- **100% (default)**
Original saturation (no change)
- **Above 100%**
Colors become increasingly vivid and intense

This parameter applies the same scaling factor to all colors.

Keyframes

The **Level** parameter can be **keyframed**, allowing saturation to change smoothly over time.

This enables:

- Gradual color fade-in or fade-out
- Animated color emphasis
- Stylized transitions between monochrome and color

Visual characteristics

Typical effects include:

- Uniform increase or decrease in color intensity
- Progressive loss of color detail below 100%
- Strong, sometimes exaggerated colors above 100%
- No protection for already saturated colors

Recommended use cases


- Global color intensity adjustments
- Stylized color effects
- Monochrome-to-color transitions
- Situations where equal treatment of all colors is desired

Limitations

- Affects all colors equally
- High values may cause color clipping or unnatural results
- Not selective by hue or saturation range



The **Scan Lines** filter simulates **horizontal scan lines** similar to those seen on CRT displays or interlaced video systems. It creates evenly spaced horizontal gaps across the image to produce a retro or display-based visual effect.



This filter does not modify color or motion.

Behavior and compositing

The Scan Lines filter does **not draw opaque black lines**. Instead, it **removes or makes transparent portions of the image** in horizontal bands.

As a result:

- When applied to a clip on the **bottom video track**, the transparent bands reveal the project background, which typically appears as **black lines**.
- When applied to a clip on a track **above another video**, the transparent bands reveal the **clip underneath**, making the scan lines appear as horizontal windows to the lower track.

In other words, the filter behaves like a **striped transparency mask**, not a simple overlay.

Visual characteristics

Typical effects include:

- Regular horizontal line pattern
- Reduced vertical resolution appearance
- Retro or display-like visual texture
- Underlying layers visible through the lines when composited

The appearance depends on what exists beneath the filtered clip.

Recommended use cases


- CRT or retro display simulation
- Stylized compositing effects
- Revealing underlying footage in a patterned way
- Combining multiple video layers with structured transparency

Limitations

- No adjustable parameters
- Fixed line spacing and thickness
- Appearance depends on track order and underlying content
- Uniform effect across the frame

The **Sepia Tone** filter applies a **two-axis color tint** inspired by traditional sepia processing.

Historically, *sepia* refers to a brownish tone produced by chemical treatment of black-and-white photographs and film, originally used to improve image longevity and later adopted for its distinctive aesthetic.



In Shotcut, this filter generalizes that concept: it can produce classic sepia looks **as well as other monochrome color tints**, depending on parameter settings.

Parameters

Yellow-Blue (0–255)

Controls the color balance along the **yellow–blue axis**.

- Lower values
Shift the image toward **blue**
- Higher values
Shift the image toward **yellow**

Cyan-Red (0–255)

Controls the color balance along the **cyan–red axis**.

- Lower values
Shift the image toward **cyan**
- Higher values
Shift the image toward **red**

Neutral point and monochrome behavior

When **both** sliders are set near **127-128**, the filter applies **no color bias**, resulting in a **neutral black-and-white image**.

Values above or below this midpoint introduce color tinting by biasing the grayscale image along the two color axes.

Color behavior and flexibility

Although named *Sepia Tone*, the filter is not limited to sepia coloration:

- The **default values** (Yellow-Blue ≈ 75 , Cyan-Red ≈ 150) produce a classic sepia-like brown tone
- Other combinations can produce:
 - Cool cyan or blue monochrome looks
 - Warm reddish or golden tones
 - Stylized monochrome color effects

The filter always starts from a **monochrome base** and then applies color bias.

Visual characteristics

Typical effects include:

- Conversion to black and white at neutral settings
- Warm or cool monochrome tinting
- Uniform color application across the frame
- No change to image geometry or motion

Recommended use cases

- Vintage or historical looks
- Period reenactments
- Stylized monochrome grading
- Creating tinted black-and-white video

Historical tone references (approximate)

The following historical photographic toning processes can be visually approximated using this filter.

1. **Sepia (Sulfur toning)**

Produces **warm brown tones** from silver-based black-and-white images.
Closely matches the filter's default behavior and naming.

2. **Gold Chloride Toning (Daguerreotype gilding)**

Used on daguerreotypes to improve contrast and permanence.
Results in **cool-to-neutral grayscale** with slight warm highlights.

3. **Selenium Toning**

Deepens blacks and shifts tones toward **cool purples or reddish-browns**.

Can be approximated with reduced Yellow-Blue and increased Cyan-Red bias.

4. **Platinum / Palladium Printing**

Known for **neutral to slightly warm gray tones** with smooth tonal transitions.

Approximated by keeping both sliders near midpoint with **subtle warm** bias.


5. **Cyanotype**

Produces **distinctive blue** monochrome images.

Approximated by lowering Yellow-Blue and Cyan-Red toward blue/cyan values.

Limitations

- Parameters **cannot be keyframed**
- Uniform effect across the frame
- Not intended for precise color grading
- No control over contrast or luminance



The [frei0r plugin](#) description says, “Unsharp masking (port from Mplayer).”

The [mplayer man](#) page says it does gaussian blur or a sharpen using an unsharp mask routine.


The underlying [plugin code](#) range is $[-1.5, 3.5]$ where 0 is the switching point between blur and sharpen. Thus, 50% in Shotcut converts to 1.0 in the plugin code, and the switch point where there is no change is 30% in Shotcut.

See also https://en.wikipedia.org/wiki/Unsharp_masking

The **Size** parameter controls the amount of blur just like nearly all blur filters have as a parameter. Sharpening still creates a blur, but it uses this as a mask to sort of subtract from the original. This also explains why it can also be used to blur and why many tools call this function “unsharp” or “unsharp mask.” However, I do not like to call it that in Shotcut because it is confusing and deceptive for non-advanced users who are generally looking to increase sharpness (while also not limiting it to that).




Location: Filters > Video > Size, Position & Rotate



Default Shotcut Installation: This filter will be located in the Favorite tab.

This Favorite can be removed/added by clicking on the * symbol.



This filter is used to manipulate the size and placement of a video/image in your video. This filter can also use Keyframes.

Size Modes:

- Fit - Will not enlarge image/video in the Rectangle Control, yet allows image to be scaled down.
- Fill - Fits any image/video vertically to the Video Mode. Constrains the rectangle control to the aspect ratio of the original source
- Distort - Allows resizing to the Rectangle Control, overriding Aspect Ratio of original source

Horizontal fit:

- Left - Aligns to the left side of Rectangle Control
- Center - Aligns to the center of Rectangle Control
- Right - Aligns to the right of Rectangle Control


Vertical fit:

- Top - Aligns to the top of Rectangle Control
- Middle - Aligns to the middle of Rectangle Control
- Bottom - Aligns to the bottom of Rectangle Control

The Rectangle Control is the box that surrounds the image/video allowing placement and size manipulation. In this rectangle you'll see the Position Handle (white/gray dot) that you can move the image/video to a desired location.

The filter can be deselected by unchecking the box, yet all of the changes remain there.


To remove the filter use the - sign.



Position Handles

The position handles in the corners allow the rectangle to be scaled proportionally.

If the Size mode is “Fit” or “Distort”, additional position handles are displayed on the top, bottom, left, and right. Those position handles can be used to change the shape of the rectangle. The top, bottom, left, and right handles are not available when the Size mode is “Fill” because the aspect ratio of the rectangle is enforced.




Snapping

When rotating, the rotation will snap to 0, 90, 180, and 270 degrees. Additionally, if a grid is enabled on the preview window, the position handles will snap to the grid.

To disable snapping, hold the Alt key while rotating or moving the position handles.


A sketch effect is a digital transformation that **makes a image look like a hand drawing**. It can mimic different artistic styles, such as pencil, ink, charcoal, or crayon.





Removes or obscures a **rectangular area** of the image by replacing it with pixels blended from the surrounding region.

The filter samples pixels around the perimeter of the defined area and uses them to fill the interior, producing a localized concealment effect.



Spot Remover is a **spatial effect** applied per frame.

Parameters

Position

Defines the location of the rectangular area to be removed.

- Can be adjusted numerically or interactively in the player preview.
- Moving the position changes which part of the image is replaced.

Size

Defines the width and height of the rectangular removal area.

- Can be adjusted numerically or by dragging handles in the player preview.
- Larger sizes obscure more content but increase the risk of visible artifacts.

Load Keyframes from [Motion Tracker](#) (button)

Imports position and size keyframes from a previously applied **Motion Tracker** filter.

Behavior notes:

- Motion Tracker must be applied and analyzed first.
- This button transfers the tracked motion data into the Spot Remover's parameters.
- After loading, the Spot Remover follows the tracked object automatically.

Keyframes

- **Position** and **Size** are keyframeable.
- Keyframes can be created manually or imported from Motion Tracker.
- Imported keyframes can be edited afterward if needed.

Parameter interaction

- **Position** determines where pixels are replaced.
- **Size** determines how much of the image is replaced.
- Larger areas require more surrounding context to blend convincingly.
- Motion Tracker integration automates Position and Size changes over time.

Visual characteristics

- Rectangular concealed region
- Filled using blended pixels from the surrounding edge
- No blur or mosaic pattern by default
- Best results when the surrounding texture is relatively uniform
- Artifacts may appear if the removed area overlaps strong edges or complex detail

Recommended use cases

- Obscuring faces, logos, or license plates
- Removing small unwanted objects
- Fast concealment of moving subjects
- Cleanup tasks where precision masking is unnecessary

Limitations

- Removal shape is always rectangular
- Not suitable for complex shapes or large areas
- Blending quality depends heavily on surrounding content
- Not a true content-aware or object-removal tool

Editor's note:

In practice, this filter combined with **Motion Tracker** is one of the **fastest and simplest-easiest** ways to hide a moving face:


1. Apply **Motion Tracker** and analyze the clip
2. Add **Spot Remover**
3. Click **Load Keyframes from Motion Tracker**

No manual animation is required for basic tracking tasks.


Stabilization reduces or eliminates unwanted camera shake, creating a smoother and more professional-looking video. This is particularly important for footage shot handheld, in motion, or in windy conditions. Without stabilization, shaky videos can be distracting, disorienting, and even cause dizziness for viewers.

The Shotcut filter uses the open source library [GitHub - georgmartius/vid.stab: Video stabilization library](https://github.com/georgmartius/vid.stab)

Here is how its UI looks initially



You must first analyze before the filter affects the video. Clicking **Analyze** creates an analysis job in **Jobs**. You will be prompted to save a `.stab` file, which stores the analysis data. You are not prompted to save a file when you have made a project folder because it automatically generates the file in the project folder. When the job completes successfully it looks something like this:



Notice it says “Analysis complete” next to the button and reports the file name.

Shotcut version 25.05 added a menu item **File > Rerun Filter Analysis** that creates a new analysis job for every **Stabilize** filter in the current project. This is useful if you did a lot of trimming or splitting clips after having done stabilization or if you change the Video Mode resolution or aspect ratio. Note that this also redoes analysis for **Normalize: Two Pass** audio filters.

Parameters

@jupiter wrote in [Stabilization Settings, Descriptions, and Tendencies](#)

After rendering a video project countless times with low & high settings of each scale, then closely studying the results, this is what I believe each setting does. The higher the setting, the more it does what is described.

- **Shakiness** seems to average out the frame movements instead of each frame migrating too much compared to the previous/next frame. Max recommended.
- **Accuracy** seems to prefer keeping the main point-of-reference figure stabilized instead of the moving background (this was for my case of a moving selfie camera; not sure what happens with a stationary camera and moving objects; in a video of walking around filming a city, this setting didn't make any difference). I found Accuracy to be more effective at stabilizing the video than Shakiness, tho there seemed no downside to maximizing both.
- **Zoom is the relative size of the output frame compared to the input frame.** If the whole input frame matters, I wouldn't set it above 0%. And if you do want to zoom in, I'd rather do it with Smoothing.
- **Smoothing is how much the frames are allowed to move around from center to stabilize the image** (this can be seen when Zoom is negative). It also zooms in, which I guess is to hide the shifting frame borders. Without allowing some Smoothing, Stabilize isn't able to do its job. I suggest between 10 to 30, depending how tolerant you are if the moving original border becomes visible in your output. If you prefer to not see any corrections along the edge and want to keep almost the entire outer frame, then set Zoom=0 and Smoothing to less than 5, but don't expect much stabilization.

By setting the Zoom out the maximum of -50% and the smoothness around 30, I was able to get a good perspective of how hard its working and what it's doing. I suggest max shakiness, max accuracy, zoom around 0%, and adjust your smoothness for how much you want the video to be stabilized.

To compare the scale of Shakiness vs Accuracy, I rendered the same video 4 times with low and high settings for each (low-low, high-low, low-high, high-high) then made a single video with each version in its own video track and displaying them in the 4 corners using the Size and Position filter (like 4 Picture-in-Pictures). Then Render/Export the whole project to have them play smoothly side-by-side and analyze the movements frame-by-frame. I also took the same approach with Zoom and Smoothing to determine their effects and preferred settings.

In my very dynamic sports video with a GoPro on a selfie stick, the Stabilize feature helped tremendously, tho the algorithm still needs improvement. Big stabilization movements seem to correct in sets of 2 frames instead of every frame (or the average of many frames) across the sample, resulting in jitteriness. Lossy vs Lossless didn't seem to correct this, nor did any combination of Stabilization settings.

Stabilize/Analyzing filter tendencies:


- When exporting the video, the stabilization filter won't apply to the output if the frame rate or video size is changed.
- When Stabilizing (or any other filter) make sure you've properly selected what you want it applied to. Videos in the Player, Playlist, and Timeline are 3 different instances, and are not connected. To apply to the other instances, you must update or add the video to them. When Exporting, you also select From which source at the top of the tile.
- Player and Playlist videos only analyze the portion selected. If its in-point/out-points are ever expanded (in the Player or Timeline), the new sections won't be stabilized.
- Timeline clips analyze the entire source video, not just the portion that's selected. If you only need a small clip of a long source video, it's much quicker to analyze a portion in the Player/Playlist, then add it and further trim it down in the Timeline.
- The Stabilization filter can only be applied to individual clips, not the entire track or project. To affect more than just a single clip, either apply stabilization before you chop a source video into many pieces, or render your tracks or final project into an output file with minimum compression, then start a new project with that output file as your source and apply your stabilization and other global processing as needed.
- If after stabilizing, you see little artifacts along the edges, this is a byproduct of it being stabilized. If you set your Zoom to a negative value with a little bit of Smoothing, you can see what's happening. Read about Zooming and Smoothing above for a solution.


The Subtitles Burn-In Filter was added in Shotcut 24.08

The Subtitle Burn-In Filter can overlay subtitles that have been authored in the **Subtitles Panel**. The filter can only be added to the **Timeline Output**.

After you add the Subtitle Burn-In Filter to the Timeline Output, select the subtitle track name from the Subtitle Panel that you want to overlay.

The subtitle overlay text can be configured similar to the Simple:Text filter - including the preview window controls for size and position.





Introduction


The Rich Text filter lets you format text (more flexible than Text: Simple) and automatically wraps lines within the rectangle (unlike Text: Simple). To get started after adding the filter simply enter the text in the rectangle that appears on top of the video preview area. The text and rectangle are only editable when the play head is over the selected clip when using the Timeline. Then, format the text using the floating toolbar along the top of the rectangle. First, you must select some text as shown in the screenshot above.

Toolbar

Each icon on the toolbar has a tool tip when you hover your mouse over it for one second:




Opens a menu with less commonly used functions explained below.




Sometimes, you want to get the toolbar out of the way of your work.

So, you can collapse it:




And click again to restore it.

You can right-click the selected text to open a menu of things to do with editing the text:



These are explained below.


Toolbar Menu



- Use **File > Open...** to open a file you created with **Save As...**. The text you enter is saved inside of the project automatically, but maybe you want to save it to use in another clip or project.
- Use **File > Save As...** to save it as rich text in [a subset of the older HTML 4 and CSS standards](#) or as plain text in a `.txt` file. Most people should not

attempt to write the HTML to be included by this filter. This is not a full browser engine.

- **Edit**
- **Undo**
- **Redo**
- **Cut** - Copy the selected text to the system clipboard and remove it.
- **Copy** - Copy the selected text to the clipboard.
- **Paste** - Paste text (with formatting if available) from the clipboard.
- **Paste Text Only** is a special version of paste that does not include formatting information. If you have copied text from a word processor or web page, sometimes the formatting does not paste nicely into Shotcut, or you simply do not care about the original formatting.
- **Select All**
- **Insert Table** You can even insert a table with or without lines. This can be useful for making columns of text, but the formatting of the table such as line style, padding, and cell background color is not available.



Parameters

- **Preset** contains a number of size and animation presents. You can also save your own as well as remove one you saved.



- **Position**
- **Background Size** - You can change the size of the text rectangle and move it using either these numeric fields or by dragging the handles of the rectangle on the video preview.
- **Text size** is not editable. These are handy for your reference, and you can select the values to copy them to the clipboard.
- **Background color** - You can change the background color of the rectangle as well as make it translucent (change Alpha or Opacity inside the color dialog).
- **Overflow** controls what happens when there is more text than fits into the rectangle.
 - **Automatic**
 - **Visible** - The text continues on below the rectangle.
 - **Hidden** - The text outside of the rectangle is clipped.
- **Load Keyframes from Motion Tracker** applies the motion tracking data to the rectangle position and optionally size as well.



Introduction

The Text: Simple filter in Shotcut lets you overlay dynamic text onto your video. This text can be customized using keywords that are replaced with specific information about the clip or system.

Keywords

- Keywords are enclosed in hash symbols (#). For example, `#timecode#` represents the frame's timecode.
- Supported keywords:
 - `#timecode#` or `#smpte_df#`: SMPTE drop-frame timecode of the frame.
 - `#smpte_ndf#`: SMPTE non-drop-frame timecode of the frame.
 - `#frame#`: Frame number of the frame.
 - `#filedate#`: Modification date of the file (GMT).
 - `#localfiledate#`: Modification date of the file (adjusted for local time zone).
 - `#localtime#`: Current system date and time.
 - `#resource#`: Full file name with path of the source file.
 - `#filename#`: File name without the path.
 - `#basename#`: File name without the extension.
 - `#createdate#`: Creation date of the source file. Set in the playlist clip, otherwise defaults to file modification date.
 - `#meta.<field>#`: Accesses metadata properties for file based clips.

Formatting Time-based Keywords

- Time-based keywords (`#createdate#`, `#filedate#`, `#localfiledate#`, and `#localtime#`) can be formatted using the `strftime` function.
- Add the desired `strftime` format string after the keyword, separated by a non- # delimiter.
 - `#localtime %I:%M:%S %p#` displays the time in 12-hour format with AM/PM.
 - Refer to the [documentation](#) for format specifiers.

Using Metadata Properties

- Use `#meta.<field>#` to include properties from the “Properties > Metadata” section.
- Example
 - `#meta.media.0.codec.frame_rate# fps`
 - displays the frame rate followed by “fps”

The best way to see what metadata properties are available for a clip is to open the project file (.mlt) in a text editor and look for the section that applies to that clip. Here is an example section:


```
<chain id="chain0" out="00:10:34.520">
  <property name="length">00:10:34.560</property>
  <property name="eof">pause</property>
  <property name="resource">bbb_sunflower_1080p_60fps_normal.mp4</property>
  <property name="mlt_service">avformat-novalidate</property>
  <property name="meta.media.nb_streams">3</property>
  <property name="meta.media.0.stream.type">video</property>
  <property name="meta.media.0.stream.frame_rate">60</property>
  <property name="meta.media.0.stream.sample_aspect_ratio">0</property>
  <property name="meta.media.0.codec.width">1920</property>
  <property name="meta.media.0.codec.height">1080</property>
  <property name="meta.media.0.codec.rotate">0</property>
  <property name="meta.media.0.codec.pix_fmt">yuv420p</property>
  <property name="meta.media.0.codec.sample_aspect_ratio">1</property>
  <property name="meta.media.0.codec.colorspace">709</property>
  <property name="meta.media.0.codec.name">h264</property>
  <property name="meta.media.0.codec.long_name">H.264 / AVC / MPEG-4 AVC / MPEG-4
  <property name="meta.media.0.codec.bit_rate">4001453</property>
  <property name="meta.attr.0.stream.creation_time.markup">2013-12-16T17:59:32.00
  <property name="meta.attr.0.stream.handler_name.markup">GPAC ISO Video Handler<
  <property name="meta.attr.0.stream.vendor_id.markup">[0][0][0][0]</property>
  <property name="meta.media.1.stream.type">audio</property>
  <property name="meta.media.1.codec.sample_fmt">fltp</property>
  <property name="meta.media.1.codec.sample_rate">48000</property>
  <property name="meta.media.1.codec.channels">2</property>
  <property name="meta.media.1.codec.name">mp3float</property>
  <property name="meta.media.1.codec.long_name">MP3 (MPEG audio layer 3)</propert
  <property name="meta.media.1.codec.bit_rate">160000</property>
```

This clip has a property called “meta.attr.composer.markup”. So, if we put “#meta.attr.composer.markup#” in the text filter, it displays “Sacha Goedegebure”.

Additional Notes

Escape a literal # character with \#.

This filter animates the text as if it is being written by a typewriter—one letter followed by two and so on. It can also do word-by-word and line-by-line.



This filter first appeared in version 25.10.31.


All of the parameters on the bottom half are the same as [Text: Simple Video Filter](#). So, this topic explains mainly the typewriter parameters on the top half:

- **Text** is the message you want it to type.
- **Typewriter rate** is the number of frames between each character, word, or line appearance. Lower is faster.
- **Rate variance** is the variation in timing (in frames).
- **Animation** is whether to show by character (letter), word, or line.
- **Cursor visibility** is whether and when a cursor is drawn. If you want the cursor to blink for a little while at the end choose **Always Visible** or choose **No Cursor** to disable it completely.
- **Cursor shape** is the style of the cursor. Underscore and block are the most common cursors, but you can choose another. In fact, you can type in something totally custom; however, only the first character will be used.
- **Blink rate** is how fast the cursor appears to be blinking where lower is faster. In version 25.10 you cannot set this below 5 to turn off blinking, but that will be fixed in the next version. For version 25.10 you can set this to its maximum 200 and try not to go over 200 frames while using **Cursor visibility** = **While Typing**.

- **Font** starts the standard text parameters. A monospace (fixed width) font is strongly recommended if you include a cursor; otherwise, the text will shift left and right as the cursor blinks. Also, typewriters were traditionally monospace.

Converts the image into a **pure black-and-white result** by comparing pixel values against a defined threshold level.

All pixels are forced to either black or white, removing grayscale detail and color information.



Threshold is a **spatial effect** applied per frame. It does not analyze motion or time.

Parameters

Level (0.0 - 100.0%)

Defines the cutoff point used to decide whether a pixel becomes black or white.

- **Low values (≈0.0 - 20.0%)**
Most pixels are classified as white. Only the darkest areas become black.
- **Mid values (≈40.0 - 60.0%)**
Balanced split between black and white. Midtones are pushed decisively to one side or the other.
- **High values (≈80.0 - 100.0%)**
Most pixels are classified as black. Only the brightest areas remain white.

Behavior note:

- The comparison is abrupt: there is no smoothing or feathering.
- Small changes around mid values can cause large visible shifts.

Invert (checkbox)

Swaps the black and white result.

- Pixels that would be white become black

- Pixels that would be black become white

This does not change how the threshold is calculated, only how the result is displayed.

Compare with alpha channel (checkbox)

Changes the source used for threshold comparison.

- **Unchecked**
Thresholding is based on the image's luminance (brightness).
- **Checked**
Thresholding is based on the **alpha channel** instead of luminance.

Important behavior notes:

- This option is meaningful only if the clip contains an alpha channel.
- When enabled, visible RGB color information is ignored.
- Useful for visualizing or isolating transparency data.

Keyframes

The **Level** parameter can be keyframed.

This enables:

- Animated transitions between black and white states
- Progressive reveals or wipes
- Time-based threshold effects driven by luminance or alpha

Visual characteristics

- Output contains only pure black and pure white
- No grayscale or color information
- Sharp, high-contrast edges
- Strong posterized appearance

Recommended use cases


- Creating matte or stencil-like visuals
- Isolating shapes based on brightness
- Graphic or abstract styling
- Preparing hard masks from luminance or alpha
Since none of the Mask filters can use a filtered clip object as a source, export the filtered clip as a frame or intermediate video file to use it with other things.
- Diagnostic visualization of alpha channels

Limitations

- No edge softness or anti-aliasing control
- Very sensitive to small Level changes
- Can produce flicker on noisy or compressed footage
- Not suitable for smooth tonal transitions



The **Timer** filter overlays a **time counter** on the video. It can count **up or down**, display time in various formats, and be styled and animated like a title.




This filter generates **visible graphics**, not metadata, and therefore must be applied in a context where video output exists.

Where the Timer must be applied

For the timer to be visible, the filter must be added to:

- A **video clip**, or
- The **Output track** on top of the stack (top-left of the timeline panel)



Warning:

If the Timer filter is added to an **empty video track**, the timer will **not be visible**, because there is no image to render onto.

When applied to:

- A **clip**: timing is relative to that clip.
- **Output**: timing is relative to the full project timeline.

Time controls

Format

Selects how time is displayed (for example: `HH:MM:SS.SSS`).

- `HH` = hours
- `MM` = minutes
- `SS` = seconds
- `SSS` = milliseconds

Different formats control the level of precision shown.

Direction

Controls how the timer progresses.

- **Up**
Counts upward from the starting value
- **Down**
Counts downward toward zero

Start Delay (default: 00:00:00.000)

Delays when the timer becomes active.

- The “>” button sets the start delay to the **current playhead position**
- When applied to a clip, the delay is **relative to the clip start**
- When applied to Output, the delay is **relative to the project start**

Duration (default: 00:00:10.000)

Defines how long the timer runs.

- Default is **10 seconds**
- The “>” button sets the duration so the timer **ends at the current playhead position**
- Like Start Delay, this is relative to the clip or Output context

Offset (default: 00:00:00.000)

Controls the **starting value** of the timer.

- **Direction = Up**

The timer counts upward starting from the Offset value

- **Direction = Down**

The timer counts downward starting from the Offset value

Speed (default: 1.00000)

Scales how fast the timer runs.

- **1.0** = real time (seconds)
- Higher values = faster counting
- Lower values = slower counting

This affects how Duration is interpreted.

Text appearance

Font

Selects the font used for the timer.

- Includes a **color selector**
- Font color can be **keyframed**
- A separate button selects the font family
- Optional **Use font size** checkbox enables manual font sizing

Outline

Controls the text outline.

- **Color** (keyframeable)
- **Thickness** (default 0)

Useful for improving readability over complex backgrounds.

Background

Adds a background behind the timer text.

- Color selector (transparent by default)
- Color can be **keyframed**
- **Padding** controls space around the text (default 0)

Opacity (0–100%)

Controls overall transparency of the timer overlay.

- Can be **keyframed**

- Affects text, outline, and background together

Position and size

Position / Size

Controls where the timer appears and how large it is.

- Both are **keyframeable**
- Can be adjusted numerically or **directly in the viewer**
- Viewer-based adjustment is often easier and more precise

Alignment

Controls how the timer fits within its bounding area:

- Horizontal: Left / Center / Right
- Vertical: Top / Middle / Bottom

Motion tracking integration

Load Keyframes from Motion Tracker

Loads motion-tracking data so the timer can **follow moving objects or people**.

This allows the timer to:

- Stay attached to a subject
- Move dynamically with the scene

Recommended use cases


- On-screen timers and countdowns
- Sports or performance timing
- Synchronization references
- Educational or instructional videos
- Motion-tracked overlays

Limitations

- Must be applied to a clip or Output to be visible
- Generates graphics, not metadata
- Timing behavior depends on clip vs Output context
- Requires motion tracking data to follow movement

Automatically applies **fade-in and fade-out** to **all video clips within a single track**.


The filter operates at the **track level**, not on individual clips, and is intended to automate consistent fades across multiple clips at once.



This filter is **not keyframeable** and applies the same fade behavior uniformly to every clip in the selected track.

Availability and scope

This filter is only visible and accessible when the **video track header** is selected, not when an individual clip is selected.



- It does **not** appear when a clip is selected.
- It applies to **all clips contained in the selected track**.
- Fades are applied automatically at:
 - the beginning of each clip (fade in)

- the beginning of each clip (fade in)
- the end of each clip (fade out)

This makes it a track-wide automation tool rather than a clip-level effect.

Parameters

Fade duration (0.01–1.00 s)

Defines the length of the fade applied at the start and end of each clip, in **seconds**.

- **Low values (≈0.01–0.20 s)**
Very fast fades. Transitions are barely noticeable.
- **Mid values (≈0.30–0.70 s)**
Smooth, conventional fades suitable for most editing workflows.
- **High values (≈0.80–1.00 s)**
Long, clearly visible fades.

Behavior notes:

- The same duration is used for both fade-in and fade-out.
- The duration is applied independently to each clip.

Fade color

Defines the color used for the fade.

Controls include:

- A **color swatch** to open the color palette
- A **color picker** to sample a color directly from the player preview
- A **Transparency** button to fade using alpha instead of a solid color

Behavior notes:

- Using a solid color produces a fade to/from that color.
- Using Transparency produces a fade to/from transparency, revealing underlying tracks.
- The chosen color applies uniformly to all clips in the track.

Fade direction indicators (non-editable)

Two radio buttons are displayed:

- **Fade in** (locked)
- **Fade out** (locked)

These are **informational only** and indicate that both fade-in and fade-out are applied automatically. They are not user controls.

Visual characteristics

- Consistent fade-in at the start of each clip
- Consistent fade-out at the end of each clip
- Uniform color or transparency behavior across the entire track
- No per-clip variation

Recommended use cases

- Applying consistent fades to many clips quickly
- Normalizing transitions in montage or compilation tracks
- Preparing tracks for insertion over background video
- Avoiding repetitive manual fade setup on each clip


Limitations

- Track-level only; cannot be applied to individual clips
- No per-clip customization
- No keyframe support
- Same duration and color for all clips
- Not suitable when clips require different fade timings



Creates a **motion trail effect** by blending the current frame with a number of previous frames.

Moving objects leave visible afterimages, while static areas remain largely unchanged.



Trails is a **temporal effect**. Its behavior depends on motion over time; evaluating it on a single frame is misleading.

Parameters

Amount (2 - 30 frames)

Defines how many previous frames are accumulated to produce the trail.

- **Low values (2 - 5 frames)**
Short, subtle trails. Motion is emphasized without excessive smearing.
- **Mid values (6 - 15 frames)**
Clearly visible trails. Moving objects leave extended afterimages that persist briefly.
- **High values (16 - 30 frames)**
Long, persistent trails. Motion can appear smeared, with multiple overlapping afterimages.

Behavior notes:

- Amount is measured in **frames**, not time.
The visual length of the trail depends on the project's frame rate.
- Higher values increase persistence, not brightness.
- The effect accumulates history; it does not extrapolate motion.

Visual characteristics

- Afterimages following moving objects
- Static regions remain mostly stable
- Trails fade gradually as new frames replace older ones
- Stronger effect on fast or high-contrast motion

Recommended use cases

- Stylized motion emphasis
- Abstract or experimental visuals
- Dance, light movement, or performance footage
- Visualizing motion paths
- Can increase the perceived speed of motion by leaving visible trails behind moving objects.
- Atmospheric or dream-like effects

Usage notes and tips

- Always preview during **playback** to judge the effect accurately.
- Consider frame rate when choosing **Amount**: higher frame rates produce shorter-looking trails for the same value.
- Combining Trails with opacity or color effects can enhance readability.
- Use modest values to avoid excessive smearing or ghosting.


Limitations

- Not keyframeable
- Can cause heavy motion smearing at high values
- Accumulates noise and compression artifacts
- Not suitable for precise motion analysis

Shotcut expects all image and video sources to have un-pre-multiplied alpha. It does not do an automatic conversion even when there is metadata indicating premultiplied.

Use this filter to convert an image or video from **premultiplied alpha** to **straight (unpremultiplied) alpha**.

This operation separates color values from the alpha channel so that RGB channels are no longer scaled by transparency. Visible changes occur only along semi-transparent edges, where color values were previously scaled by alpha.



Unpremultiply Alpha is a **technical correction filter**, not a visual effect.

Parameters

This filter has **no controls**.

What this filter does (conceptual)

In premultiplied alpha:

- RGB values are already multiplied by the alpha channel
- Semi-transparent pixels contain darker RGB values
- This format is common in compositing pipelines and intermediate processing

Unpremultiply Alpha reverses this process:

- RGB values are divided by alpha
- Color intensity is restored independently of transparency
- The image is converted to straight alpha representation

When the effect is visible (important)

If the clip **does not contain an alpha channel**, this filter produces **no visible change**.

Visible changes occur only when:

- The clip has an alpha channel **and**
- RGB values are premultiplied by that alpha

Note:

Visible changes occur only along semi-transparent edges, where color values were previously scaled by alpha.

Typical workflows

This filter is usually used:

- After operations that require premultiplied alpha processing
- Before exporting or compositing where straight alpha is expected
- To correct dark edges or halos around transparent areas
- In combination with **Premultiply Alpha** in advanced workflows

Example workflow:

1. Source with alpha
2. Premultiply Alpha
3. Apply effects that expect premultiplied data
4. **Unpremultiply Alpha**
5. Composite or export

Visual characteristics

- No change on opaque footage
- No change if no alpha channel is present
- On affected clips, restores correct color intensity near transparent edges
- Can remove dark fringes caused by premultiplied RGB values

Recommended use cases


- Correcting alpha handling in compositing pipelines
- Preparing clips with transparency for export
- Fixing edge darkening around semi-transparent regions
- Advanced workflows involving alpha-aware effects

Limitations

- No effect on clips without alpha
- No visual feedback if used incorrectly
- Requires understanding of alpha handling concepts
- Not intended for general-purpose editing

Creates a **temporal zoom feedback effect** by reusing and scaling previous frames.

The result resembles a visual echo where the image appears to zoom in or out continuously, producing a disorienting or hypnotic motion.



Vertigo is a **temporal effect**. Its appearance depends on motion over time; evaluating it on a single frame is misleading.

Parameters

Speed (0.00 - 100.00%)

Controls how quickly the feedback zoom evolves over time.

- **Low values**
Slow, subtle zoom drift. The effect builds gradually and is less distracting.
- **Mid values**
Clearly visible motion. The zoom progression becomes noticeable and rhythmic.
- **High values**
Fast zoom progression. The image rapidly expands or collapses, producing strong visual instability.

Behavior notes:

- Speed affects **temporal progression**, not zoom amount.
- Higher values increase how quickly the effect changes from frame to frame.
- Visual intensity depends on both Speed and Zoom together.

Zoom (0.00 - 100.00%)

Controls the **strength of the zoom applied to the feedback image**.

- **Low values**
Minimal scaling. The effect is subtle and may be hard to notice without motion.
- **Mid values**
Moderate scaling. The zoom feedback becomes clearly visible.
- **High values**
Strong scaling. The image rapidly grows or shrinks within itself, often revealing borders or heavy repetition.

Behavior notes:

- Zoom defines **how much** scaling is applied, not how fast.
- Extreme values can cause strong image repetition or cropping artifacts.

Keyframes

Both **Speed** and **Zoom** can be keyframed.

This enables:

- Gradual buildup or release of the vertigo effect
- Pulsing or oscillating zoom feedback
- Time-based transitions between stable and unstable visuals

Parameter interaction

- **Zoom** sets the magnitude of scaling applied to previous frames.
- **Speed** controls how quickly that scaling evolves over time.
- High Zoom with low Speed produces slow, dramatic distortion.
- High Speed with low Zoom produces fast but subtle motion.
- High values on both produce aggressive, disorienting results.

Visual characteristics

- Repeated zooming of the image into itself
- **Infinity mirror-like effect** caused by recursive frame feedback
- Motion trails caused by temporal feedback
- Progressive scaling artifacts
- Disorienting or hypnotic visual motion

Recommended use cases

- Stylized or experimental visuals
- Dream, hallucination, or disorientation effects
- Abstract motion backgrounds
- Emphasizing instability or psychological tension
- Transitional effects between scenes


Usage notes and tips

- Always evaluate the effect during **playback**, not on a still frame.
- Start with low Zoom values and increase gradually.
- Keyframing Speed often produces smoother results than abrupt changes.
- The effect is sensitive to frame rate and clip length.

Limitations

- Can become visually overwhelming at high values
- No control over direction (inward vs outward is implicit)
- Accumulates compression artifacts over time
- Not suitable for subtle correction tasks

The **Vibrance** filter adjusts color intensity in a **selective and adaptive way**. Unlike **saturation**, which affects all colors equally, vibrance primarily increases the intensity of **less-saturated colors** while limiting changes to colors that are already strong.



This helps enhance color richness while preserving natural-looking skin tones and avoiding clipping or oversaturation.

Parameters

Intensity (–100 to 100)

Controls the **overall strength** of the vibrance effect.

- Negative values
Reduce color intensity, especially in less-saturated areas
- Positive values
Increase color intensity selectively

At higher values, muted colors become more vivid while already saturated colors are affected less aggressively than with saturation.

Red (–100 to 100)

Adjusts vibrance **specifically in red hues**.

- Negative values
Reduce vibrance in reds
- Positive values
Increase vibrance in reds

Useful for fine-tuning skin tones, warm highlights, or red-dominant elements.

Green (–100 to 100)

Adjusts vibrance **specifically in green hues**.

- Negative values
Reduce vibrance in greens
- Positive values
Increase vibrance in greens

Commonly used for vegetation or correcting green casts.

Blue (–100 to 100)

Adjusts vibrance **specifically in blue hues**.

- Negative values
Reduce vibrance in blues
- Positive values
Increase vibrance in blues

Useful for skies, water, or cool-toned scenes.

Keyframes

All parameters of the **Vibrance** filter (**Intensity**, **Red**, **Green**, and **Blue**) can be **keyframed**, allowing vibrance strength and color emphasis to change over time.

This enables gradual color enhancement, animated color emphasis, or dynamic color correction across a clip.

Note:

The per-channel controls modify the vibrance effect and require a **non-zero “Intensity” value to have an effect**.

“**Intensity**” controls the overall strength of the vibrance effect.

The **Red, Green, and Blue** controls adjust how that vibrance is distributed across color channels and have **no effect** unless **Intensity** is **non-zero**.

Vibrance vs Saturation

Although related, vibrance and saturation behave differently:

- **Saturation**

Increases or decreases color intensity **uniformly across all colors**, including already saturated areas. This can quickly lead to clipped colors or unnatural skin tones.

- **Vibrance**

Adjusts color intensity **selectively**, focusing on colors that are less saturated and protecting already intense colors and skin tones.

As a result, vibrance produces a **more natural and controlled color boost**.

Visual characteristics

Typical effects include:

- Increased richness in muted colors
- Better preservation of highlights and skin tones
- Reduced risk of color clipping compared to saturation
- Subtle, adaptive color enhancement

Recommended use cases

- Enhancing flat or washed-out footage
- Improving color presence without oversaturation
- Correcting dull lighting conditions
- Fine-tuning specific color ranges after primary color correction


Vibrance is often best applied **before or instead of saturation**.

Limitations

- Not a replacement for full color grading
- Extreme values can still produce unnatural colors
- Per-channel adjustments affect only vibrance, not hue

The **Vignette** filter darkens the edges of the image while preserving a brighter area toward the center.

Historically, vignetting originated as an **optical artifact** in photography and cinema, caused by lens design limitations, mechanical obstructions, or light falloff toward the edges of the image circle. Over time, it became a **deliberate stylistic effect** used to draw attention to the subject.



In Shotcut, the vignette is applied as a **radial darkening mask** centered on the image.

Parameters

Radius (0-100%)

Controls the **size of the unaffected central area** of the vignette.

- **0%**
No clear center remains; the entire image is darkened
(At 100% Opacity, the frame becomes fully black)
- **Higher values**
Larger central area remains visible
Darkening is pushed further toward the edges

This parameter defines the radius of the vignette's inner region.

Feathering (0-500%)

Controls the **softness of the transition** between the central area and the darkened edges.

- **0%**
Hard, abrupt edge

The vignette boundary is clearly visible

- **Higher values**

Progressively softer and smoother transition

This controls the falloff of the vignette rather than its size.

Opacity (0-100%)

Controls the **strength of the darkening** applied by the vignette.

- **0%**

No visible effect

- **100%**

Full vignette strength

Opacity affects how dark the edges become but does not change the vignette's shape.

Keyframes

All three parameters (**Radius**, **Feathering**, and **Opacity**) can be **keyframed**, allowing the vignette to change over time.

This enables:

- Animated focus shifts
- Gradual darkening or brightening
- Stylized transitions or emphasis effects

Non-linear feathering

When enabled, the feathering follows a **non-linear curve**.

- Disabled

Linear, even falloff from center to edges

- Enabled

More natural, film-like falloff

The transition accelerates or decelerates instead of changing evenly

This option produces a smoother, more organic vignette similar to optical lens behavior.

Visual characteristics

Typical effects include:

- Darkened corners and edges
- Focused attention toward the center

- Hard or soft vignette boundaries depending on Feathering
- Natural-looking falloff when Non-linear feathering is enabled

Recommended use cases

- Directing viewer attention to a subject
- Simulating lens or optical characteristics
- Creating mood or atmosphere
- Subtle framing or stylized transitions using keyframes


Limitations

- Radial shape only (no custom mask)
- Uniform application across the frame
- Does not affect color, only luminance
- Extreme settings may obscure important details

Shotcut expects all image and video sources to have un-pre-multiplied alpha. It does not do an automatic conversion even when there is metadata indicating premultiplied.

Use this filter to convert an image or video from **premultiplied alpha** to **straight (unpremultiplied) alpha**.

This operation separates color values from the alpha channel so that RGB channels are no longer scaled by transparency. Visible changes occur only along semi-transparent edges, where color values were previously scaled by alpha.



Unpremultiply Alpha is a **technical correction filter**, not a visual effect.

Parameters

This filter has **no controls**.

What this filter does (conceptual)

In premultiplied alpha:

- RGB values are already multiplied by the alpha channel
- Semi-transparent pixels contain darker RGB values
- This format is common in compositing pipelines and intermediate processing

Unpremultiply Alpha reverses this process:

- RGB values are divided by alpha
- Color intensity is restored independently of transparency
- The image is converted to straight alpha representation

When the effect is visible (important)

If the clip **does not contain an alpha channel**, this filter produces **no visible change**.

Visible changes occur only when:

- The clip has an alpha channel **and**
- RGB values are premultiplied by that alpha

Note:

Visible changes occur only along semi-transparent edges, where color values were previously scaled by alpha.

Typical workflows

This filter is usually used:

- After operations that require premultiplied alpha processing
- Before exporting or compositing where straight alpha is expected
- To correct dark edges or halos around transparent areas
- In combination with **Premultiply Alpha** in advanced workflows

Example workflow:

1. Source with alpha
2. Premultiply Alpha
3. Apply effects that expect premultiplied data
4. **Unpremultiply Alpha**
5. Composite or export

Visual characteristics

- No change on opaque footage
- No change if no alpha channel is present
- On affected clips, restores correct color intensity near transparent edges
- Can remove dark fringes caused by premultiplied RGB values

Recommended use cases

- Correcting alpha handling in compositing pipelines
- Preparing clips with transparency for export
- Fixing edge darkening around semi-transparent regions
- Advanced workflows involving alpha-aware effects

Limitations

- No effect on clips without alpha
- No visual feedback if used incorrectly
- Requires understanding of alpha handling concepts
- Not intended for general-purpose editing



The White Balance filter works well by following these steps:

Restore “Neutral color” and “Color temperature” to their defaults (6500 degrees).

Click on the eyedropper and position the crosshairs over an object which is supposed to be white, e.g. a white card, white shirt, etc. and click.

Decodes Ambisonic audio (a spatial/3D audio format) into the project's configured output channel layout and allows rotation of the sound field in 3D space.

The filter's available controls and modes depend on **Settings > Audio Channels**.


Ambisonic Decoder is a spatial audio filter. Its effect is temporal and should be evaluated during playback, but also some of the audio scopes in **View > Scopes** can help by visualization.

Availability and mode behavior

The **Mode** control and its options depend on the project setting: **Settings > Audio Channels**.


1 (Mono)

- **Mode:** not displayed
- The filter exposes orientation controls only.




2 (Stereo)

- **Mode:**
 - Stereo
 - Binaural




4 (Quad / Ambisonics)

- **Mode:**
 - Binaural
 - Quad
 - Ambisonic



6 (5.1 Surround)

- **Mode:** not displayed



Binaural mode restriction

“Binaural” audio provides an immersive, 3D audio experience designed specifically for headphones.

When **Binaural** is selected (in any configuration where it is available):

- **Yaw, Pitch, Roll, and Zoom** are locked
- Sliders remain visible but cannot be adjusted

This reflects that binaural decoding uses a fixed head-related transfer function (HRTF) and does not expose orientation controls in this mode.

Parameters

Mode

Defines how the ambisonic sound field is decoded to the project’s output format.

- **Stereo**
Decodes spatial information into a conventional two-channel stereo image.
- **Binaural**
Produces headphone-oriented spatial audio using binaural decoding. Orientation controls are disabled.
- **Quad**
Decodes to four-channel surround output.
- **Ambisonic**
To use this filter useful to transform or rotate the spatial audio.

Yaw (–360° to +360°)

Rotates the sound field **horizontally** around the listener.

- Equivalent to turning left or right.

- Full $\pm 360^\circ$ rotation allows continuous or looping motion.

Pitch (-180° to $+180^\circ$)

Rotates the sound field **vertically**.

- Tilts the perceived sound environment upward or downward.
- Affects perceived elevation of sound sources.

Roll (-180° to $+180^\circ$)

Rotates the sound field around the **forward axis**.

- Tilts the entire sound field clockwise or counterclockwise.
- Useful when audio orientation must match camera roll.

Zoom (-100.0% to $+100.0\%$)

Controls the visual **scale** of the orientation reference displayed in the preview player.

- **-100%** — very small reference cross
- **0%** — medium (default) size
- **+100%** — large reference cross

Important clarification:

Zoom affects the visual reference overlay only.

It does not change audio loudness or spatial intensity.

Paste Parameters (button)

Applies previously copied Ambisonic orientation parameters to this filter.

Typical use cases:

- Synchronizing orientation across multiple clips
- Reusing orientation data from another Ambisonic Decoder instance
- Matching audio rotation to externally generated metadata

Viewer overlay (orientation reference)

When the filter is active, a cross-shaped overlay appears in the preview player:

- **Vertical axis:**
 - Blue at the top → Green at the bottom (gradient)
- **Horizontal axis:**
 - Red on the left → Yellow on the right (gradient)

This overlay:


- Visualizes the current orientation reference
- Updates when Yaw, Pitch, Roll, or Zoom change
- Is not rendered into exports

Interactive orientation control (viewer)

The orientation cross displayed in the preview player is interactive.

- The cross can be rotated directly with the mouse in the preview player.
- Horizontal dragging adjusts **Yaw**.
- Vertical dragging adjusts **Pitch**.
- The mouse wheel adjust **Zoom**.
- Changes made via the overlay are immediately reflected in the corresponding sliders and keyframes.

This interaction provides a direct, visual way to adjust sound field orientation without relying solely on numeric controls.



Although the overlay itself is not exported, **Zoom** is keyframeable because:

- It allows the reference to remain readable as orientation changes
- Large rotations may require scaling the reference for clarity
- It supports animation workflows where orientation evolves over time

The keyframes affect how the orientation reference is displayed, not the audio signal.

Keyframes

- **Yaw, Pitch, Roll, and Zoom** are keyframeable
- Keyframes enable:
 - Rotating the sound field over time
 - Matching audio orientation to camera movement
 - Creating dynamic spatial audio motion

Keyframes are ignored when Binaural mode locks the controls.

Parameter interaction

- **Yaw, Pitch, Roll** define listener orientation within the Ambisonic sound field
- **Zoom** scales the orientation reference overlay
- **Mode** determines whether orientation controls are active and how decoding is performed

Auditory characteristics

- No **visual** effect is exported
- Audible rotation and repositioning of sound sources
- Strongest spatial cues in multi-channel and Ambisonic modes
- Limited directional cues in stereo
- Fixed orientation in binaural mode

Recommended use cases

- Decoding ambisonic recordings
- Converting Ambisonics to surround sound
- Spatial audio for 360° or VR video
- Matching audio orientation to camera motion
- Controlled rotation of sound environments
- Advanced spatial sound design


Limitations

- Requires Ambisonic or spatial audio sources to be meaningful
- Binaural mode disables orientation controls
- Stereo output cannot reproduce full 3D spatial cues



Encodes a **mono or conventional audio signal** into an **ambisonic sound field** by assigning it a virtual position in 3D space.

Instead of decoding spatial audio for playback, this filter performs the opposite operation: it **places a sound source into an ambisonic environment**.



Ambisonic Encoder is a spatial audio generation filter. Its effect is temporal and must be evaluated during playback.

Parameters

Azimuth (-360° to +360°)

Controls the **horizontal angle** of the virtual sound source around the listener.

- **0°**
Sound is positioned directly in front.
- **Positive / negative values**
Rotate the sound source clockwise or counterclockwise around the listener.
- The full $\pm 360^\circ$ range allows continuous rotation or looping motion.

Azimuth corresponds to left–right positioning around the listener.

Elevation (-360° to +360°)

Controls the **vertical angle** of the virtual sound source.

- **0°**
Sound is positioned at ear level.
- **Positive values**
Move the sound source upward.

- **Negative values**
Move the sound source downward.

Elevation corresponds to up–down positioning.

Keyframes

Both **Azimuth** and **Elevation** are keyframeable.

This enables:

- Moving a sound source through space over time
- Circular or spiral motion paths
- Synchronizing sound movement with on-screen action
- Static placement or animated spatial motion

Viewer overlay (position reference)

When the Ambisonic Encoder filter is active, a rectangular overlay appears in the preview player.

- A semi-transparent dot represents the current virtual sound position.
- The horizontal axis corresponds to **Azimuth**.
- The vertical axis corresponds to **Elevation**.

Behavior and interaction

- Moving **Azimuth** shifts the dot left or right.
- Moving **Elevation** shifts the dot up or down.
- The dot can move outside the rectangle, indicating angles beyond the primary reference range.
- The dot can also be dragged directly with the mouse in the preview player:
 - Horizontal dragging updates **Azimuth**
 - Vertical dragging updates **Elevation**
 - Changes are reflected immediately in the corresponding sliders

Note:

The overlay is a visual and interactive reference only and is not rendered into exports.

What the overlay represents

The rectangle and dot form a **visual positioning reference** for the encoded sound source:

- It shows **directional placement**, not distance or loudness.
- It updates in real time as parameters change.
- It is a **guide only** and is **not rendered into exports**.

The overlay helps correlate numeric angle values with perceived spatial movement.

Purpose and typical workflow

The Ambisonic Encoder is used to:

- Convert a mono sound into ambisonic space
- Define where a sound exists within a 3D sound field
- Prepare audio for later decoding with [Ambisonic Decoder](#)
- Create spatial motion before mixing or rendering

Typical workflow:

1. Encode individual sound sources using **Ambisonic Encoder**
2. Animate their positions using keyframes
3. Decode the resulting ambisonic mix using **Ambisonic Decoder**

Auditory characteristics

- Audible changes in perceived sound direction
- No visual output is rendered
- Motion is smooth when keyframes are interpolated
- Perceptual clarity depends on the decoding method used later

Recommended use cases

- Spatial and surround sound design
- 360° or VR audio workflows
- Placing sound effects in a virtual environment
- Animated sound movement
- Preparing sources for Ambisonic mixing

Limitations

- Does not control distance or attenuation
- No control over spread or width
- Requires Ambisonic decoding later to be audible as spatial audio



Applies a **band-pass filter** that allows only a selected range of frequencies to pass while attenuating frequencies below and above that range.

This filter is used to isolate specific frequency content, shape tone, or create stylized audio effects.

Band Pass is a **frequency-domain audio filter**. Its effect is evaluated continuously over time and can be animated using keyframes.

Parameters

Center frequency (5 - 21600 Hz)

Defines the **center point** of the frequency range that will pass through the filter.

- **Low values (≈5 - 300 Hz)**
Emphasize bass and low-frequency content.
- **Mid values (≈300 - 3000 Hz)**
Emphasize midrange frequencies, including most speech intelligibility.
- **High values (≈3 000 - 21600 Hz)**
Emphasize high-frequency detail such as sibilance or hiss.

Behavior note:

- This does not define a single frequency, but the midpoint of the allowed band.

Bandwidth (5 - 21600 Hz)

Defines the **width of the frequency band** centered around the Center frequency.

- **Low values**
Narrow band. Only a small frequency range is audible.
- **High values**
Wide band. More of the surrounding frequencies are included.

Important clarification:

Bandwidth controls *how wide* the pass region is, not how strong the filter sounds.

Rolloff (1 - 10)

Controls how sharply frequencies outside the band are attenuated.

- **Low values ($\approx 1 - 3$)**
Gentle slope. Frequencies outside the band fade gradually.
- **High values ($\approx 7 - 10$)**
Steep slope. Frequencies outside the band are cut more aggressively.

Rolloff affects the **transition sharpness**, not the band width itself.

Dry - Wet (0.0 - 100.0%)

Controls the mix between the original signal and the filtered signal.

- **0.0%**
Original (unfiltered) audio only.
- **50.0%**
Equal mix of original and filtered audio.
- **100.0%**
Filtered audio only.

Channel selection (buttons)


Determines **which audio channels** the filter is applied to.
These buttons are **not keyframeable**.

Mono / Stereo projects

Settings → **Audio Channels** → **1 (Mono)** or **2 (Stereo)**

Channels

- **L** — Enables or disables the filter on the left channel
- **R** — Enables or disables the filter on the right channel




In mono projects, enabling either channel affects how the mono signal is routed into stereo output.

Quad / Ambisonics projects

Settings → Audio Channels → 4 (Quad / Ambisonics)

Channels

- **L** — Front Left
- **R** — Front Right
- **Ls** — Surround Left
- **Rs** — Surround Right




Each channel can be filtered independently.

5.1 Surround projects

Settings → Audio Channels → 6 (5.1 Surround)

Channels

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right



This allows precise frequency control per speaker channel.

Keyframes

All four sliders (**Center frequency**, **Bandwidth**, **Rolloff**, **Dry - Wet**) are keyframeable.

This enables:

- Frequency sweeps
- Moving “radio” or “telephone” effects
- Time-based tonal changes
- Dynamic emphasis of different frequency ranges

Channel selection buttons are static.

Auditory characteristics

- Audible isolation of a specific frequency range
- Removal of bass and treble outside the selected band
- Narrow bandwidth produces a thin, focused sound
- Wide bandwidth sounds more natural but still filtered

Recommended use cases

- Telephone or radio voice effects
- Isolating dialogue frequencies
- Sound design and special effects
- Removing unwanted low or high frequencies
- Frequency sweeps and transitions
- Channel-specific filtering in surround mixes

Usage notes and tips

- Start with a **midrange center frequency** ($\approx 1 - 2$ kHz) for speech effects.
- Narrow bandwidth + high rolloff produces classic “band-limited” sounds.
- Use **Dry - Wet** to soften the effect instead of widening the band.
- In surround projects, filtering only the **center channel** is useful for dialogue shaping.

Limitations

- Not a full equalizer
- No visual frequency display
- Extreme settings can sound artificial
- LFE channel filtering may have limited audible effect on some systems



Adjusts the **relative levels between audio channels** to reposition sound within a multi-channel field.

The filter redistributes existing audio energy between channels; it does not add spatialization or effects.

The available controls and layout depend on the project's audio channel configuration.

Balance is a **spatial level-mapping filter**. It is evaluated per frame and can be animated with keyframes.


Availability and UI behavior

The Balance filter adapts its interface based on:

Settings → **Audio Channels**

1 (Mono) or 2 (Stereo)

- The filter exposes one slider:
 - **Left - Right** (keyframeable)




This provides a simplified balance control for projects without discrete surround channels.

4 (Quad / Ambisonics) or 6 (5.1 Surround)

- The filter exposes:
 - An **Operation** dropdown
 - One slider at a time, whose meaning depends on the selected operation
- All sliders are **keyframeable**

- Only the slider relevant to the selected operation is visible



Parameters (Mono / Stereo projects)

Settings → **Audio Channels** → **1 (Mono)** or **2 (Stereo)**

The filter has **one slider**:

Left - Right (0.00 - 1.00)

Balances audio between the left and right channels.

- **0.00**
Full emphasis on the left channel
- **0.50**
Centered balance
- **1.00**
Full emphasis on the right channel

Notes:

- In **mono** projects, this control determines how the mono signal is distributed between left and right output.
- In **stereo** projects, it adjusts the relative level between the existing left and right channels.

In mono or stereo projects, this control provides a conceptual balance rather than discrete channel routing.

Parameters (Quad / 5.1 projects)

Operation

Selects which channel relationship is being adjusted.
The visible slider and its meaning change based on this selection.

Front Balance → **Left - Right (0.00 - 1.00)**

Balances left vs right levels within the front channels.

- **0.00** → Front-left emphasis
- **0.50** → Centered
- **1.00** → Front-right emphasis

Surround Balance → Left - Right (0.00 - 1.00)

Balances left vs right levels within the surround channels.

- **0.00** → Surround-left emphasis
- **0.50** → Centered
- **1.00** → Surround-right emphasis

Front + Surround... → Left - Right (0.00 - 1.00)

Balances left vs right across all channels, front and surround together.

- **0.00** → Left-side emphasis
- **0.50** → Centered
- **1.00** → Right-side emphasis

Left Fade → Front - Surround (0.00 - 1.00)

Adjusts front vs surround balance for the left-side channels only.

- **0.00** → Front-left emphasis
- **0.50** → Balanced
- **1.00** → Surround-left emphasis

Right Fade → Front - Surround (0.00 - 1.00)

Adjusts front vs surround balance for the right-side channels only.

- **0.00** → Front-right emphasis
- **0.50** → Balanced
- **1.00** → Surround-right emphasis

Left + Right Fade → Front - Surround (0.00 - 1.00)

Adjusts front vs surround balance for all channels simultaneously.

- **0.00** → Front channels emphasized
- **0.50** → Balanced
- **1.00** → Surround channels emphasized

Keyframes

All sliders are keyframeable.

This enables:

- Animated panning across left/right channels
- Moving sound from front to surround over time
- Dynamic spatial rebalancing during a scene

Parameter interaction

- **Operation** determines *which channel relationship* is affected.
- The visible slider always operates on a 0.00 - 1.00 normalized scale.
- Changes redistribute level between channels without altering overall timing or adding effects.

Auditory characteristics

- Smooth level transitions between channels
- No added reverb or spatial modeling
- Perceived sound position shifts within the existing channel layout
- Most noticeable on discrete multi-channel systems

Recommended use cases


- Correcting imbalanced mixes
- Repositioning dialogue or effects
- Creating motion across speaker layouts
- Adapting audio for different surround configurations
- Fine-tuning spatial placement without re-encoding audio

Limitations

- Does not create spatial cues from mono sources
- No distance, depth, or elevation control
- Effectiveness depends on playback system
- Stereo playback cannot reproduce full surround intent



Copies the audio signal from one channel to another within the same clip. This filter performs a **direct channel duplication**, replacing the destination channel with the source channel's signal.



Copy Channel is a **routing utility filter**. It does not mix, sum, pan, or process audio levels or frequencies.

Parameters

Copy from

Selects the **source channel** whose signal will be copied.

The available options depend on the project's **Audio Channels** setting.

To

Selects the **destination channel** that will receive the copied signal.

The destination channel's original content is *replaced* by the source channel.

Channel availability by project configuration

Check **Settings > Audio Channels**

1 (Mono)

- **Copy from:** inactive
- **To:** inactive

With only one channel, there is nothing to copy or replace.

2 (Stereo)

- **Copy from:**

- Left
- Right

- **To:**

- Left
- Right

Allows copying one stereo channel into the other, which is useful when one channel is empty or silent.

4 (Quad / Ambisonics)

- **Copy from / To:**

- Front left
- Front right
- Left surround
- Right surround

Enables duplication between any of the four channels.

6 (5.1 Surround)

- **Copy from / To:**

- Front left
- Front right
- Center
- Low frequency (LF)
- Left surround
- Right surround

Allows precise channel reassignment within a surround mix.

What this filter does

- Copies audio one-to-one from a source channel
- Replaces the destination channel's content
- Does not blend or sum channels
- Does not change loudness, timing, or tone
- Operates only on the selected clip

Auditory characteristics

- Destination channel becomes identical to the source channel

- Original destination audio is lost for the duration of the clip
- No change in overall loudness unless combined with other filters

Recommended use cases

- Fixing incorrectly routed channels
- Duplicating a mono signal across multiple channels
- Replacing a missing or silent channel
- Copying dialogue into the center channel in surround projects
- Preparing material for further channel-specific processing

Usage notes and tips


- Use **Mute** to silence unwanted channels before copying, if needed.
- To copy one channel into multiple channels, add multiple instances of the filter.
- This filter is often used early in the filter chain for routing corrections.
- For blending channels instead of replacing them, use mixing or **Downmix** filters instead.

Limitations

- No mixing or gain control
- No keyframing
- No time-based behavior

Reduces the **dynamic range** of an audio signal by attenuating levels that exceed a defined threshold.

Loud sounds are reduced, while quieter sounds are left unchanged, resulting in more consistent overall loudness.



This compressor operates in the **amplitude domain** and is applied continuously over time.

Parameters

RMS → Peak (0.0 - 100.0%)

Controls how the compressor detects signal level.
Lower values = smoother, more musical compression.
Higher values = tighter control of peaks.

- **0.0% (RMS)**
Detection is based on average signal energy.
Compression reacts smoothly and is less sensitive to short transients.
- **100.0% (Peak)**
Detection is based on instantaneous peaks.
Compression reacts quickly to sharp transients.
- **Intermediate values**
Blend between RMS and Peak behavior.

Attack (2 - 400 ms)

Controls how quickly compression begins **after** the signal exceeds the threshold.

- **Short attack**
Quickly reduces transients (drums, plosives).
- **Long attack**
Allows transients to pass before compression engages.

Release (2 - 800 ms)

Controls how quickly compression stops **after** the signal falls below the threshold.

- **Short release**
Compression disengages quickly; can sound aggressive or “pumping”.
- **Long release**
Smoother, more natural recovery.

Threshold (-30 to 0 dB)

Sets the level above which compression is applied.

- **Lower values**
More of the signal is compressed.
- **Higher values**
Only loud peaks are affected.

Ratio (1:1 to 1:20)

Defines how much the signal is reduced once it exceeds the threshold.

- **1:1**
No compression.
- **1:2 to 1:4**
Gentle to moderate compression.
- **1:8 to 1:20**
Strong compression, approaching limiting.

Knee radius (1 - 10 dB)

Controls how smoothly compression transitions around the threshold.

- **Low values**
Hard knee: compression engages abruptly.
- **High values**
Soft knee: compression engages gradually.

Soft knees are often perceived as more natural.

Makeup gain (0.0 - 24.0 dB)

Boosts the output level after compression.

- Used to restore loudness lost due to gain reduction.
- Does not affect how compression is calculated.

Channel selection (toggle buttons)

Determines which audio channels the compressor is applied to. These buttons are toggle switches.

Note

Available buttons depend on the project's audio channel configuration in **Settings > Audio Channels**.

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Each button enables or disables compression on that channel independently.

Gain Reduction meter (visual indicator)

A visual-only meter labeled **Gain Reduction** (-24 dB to 0 dB) shows how much level is being reduced by the compressor.

Important behavior notes:

- It reflects reduction, not output level.
- To see meaningful movement, sufficient input level is required; a **Gain / Volume** or other level-adjusting filter can appear earlier in the filter chain if needed.
- Without sufficient input level, the meter may remain static even if the compressor is active.

The meter is informational only.

Visual / auditory characteristics

- Reduced volume differences between loud and quiet sections
- More consistent perceived loudness

- Potential reduction of transients depending on settings
- Overuse can produce pumping or flattened dynamics

Recommended use cases

- Leveling dialogue
- Controlling peaks in music or sound effects
- Reducing dynamic swings in narration
- Preparing audio for broadcast or online delivery
- Channel-specific dynamics control in surround mixes

Usage notes and tips

- Start by setting **Threshold** and **Ratio** first.
- Use **Attack** and **Release** to shape transient behavior.
- Prefer **RMS**-heavy detection for dialogue and music.
- Use **Makeup gain** sparingly to restore loudness.
- Watch the **Gain Reduction** meter to avoid over-compression.

Limitations

- No sidechain input
- No lookahead control
- No keyframe support
- Meter depends on upstream signal level

More about dynamic range compression


Dynamic range compression

Dynamic range compression (DRC) or simply compression is an audio signal processing operation that reduces the volume of loud sounds or amplifies quiet sounds, thus reducing or compressing an audio signal's dynamic range. Compression is commonly used in sound recording and reproduction, broadcasting, live sound reinforcement and some instrument amplifiers. A dedicated electronic hardware unit or audio software that applies compression is called a compressor. In the 2000s, compressors became a...



Adds a **time-delayed repetition** of the audio signal, creating echo-like effects ranging from subtle spatial depth to pronounced rhythmic repeats.

The filter mixes delayed audio back into the original signal and can optionally feed repeats back into the delay line.



Delay is a **temporal** audio effect evaluated over time.

Parameters

Delay (0.00 - 4.00 s (second))

Sets the time offset between the original signal and the delayed signal.

- **Very low values (≈ 0.00 - 0.05 s)**
Short delays that can thicken sound or create a doubling/comb-filter-like character.
- **Short to mid values (≈ 0.05 - 0.50 s)**
Audible echoes that add space or rhythmic emphasis.
- **Long values (≈ 0.50 - 4.00 s)**
Distinct, separated repeats suitable for dramatic echo effects.

Behavior note:

- Delay time determines *when* repeats occur, not how loud they are.

Feedback (-70.0 - 0.0 dB)

Controls the gain applied to the delayed signal before it is fed back into the delay line.

The value is expressed in decibels (dB), meaning it represents a gain reduction, not a percentage.

- **-70.0 dB**

The feedback signal is effectively muted.

The delayed signal is **not fed back**, resulting in a **single echo only**.

- **Lower negative values (e.g. -40 to -20 dB)**

The feedback signal is strongly attenuated.

Each repeat is much quieter than the previous one, producing **a few rapidly decaying echoes**.

- **Higher values (e.g. -10 to -3 dB)**

The feedback signal is only slightly attenuated.

Repeats decay slowly, producing **many audible echoes**.

- **0.0 dB**

No attenuation is applied to the feedback signal.

Each repeat is fed back at full level, which can result in **very long decay or sustained resonance**.

Note:

- **Feedback** is expressed in decibels because it controls the gain applied to the delayed signal before it is re-injected into the delay line.
- High feedback values can cause prolonged ringing or runaway echo, especially with short delay times.

Dry → Wet (0.0 - 100.0%)

Controls the mix between the original (dry) signal and the delayed (wet) signal.

- **0.0%**

Original signal only.

- **50.0%**

Equal mix of dry and delayed audio.

- **100.0%**

Delayed signal only.

Channel selection (toggle buttons)

Determines **which audio channels** the delay is applied to.

These buttons are toggle switches.

Note

Available buttons depend on the project's audio channel configuration in **Settings > Audio Channels**.

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Each channel can be enabled or disabled independently.

Auditory characteristics

- Audible echoes or repeats following the original sound
- Increased sense of space or depth at moderate settings
- At short delay times with high feedback, a resonant or metallic tone may appear
- High feedback can produce long decays or continuous echo tails

What the Delay filter produces

- Echoes / repeats of the original signal
- Resonance or ringing at short delay times combined with high feedback
- Metallic or comb-filter-like coloration in extreme settings

What it does *not* produce

- No random noise is generated
- No hiss or grain is added by the filter itself

Any harsh or unstable sound comes from feedback accumulation, not noise generation.

Recommended use cases

- Adding depth or ambience to dialogue or effects
- Creating rhythmic echo effects on music or sound design
- Simulating distance or reflections
- Stylized or dramatic audio emphasis
- Channel-specific echoes in surround mixes

Usage notes and tips


- Start with short delay times and low feedback for subtle enhancement.
- Use **Dry** → **Wet** to control presence without changing delay timing.
- Be cautious with high feedback; monitor output levels.
- In surround projects, applying delay only to surround channels can enhance immersion without affecting clarity.

Limitations

- No tempo synchronization
- No keyframe support
- High feedback can cause excessive ringing
- Delay time is fixed during playback

Converts **multi-channel audio into a single mixed signal** and outputs that signal uniformly to the project's configured output channels.

All input channels are summed together without regard to channel role or frequency content.



Downmix is a channel-routing and summing filter, not a spatial or frequency-aware processor.

What this filter does

- Sums all input channels into one mono signal
- Sends the same mixed signal to the selected output channel(s)
- Does not preserve surround positioning, center emphasis, or low frequency (LF) separation
- Does not perform weighting, panning, or frequency-based routing

This is a direct channel-summing operation, not a surround-aware conversion.

Parameters

Target channels

Selects which output channel(s) receive the down-mixed mono signal.

Note

Available buttons depend on the project's audio channel configuration:
Settings > Audio Channels

Typical targets include:

- **L / R** — Front left and right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls / Rs** — Surround left and right

The same mono signal is sent to the selected target channel(s).

Channel selection behavior (source-dependent)

When the **Downmix** filter is applied:

- If the **project** uses more than two audio channels but, for instance, the source clip is stereo (2 channels):
 - **L** and **R** target channels are selected automatically
 - Other channels (**C**, **LF**, **Ls**, **Rs**) are not selected by default



- All target channels remain manually selectable

This default reflects the source channel layout rather than the project layout.

Note:

The filter does not up-mix automatically. Only channels present in the source are preselected; additional target channels must be enabled manually.

Channel behavior and limitations

- All input channels are treated equally and summed together.
- The filter does not:
 - detect dialogue
 - prioritize the center channel
 - isolate bass for **LF**
 - apply loudness compensation
- Output level depends on the combined input level and may require gain adjustment downstream.

Auditory characteristics

- Audible collapse of spatial separation
- Identical signal on selected output channels
- Surround information is lost
- Potential increase in overall level due to summing

Typical use

- Typical use cases include routing down-mixed stereo recordings to the **Center** channel for dialogue or spoken voice, or targeting the **LF** channel for bass-heavy material such as effects or music.

Recommended use cases

- Converting stereo audio to mono
- Creating dual-mono output (same signal on left and right)
- Preparing dialogue tracks for center-channel placement
- Routing summed audio to **LF** for bass-heavy material
- Simplifying complex channel layouts for compatibility


Usage notes and tips

- After down-mixing, use **Gain / Volume** or **Limiter** to control output level.
- When targeting the **Center** channel, this filter is useful for spoken content.
- When targeting **LF**, ensure the material is appropriate for low-frequency playback.
- For surround-aware or weighted down-mixing, use dedicated mixing workflows instead.

Limitations

- No channel weighting or intelligent mixing
- No frequency-aware routing
- No preservation of spatial cues
- Can cause clipping if summed levels are high

Applies a **graphic equalizer** with 15 fixed frequency bands, allowing precise adjustment of tonal balance across the audible spectrum. Each band boosts or attenuates a specific frequency range, shaping how bass, mids, and treble are perceived.



This filter operates in the **frequency domain**. All band controls are **keyframeable as a group**, not individually.

Parameters

Frequency bands

Each slider controls gain (± 20 dB) for a fixed frequency band. Positive values boost that range; negative values attenuate it.

Below is an *explanation* of what each band typically affects. These descriptions are indicative only; real content often spans multiple bands.

50 Hz

Very low bass.

- Felt more than heard
- Sub-bass, rumble, impact
- Too much can sound boomy or cause speaker stress

Examples: deep bass effects, explosions, very low musical notes

100 Hz

Bass body.

- Weight and fullness
- Excess can sound muddy

Examples: bass guitar, kick drum body, low warmth in voices

156 Hz

Upper bass / low warmth.

- Adds thickness
- Too much reduces clarity

Examples: male voice warmth, lower instruments

220 Hz

Low-mid transition.

- Body of many sounds
- Can easily sound boxy

Examples: lower speech resonance, room tone

311 Hz

Low mids.

- Important for natural tone
- Excess often causes muddiness

Examples: voice body, acoustic instruments

440 Hz

Midrange foundation.

- Tonal balance area
- Can sound nasal if over-boosted

Examples: general speech tone, many instruments

622 Hz

Midrange clarity.

- Improves intelligibility
- Overuse can sound honky

Examples: spoken voice definition

880 Hz

Upper midrange.

- Presence and articulation
- Too much can sound harsh

Examples: consonants in speech, instrument attack

1250 Hz

Presence region.

- Speech intelligibility
- Helps sounds “cut through”

Examples: dialogue clarity, lead instruments

1750 Hz

Upper presence.

- Detail and bite
- Excess can fatigue the ear

Examples: crisp speech, edge in vocals

2500 Hz

High-mid presence.

- Sharpness and attack
- Sensitive range for human hearing

Examples: speech intelligibility, clarity boost

3500 Hz

Definition and bite.

- Enhances detail
- Can become aggressive quickly

Examples: dialogue articulation, string attack

5000 Hz

High presence.

- Brilliance and detail
- Too much causes sibilance

Examples: “s” sounds, cymbal detail

10 000 Hz

Treble.

- Air and brightness
- Excess sounds hissy

Examples: sparkle in music, vocal air

20 000 Hz

Extreme high frequencies.

- Mostly air and sheen
- Often inaudible on many systems

Examples: subtle brightness, high-end polish

Keyframes (grouped)

All 15 sliders are keyframed together.

- Individual bands cannot be animated independently.
- Keyframes store the entire EQ curve at each point in time.

This enables:

- Tonal changes over time
- Scene-dependent EQ
- Gradual transitions between different EQ shapes

Channel selection

Determines which audio channels the equalizer is applied to. They are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings >**

Audio Channels

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent equalization per channel requires adding separate instances of the Equalizer filter, each targeting different channels.

Auditory characteristics

- Changes tonal balance, not loudness alone
- Boosting multiple adjacent bands increases perceived fullness
- Excessive boosts can introduce harshness or resonance
- Large cuts can make audio sound thin or distant

Recommended use cases

- Improving dialogue clarity
- Reducing muddiness or harshness
- Enhancing bass or brightness
- Matching tonal balance between clips
- Channel-specific EQ in surround mixes

Usage notes and tips

- Start with small adjustments (± 2 - 4 dB).
- For dialogue:
 - Reduce low bands (50 - 200 Hz) to remove rumble
 - Gently boost 1 - 3 kHz for intelligibility
- For bass-heavy content:
 - Boost low bands carefully
 - Avoid excessive overlap into low mids
- If the sound becomes harsh or fatiguing, reduce high-mid bands first.
- Always evaluate changes during playback.

Limitations


- Fixed frequency bands (no custom frequencies)
- No Q (bandwidth) control
- Grouped keyframes only

- Extreme boosts can cause distortion or listener fatigue



Equalizer: 3-Band (Bass & Treble) Audio Filter

Applies a **simple three-band equalizer** designed for quick tonal adjustments. Unlike multi-band equalizers, this filter focuses on broad frequency regions, making it suitable for fast, intuitive correction rather than detailed shaping.



This filter operates in the **frequency domain**. All three controls are keyframeable as a group, not individually.

Parameters

Low (Bass)

Controls low-frequency content.

- Adds or removes bass weight
- Too much boost can sound boomy or muddy
- Cutting reduces rumble and low-end clutter

Examples: bass guitar, kick drum, low warmth in voices

Mid

Controls the midrange, where most speech and musical detail lives.

- Boosting improves clarity and presence
- Excessive boost can sound nasal or harsh
- Cutting can make audio sound distant or hollow

Examples: dialogue intelligibility, body of instruments

High (Treble)

Controls high-frequency content.

- Boosting adds brightness and air
- Too much boost causes hiss or sibilance
- Cutting softens sharp or brittle sounds

Examples: consonants in speech, cymbals, high detail

Keyframes (grouped)

All three sliders are keyframed together.

- Individual bands cannot be animated independently
- Each keyframe stores the full three-band EQ state

This enables gradual tonal shifts over time, such as moving from muffled to clear audio.

Channel selection (toggle buttons)

Determines which audio channels the equalizer is applied to. Buttons are toggle switches and are not keyframeable.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Important note:

Independent equalization per channel requires adding separate instances of the filter.

Auditory characteristics

- Broad tonal changes rather than precise adjustments
- Boosting Low and High together can create a “smile” curve
- Excessive boosts may introduce resonance or listener fatigue
- Cuts are often more natural than large boosts

Recommended use cases

- Quick tonal correction
- Improving dialogue clarity

- Adding warmth or brightness
- Matching tonal balance between clips
- Simple adjustments in multi-channel projects

Usage notes and tips

- Start with small changes (± 2 - 3 dB).
- For dialogue:
 - Slightly reduce **Low** to remove rumble
 - Gently boost **Mid** for clarity
- For music:
 - Adjust **Low** for weight
 - Use **High** for brightness and air
- If audio becomes harsh, reduce **High** before adjusting **Mid**.

Limitations

- No precise frequency control
- No Q (bandwidth) adjustment
- Grouped keyframes only



Applies a fully parametric equalizer with five adjustable sections, allowing precise control over specific frequency ranges.

Unlike graphic [equalizers](#), this filter lets you choose which frequencies are affected, how much they are boosted or cut, and how wide each adjustment is.



This filter operates in the frequency domain.

Structure overview

The filter is divided into five EQ sections:

1. **Low Shelf** – broad control of low frequencies
2. **Band 1** – fully parametric mid band
3. **Band 2** – fully parametric mid band
4. **Band 3** – fully parametric mid band
5. **High Shelf** – broad control of high frequencies

Each section affects a different part of the spectrum and serves a distinct purpose.

Parameters

Low Shelf

Affects all frequencies below the selected cutoff frequency.

Frequency (20.0 - 20 000.0 Hz)

Sets the point below which frequencies are boosted or attenuated.

- Lower values affect deep bass
- Higher values affect upper bass and low mids

Examples: rumble, bass weight, low-end warmth

Gain (-30.0 to +30.0 dB)

Controls how much the low frequencies are boosted or cut.

- Positive values add bass
- Negative values reduce bass

Slope (0.0 - 1.0)

Controls how gradually the shelf transitions.

- **Low values** – gentle, smooth transition
- **High values** – steeper, more pronounced shelf

Band 1 / Band 2 / Band 3 (Parametric bands)

Band 1, **Band 2**, and **Band 3** are three independent parametric mid-frequency filters.

Having multiple mid bands allows you to:

- Remove several problematic frequencies at once
- Boost one frequency area while cutting another
- Shape voice or music using multiple focused adjustments
- Avoid overusing shelves for mid-range corrections
- Target different frequency areas, or
- Overlap intentionally to shape complex tonal curves

For example:

- **Band 1:** reduce low-mid muddiness
- **Band 2:** enhance speech intelligibility
- **Band 3:** tame harshness or resonance

Each band contributes one “piece” to the final EQ curve.

These are fully adjustable midrange bands, typically used for detailed tone shaping. Each band has the same controls.

Frequency (20.0 - 20 000.0 Hz)

Selects the **center frequency** affected by the band.

- Low values: warmth and body

- Mid values: clarity and presence
- High values: sharpness and detail

Gain (-30.0 to +30.0 dB)

Controls how much the selected frequency range is boosted or attenuated.

- Boosting emphasizes that frequency
- Cutting reduces it

Bandwidth (0.0 - 4.0 octaves)

Controls **how wide** the affected frequency range is.

- **Low values** – narrow band (precise, surgical adjustment)
- **High values** – wide band (broad tonal shaping)

Audible note:

Narrow bandwidth with high gain can introduce audible resonance or ringing.

High Shelf

Affects all frequencies above the selected cutoff frequency.

Frequency (20.0 - 20 000.0 Hz)

Sets the point above which frequencies are boosted or attenuated.

- Lower values affect upper mids
- Higher values affect treble and air

Gain (-30.0 to +30.0 dB)

Controls brightness and high-frequency energy.

- Positive values add clarity and air
- Negative values soften harsh or hissy sounds

Slope (0.0 - 1.0)

Controls how sharply the shelf transitions.

- Low values sound smoother
- High values sound brighter but more aggressive

Keyframes (grouped)

All parameters are keyframed together.

- Individual bands cannot be animated independently
- Each keyframe stores the entire EQ curve


This allows:

- Gradual tonal changes
- Scene-dependent EQ
- Smooth transitions between different EQ settings

Channel selection (toggle buttons)

Determines **which audio channels** the equalizer is applied to.

Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**.

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Important note:

Independent equalization per channel requires adding separate instances of the filter.

Auditory characteristics

- Precise tonal shaping
- Can subtly correct issues or drastically reshape sound
- Narrow boosts may introduce resonance
- Large boosts can cause harshness or listener fatigue
- Broad cuts often sound more natural than boosts

Recommended use cases

- Cleaning up dialogue (removing muddiness or harshness)
- Removing resonant frequencies
- Enhancing speech intelligibility
- Detailed music tone shaping
- Channel-specific EQ in surround mixes
- Corrective audio work where precision matters

Usage notes and tips

- Start with cuts before boosts.


- Use narrow bandwidth to remove problematic frequencies.
- Use wide bandwidth for natural tonal shaping.
- For dialogue:
 - Cut low frequencies with the Low Shelf to remove rumble
 - Use mid bands around 1–3 kHz to improve intelligibility
- For brightness, prefer the **High Shelf** over narrow high boosts.

Limitations

- Grouped keyframes only
- Extreme settings can cause distortion
- Requires careful listening and iteration

Reduces the level of audio signals that fall **below a defined threshold**, increasing the difference between quiet and loud parts.

Unlike a compressor, which reduces dynamic range, an expander **increases dynamic range** by attenuating low-level material.



This filter operates in the **amplitude domain** and is evaluated continuously over time.

None of its parameters are keyframeable.

What this filter does

- Low-level signals are reduced further
- Loud signals remain largely unchanged
- Background noise, room tone, or bleed can be suppressed
- Transients and intentional sounds are preserved

The **Expander** is often used as a gentle noise-reduction or cleanup tool, not as a hard gate.

Parameters

RMS → Peak (0.0 - 100.0%)

Controls how the expander detects signal level.

- **0.0% (RMS)**
Detection is based on average signal energy.
Smoother, more natural behavior.

- **100.0% (Peak)**
Detection reacts to instantaneous peaks.
Faster, more aggressive response.
- **Intermediate values**
Blend between **RMS** and **Peak** detection.

Practical note:

RMS detection is usually preferred for dialogue and music.

Attack (2 - 400 ms)

Controls how quickly attenuation begins **after** the signal drops below the threshold.

- **Short attack**
Rapid suppression of low-level sounds.
- **Long attack**
More gradual reduction, preserving natural decay.

Release (2 - 800 ms)

Controls how quickly attenuation stops **after** the signal rises above the threshold.

- **Short release**
Fast recovery; can sound abrupt.
- **Long release**
Smoother transitions; more natural sound.

Threshold (-30 to 0 dB)

Defines the level below which expansion occurs.

- **Lower values**
Only very quiet sounds are affected.
- **Higher values**
More of the signal is treated as “low level”.

Ratio (1:1 to 1:20)

Controls how strongly low-level signals are attenuated once they fall below the threshold.

- **1:1**
No expansion.
- **1:2 to 1:4**
Gentle expansion; subtle noise reduction.

- **1:8 to 1:20**
Strong expansion; approaches gate-like behavior.

Higher ratios result in more aggressive suppression.

Knee radius (1 - 10 dB)

Controls how smoothly the **Expander** transitions around the threshold.

- **Low values**
Hard knee: abrupt onset of attenuation.
- **High values**
Soft knee: gradual, less noticeable transition.

Soft knees generally sound more natural.

Attenuation (0.0 - 24.0 dB)

Defines the maximum amount of gain reduction applied to signals below the threshold.

- **Low values**
Subtle reduction of quiet material.
- **High values**
Strong suppression of background noise or bleed.

This sets a ceiling on how quiet low-level signals can become.

Channel selection (toggle buttons)

Determines which audio channels the **Expander** is applied to.
Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent processing per channel requires adding separate instances of the filter.

Auditory characteristics

- Reduced background noise between sounds
- Clearer separation between silence and signal
- Can sound cleaner and more focused
- Overuse can make audio feel unnatural or “choppy”

Recommended use cases

- Reducing room noise in dialogue
- Cleaning up microphone bleed
- Suppressing low-level background sounds
- Improving clarity without fully muting silence
- Pre-processing before compression

Usage notes and tips

- Start with low ratios and moderate attenuation.
- Use soft knees and longer release times for natural results.
- Avoid extreme settings unless a gate-like effect is desired.
- Monitor quiet passages carefully to avoid unnatural dropouts.

Relation to the Compressor

The **Expander** is conceptually the inverse of the **Compressor**:

- **Compressor**
Reduces levels *above* a threshold.
- **Expander**
Reduces levels *below* a threshold.

They are often used together:

1. **Expander** to reduce background noise
2. **Compressor** to control loudness and dynamics


This combination improves clarity while maintaining consistent output levels.

Limitations

- Not a true noise remover
- No lookahead or sidechain
- No keyframe support
- Extreme settings can sound unnatural

The Fade In Audio Filter has two parameters: duration and type.

The type parameter was added in Shotcut 25.01.03



The following waveform images illustrate each fade of a sine wave over 1 second.


Natural: A log based curve (linear in dB). This type is considered the most natural in audio mixing and is sometimes called “Logarithmic”.




S-curve: Slow at the start, fast in the middle, slow at the end



Fast-Slow: Fast at the start, then ends slow



Slow-Fast: Slow at the start, then ends fast




For more complex fade types, consider using the [Gain / Volume Audio Filter](#) with custom [Keyframe Types and Easing](#).



The Fade Out Audio Filter has two parameters: duration and type.

The type parameter was added in Shotcut 25.01.03



The following waveform images illustrate each fade of a sine wave over 1 second.

Natural: A log based curve (linear in dB). This type is considered the most natural in audio mixing and is sometimes called “Logarithmic”.




S-curve: Slow at the start, fast in the middle, slow at the end



Fast-Slow: Fast at the start, then ends slow




Slow-Fast: Slow at the start, then ends fast



For more complex fade types, consider using the [Gain / Volume Audio Filter](#) with custom [Keyframe Types and Easing](#).

Adjusts the **signal level** of audio by applying a uniform gain change. This filter increases or decreases loudness without altering tone, dynamics, or timing.



Gain / Volume is a level-control filter. It is often used as a building block before or after other audio processing.

Parameters

Level (-70 to +24 dB)

Controls the amount of gain applied to the selected audio channels.

- **Negative values** reduce level (attenuation).
Useful for preventing clipping or lowering overly loud material.
- **0 dB**
No change to the original level.
- **Positive values** increase level (amplification).
Useful for quiet recordings or restoring level after processing.

Behavior notes:

- The control is expressed in decibels (dB), representing gain, not perceived loudness directly.
- Large positive values can cause clipping if downstream headroom is insufficient.

Keyframes


The **Level** parameter is keyframeable. This enables:

- Smooth fades in or out
- Automated level rides

- Scene-by-scene loudness adjustments

Channel selection (toggle buttons)

Determines which audio channels the gain is applied to. Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent gain per channel requires adding separate instances of the filter.

Auditory characteristics

- Uniform increase or decrease in loudness
- No change in tone, frequency balance, or dynamics
- Can expose noise when boosting quiet material
- Can cause distortion if over-driven

Recommended use cases

- Normalizing clip levels by ear
- Creating fades and level automation
- Compensating level changes introduced by other filters
- Balancing dialogue, music, or effects
- Pre-gain before compression or limiting
- Channel-specific level adjustment in surround projects

Usage notes and tips


- Prefer small adjustments (± 1 –6 dB) for fine control.
- If boosting significantly, consider adding a **Limiter** downstream.
- Use negative gain before dynamics processing to preserve headroom.
- For consistent loudness across clips, combine with meters and listening checks.

Limitations

- No automatic loudness normalization
- No peak protection or limiting
- Can clip if pushed too far
- Does not address dynamic range issues by itself

Attenuates low frequencies below a selected cutoff, allowing higher frequencies to pass through.

This filter is commonly used to remove rumble, handling noise, or unwanted low-frequency energy while preserving clarity.



High Pass is a **frequency-domain filter** evaluated continuously over time.

Parameters

Cutoff frequency (5 - 21 600 Hz)

Defines the frequency below which audio is progressively attenuated.

- **Very low values (≈ 5 - 40 Hz)**
Remove subsonic rumble without affecting audible content.
- **Low values (≈ 60 - 120 Hz)**
Clean up low-end noise while preserving most voice and music.
- **Mid values (≈ 150 - 300 Hz)**
Thin out sound deliberately; remove bass content.
- **High values (> 500 Hz)**
Strong filtering effect; useful for special effects.

Note:

Raising the cutoff removes more bass.

Rolloff rate (1 - 10)

Controls how steeply frequencies below the cutoff are reduced.

- **Low values (1 - 3)**
Gentle slope. Bass fades out smoothly.
- **Mid values (4 - 7)**
Clear cutoff with controlled transition.
- **High values (8 - 10)**
Aggressive filtering. Low frequencies are strongly suppressed.

Behavior note:

- Higher rolloff values can sound unnatural if set too high.

Dry → Wet (0.0 - 100.0%)

Controls the mix between the original signal and the filtered signal.

- **0.0%**
Original audio only.
- **50.0%**
Balanced mix of filtered and unfiltered audio.
- **100.0%**
Fully filtered audio.

This allows subtle filtering without fully removing low frequencies.

Channel selection (toggle buttons)

Determines which audio channels the filter is applied to.
Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent filtering per channel requires adding separate instances of the filter.

Auditory characteristics

- Reduced bass and low-frequency energy
- Improved clarity for dialogue
- Cleaner mixes with less rumble
- Excessive settings can make audio sound thin

Recommended use cases

- Dialogue cleanup
- Removing wind or handling noise
- Reducing low-frequency hum
- Preventing bass buildup in mixes
- Preparing audio before compression

Example settings

- **Dialogue cleanup**
 - Cutoff: 80 - 120 Hz
 - Rolloff: 3 - 5
 - Dry → Wet: 100%
- **Subtle rumble reduction**
 - Cutoff: 40 - 60 Hz
 - Rolloff: 2 - 3
 - Dry → Wet: 50 - 100%
- **Special effect (thin / radio-like sound)**
 - Cutoff: 300 - 600 Hz
 - Rolloff: 6 - 8
 - Dry → Wet: 100%

Usage notes and tips

- Start with the lowest effective cutoff.
- Increase rolloff only if low frequencies remain audible.
- For natural results, avoid extreme cutoff values.
- Always evaluate changes during playback.


Limitations

- Not a substitute for detailed equalization
- Cannot selectively remove specific bass notes
- Extreme settings can reduce audio naturalness



Inverts the **polarity** of the audio signal.

This means the waveform is flipped vertically: positive values become negative, and negative values become positive.



Invert is a **phase / polarity correction filter**, not an audible effect by itself.

Parameters

This filter has no controls.

There is no channel selection; the operation applies to the entire audio clip. It inverts signal polarity only and does not flip or reverse the audio playback.

What this filter does

- The audio waveform is multiplied by -1
- Loudness, frequency content, and timing remain unchanged
- Only the polarity of the signal is reversed

This is also commonly referred to as phase inversion or polarity flip.

Why you usually hear no difference

When listening to a single audio track:

- Polarity inversion is not perceptible
- Human hearing cannot detect absolute waveform direction
- Music, dialogue, and effects sound the same

This is expected behavior and does not indicate that the filter is broken.

When this filter matters

Invert becomes audible only in relation to other audio signals, for example:

Phase cancellation issues

If two similar signals are combined (e.g. two microphones recording the same source):

- If they are out of phase, parts of the sound may cancel out
- Applying Invert to one signal can restore proper alignment

Multi-microphone recordings

- Overhead + close mics
- Stereo recordings with wiring issues
- Imported audio with incorrect polarity

Layered or duplicated tracks

- Music stems
- Sound effects layered together
- Parallel processing chains

In these cases, Invert can dramatically change the result.

Auditory characteristics

- No audible change on isolated tracks
- Can significantly alter combined signals
- Can restore bass or clarity lost due to phase cancellation
- Can also introduce cancellation if misused

Recommended use cases

- Fixing phase cancellation between tracks
- Correcting inverted polarity recordings
- Testing polarity relationships
- Aligning layered audio sources

Usage notes and tips

- Use Invert only when comparing two or more signals
- Toggle the filter on/off while listening to combined tracks
- Pay attention to bass and low frequencies; phase issues are most noticeable there
- If Invert makes the sound thinner, revert it


Limitations

- No effect on timing or frequency balance
- No visual waveform display

- No automatic phase detection
- Requires contextual listening to be meaningful

Restricts audio level so it **does not exceed a defined maximum**, preventing clipping and controlling peaks.

A limiter is a **dynamics** processor similar to [Compressor](#), but designed to act very strongly once the limit is reached.



Limiter is a peak-control filter evaluated continuously over time.

Parameters

Input gain (-20.0 to +20.0 dB)

Adjusts the level **before** limiting is applied.

- **Negative values**
Reduce input level, decreasing how often the limiter engages.
- **0 dB**
No change to the incoming signal.
- **Positive values**
Increase input level, driving the signal into the limiter more aggressively.

Important concept:

Input gain determines *how hard the signal hits the limiter*, not the final output level.

Limit (-20.0 to 0 dB)

Sets the maximum allowed output level.

- **0 dB**
Allows peaks up to full scale (digital maximum).
- **Negative values**
Lower the ceiling, leaving headroom and reducing peak loudness.

Once the signal reaches this level, the limiter reduces gain to prevent further increase.

Release (0.01 to 2.00 s)

Controls how quickly gain reduction is released **after** the signal falls below the limit.

- **Short release (≈0.01 - 0.10 s)**
Fast recovery; can sound aggressive or distorted on sustained material.
- **Medium release (≈0.10 - 0.50 s)**
Balanced behavior for most content.
- **Long release (≈0.50 - 2.00 s)**
Smooth recovery; preserves natural dynamics but may reduce overall loudness.

Channel selection

Determines which audio channels the limiter is applied to. Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent *limiting* per channel requires adding separate instances of the filter.

Gain Reduction meter (visual indicator)

A visual-only meter labeled **Gain Reduction** (-24 dB to 0 dB) displays how much level is being reduced by the limiter.

- The meter moves only when the signal exceeds the **Limit**.

- It shows reduction amount, not output level.
- It is informational only.

Auditory characteristics

- Prevents clipping and digital distortion
- Reduces dynamic peaks
- Can increase perceived loudness
- Excessive limiting can sound flat or distorted

Recommended use cases

- Preventing clipping on dialogue, music, or effects
- Final level control before export
- Protecting against unexpected peaks
- Increasing loudness safely
- Channel-specific peak control in surround mixes

Example settings

- **Peak protection (transparent)**

- Input gain: 0 dB
- Limit: -1 dB
- Release: 0.20-0.40 s

- **Loudness boost (controlled)**

- Input gain: +4 to +8 dB
- Limit: -1 dB
- Release: 0.10-0.30 s

- **Dialogue safety limiter**

- Input gain: +2 to +4 dB
- Limit: -2 dB
- Release: 0.30 - 0.60 s

Usage notes and tips


- Set **Limit** first, then adjust **Input gain**.
- Watch the **Gain Reduction** meter; constant heavy reduction indicates over-limiting.
- Use longer release times for natural material like dialogue.
- For best results, place the limiter near the end of the audio filter chain.

Limitations

- No look-ahead control
- No soft/hard knee adjustment
- Can introduce distortion if pushed too hard
- Not a substitute for careful gain staging

Attenuates **high frequencies above a selected cutoff**, allowing lower frequencies to pass through.

This filter is commonly used to soften harsh sounds, remove high-frequency noise, or deliberately reduce clarity for stylistic effects.



Low Pass is a **frequency-domain filter** evaluated continuously over time.

Parameters

Cutoff frequency (5 - 21 600 Hz)

Defines the frequency above which audio is progressively attenuated.

- **High values (≈8 000 - 20 000 Hz)**
Very subtle effect; only extreme highs are reduced.
- **Mid values (≈3 000 - 8 000 Hz)**
Noticeable softening of clarity and brightness.
- **Low values (≈500 - 3 000 Hz)**
Strong filtering; audio becomes muffled or “covered”.
- **Very low values (<500 Hz)**
Extreme effect; only bass and low mids remain.

Note:

Lowering the cutoff removes more treble.

Rolloff rate (1 - 10)

Controls how **steeply** frequencies above the cutoff are reduced.

- **Low values (1 - 3)**
Gentle slope; natural-sounding attenuation.
- **Mid values (4 - 7)**
Clear cutoff with controlled transition.
- **High values (8 - 10)**
Aggressive filtering; high frequencies are sharply suppressed.

Behavior note:

- High rolloff combined with low cutoff can sound unnatural or “boxed”.

Dry → Wet (0.0 - 100.0%)

Controls the mix between the original signal and the filtered signal.

- **0.0%**
Original audio only.
- **50.0%**
Partial filtering with retained clarity.
- **100.0%**
Fully filtered audio.

This allows subtle softening without fully removing high frequencies.

Channel selection (toggle buttons)

Determines which audio channels the filter is applied to.
Buttons are toggle switches.



Available buttons depend on the project’s audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent filtering per channel requires adding separate instances of the filter.

Relation to the Low & High Pass filters

Low Pass and High Pass are complementary filters:

- **High Pass** removes low frequencies and preserves highs.
- **Low Pass** removes high frequencies and preserves lows.

Used together, they can:

- Isolate a specific frequency band
- Create band-limited or “telephone” effects
- Clean both rumble and hiss from dialogue

Auditory characteristics

- Reduced brightness and detail
- Softer, darker sound
- Suppressed hiss or sharp transients
- Extreme settings can remove intelligibility

Recommended use cases

- Removing hiss or high-frequency noise
- Softening harsh dialogue or effects
- Creating distance or obstruction effects
- Stylized audio (radio, underwater, behind walls)
- Preparing audio before further processing

Example settings

- **Subtle softening**
 - Cutoff: 10 - 14 kHz
 - Rolloff: 2 - 4
 - Dry → Wet: 100%
- **Dialogue muffling / obstruction effect**
 - Cutoff: 2 - 4 kHz
 - Rolloff: 4 - 6
 - Dry → Wet: 100%
- **Band-limited effect (with High Pass)**
 - Low Pass cutoff: 3 - 4 kHz
 - High Pass cutoff: 300 - 500 Hz

Usage notes and tips

- Start with high cutoff values and lower gradually.
- Avoid extreme rolloff unless a strong effect is desired.
- Combine with **Dry** → **Wet** for subtle results.
- Always evaluate during playback.


Limitations

- Cannot target specific high frequencies selectively
- No resonance or Q control
- Extreme settings reduce clarity significantly



Processes a left/right (L/R) stereo signal internally using a mid/side (M/S) model to control perceived stereo width. This filter redistributes energy between the center (mid) and the sides without changing overall timing or pitch.

This filter is typically used along with a special, **dual microphone recording technique**.



Mid-Side Matrix is a **stereo spatial processing filter**. It operates only on stereo material.

Parameters

Width (0 - 100%)

Controls the balance between the mid (center) and side (stereo difference) components.

- **0%**
Side information is removed.
The result is effectively mono (center only).
- **50% (default)**
Neutral setting.
Stereo width is preserved as in the original signal.
- **>50%**
Side information is amplified.
Stereo image becomes wider.
- **100%**
Maximum widening.
Can sound exaggerated or unstable on some material.

Note:

Width changes *how wide the sound feels*, not how loud it is.

Dry → Wet (0.0 - 100.0%)

Controls the mix between the original stereo signal and the mid-side processed signal.

- **0.0%**
Original audio only.
- **50.0%**
Blend of original and processed audio.
- **100.0%**
Fully Mid-Side processed audio.

This allows subtle widening or narrowing without fully replacing the original stereo image.

Auditory characteristics

- Changes perceived stereo width
- Center content (vocals, dialogue) remains stable
- Side content (ambience, reverb, stereo effects) is emphasized or reduced
- Excessive widening may cause phase issues or poor mono compatibility

Recommended use cases

- Widening music or ambient tracks
- Narrowing overly wide stereo recordings
- Improving focus of dialogue by reducing side content
- Creative stereo effects
- Matching stereo width between clips

Example settings

- **Subtle widening (music)**
 - Width: 60 - 70%
 - Dry → Wet: 100%
- **Strong widening (ambient effects)**
 - Width: 80 - 90%
 - Dry → Wet: 100%
- **Dialogue focus / near-mono**
 - Width: 0 - 20%

- Dry → Wet: 100%
- **Gentle adjustment**
- Width: 70%
- Dry → Wet: 40 - 60%

Usage notes and tips


- Always check mono compatibility after widening.
- Small changes in **Width** are often sufficient.
- Use **Dry** → **Wet** for fine control rather than extreme **Width** values.
- Best suited for stereo music and ambience, not centered narration.

Limitations

- Operates only on stereo audio
- Does not create true spatial depth
- Extreme values can introduce phase cancellation
- Not a replacement for surround mixing

Silences audio by setting its level to **absolute zero**.

Unlike the track header mute button, which affects an entire track, this filter operates **at the clip level** and can target individual channels.



Mute applies absolute silence to selected audio channels at the clip level.


Parameters

This filter has **no sliders or numeric controls**.

Channel selection

Determines which audio channels are muted.

Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Only the selected channels are silenced; unselected channels remain unaffected.

What this filter does

- Completely mutes selected channels for the duration of the clip
- Leaves timing, sync, and other channels unchanged
- Does not alter audio data permanently
- Applies only where the filter is present

How this differs from trackmute

- **Track header mute**
 - Mutes all clips on the track
 - Affects all channels
 - Cannot target individual clips or channels
- **Mute filter**
 - Applies to a specific clip
 - Can mute one or multiple channels
 - Leaves other channels audible

Auditory characteristics

- Selected channels produce no sound
- No fade or attenuation; muting is absolute
- No effect on remaining channels

Recommended use cases

- Muting a specific clip without affecting the rest of the track
- Removing unwanted channels from imported multi-channel audio
- Silencing LF (Low Frequency) or surround channels selectively
- Creating intentional channel dropouts
- Cleaning up recordings with unused or noisy channels

Usage notes and tips

- Use **Mute** instead of **Gain / Volume** when silence is required.
- Combine with channel selection to surgically remove problem channels.
- For fades or gradual silencing, use **Gain / Volume** instead.


Limitations

- No fades or transitions
- No keyframe support
- No partial attenuation
- Must be combined with other filters for dynamic behavior



This filter uses the [Steve Harris LADSPA Gate plugin](#), whose parameters are based on the [Drawmer DS201](#).

This filter is used to reduce some noise like a hum or echo or to reduce cross-talk between microphones that were targeting different sources.



Key Filter: Low Frequency

Controls the cutoff of the low frequency filter (highpass).

Key Filter: High Frequency

Controls the cutoff of the high frequency filter (lowpass).

Output Key Only

Controls output monitor. When checked, it is the output of the key filters (so you can check what is being gated on). Otherwise, it is the normal, gated output. Disable the filter in Shotcut's list of filters to achieve bypass (or [trim the filter](#) to bound it by a time range).

Threshold

Controls the level at which the gate will open.

Attack

Controls the time the gate will take to open fully.

Hold

Controls the minimum time the gate will stay open for.

Decay


Controls the time the gate will take to close fully.

Range

Controls the difference between the gate's open and closed state.

Adjusts audio loudness **dynamically during playback** to approach a target loudness level, without a prior analysis pass.

Unlike offline or two-pass normalization, this filter reacts in real time and continuously adapts gain as the audio plays.



Normalize: One Pass is a **loudness-based automatic gain control**, not a static normalization tool.

What this filter does

- Measures integrated loudness in real time (LUFS)
- Gradually adjusts gain to approach a target loudness
- Applies changes smoothly over time, not instantaneously
- Does **not** scan the entire clip in advance

Because it works in one pass, results depend on playback order and analysis window.

Parameters

Target Loudness (-50.0 to -10.0 LUFS)

Defines the desired loudness level the filter attempts to reach.

- **Lower values (e.g. -30 to -23 LUFS)**
Quieter overall output, more headroom.
- **Higher values (e.g. -16 to -12 LUFS)**
Louder output, typical for online content.

Note:

This is a *target*, not a guaranteed final value.

Analysis Window (2 to 600 s)

Defines the **time window** over which loudness is measured.

- **Short window (≈2 - 10 s)**
Faster reaction to changes, more noticeable gain movement.
- **Medium window (≈10 - 60 s)**
Balanced behavior for dialogue and mixed content.
- **Long window (≈60 - 600 s)**
Very smooth, slow adaptation; closer to long-term loudness matching.

Important clarification:

The value is expressed in seconds, not milliseconds.

Maximum Gain (0.0 to +30.0 dB)

Sets the maximum amount of gain increase the filter is allowed to apply.

- Prevents excessive boosting of very quiet material
- Limits noise amplification

Minimum Gain (-30.0 to 0.0 dB)

Sets the maximum amount of attenuation the filter is allowed to apply.

- Prevents excessive reduction of loud sections
- Limits how quiet the output can become

Maximum Rate (0.5 to 9.0 dB/s)

Limits how fast gain is allowed to change over time.

- **Low values**
Very smooth, gradual loudness changes.
- **High values**
Faster adaptation, more audible level movement.

Note:

This controls how *quickly* the filter reacts, not how *much* it reacts.

Reset on discontinuity

When enabled, the loudness analysis resets at detected discontinuities such as:

- Cuts
- Seeks
- Playback jumps

This prevents loudness history from one section influencing another.

Visual indicators

Input Loudness (-50 to 0 LUFS)

Displays the measured incoming loudness over time.

- Updates only during playback
- Displayed for monitoring purposes only.

Output Gain (0.0 to 30.0 dB)

Displays the current gain being applied by the filter.

- Shows how aggressively the filter is acting
- Useful for diagnosing over-correction

Reset

A visual, radio-style indicator that lights up when the internal loudness state is reset.

- Informational only
- Indicates analysis restart, not a user action

Auditory characteristics

- Gradual loudness leveling
- No sudden jumps when properly configured
- Can pump or breathe if reaction is too fast
- Quiet sections may still sound noisy if boosted heavily

Recommended use cases

- Long-form dialogue with varying levels
- Live recordings with inconsistent loudness

- Content where a single loudness target is desired
- Preview normalization without offline analysis

What this filter is not

- Not true offline normalization
- Not guaranteed to hit an exact LUFS value
- Not a **limiter** (peaks can still clip)
- Not ideal for short clips with rapid cuts

Usage notes and tips

- Set **Target Loudness** first, then constrain behavior with gain limits.
- Use a longer **Analysis Window** for natural results.
- Keep **Maximum Rate** low to avoid audible pumping.
- Combine with a **Limiter** downstream for peak safety.
- Always evaluate during real-time playback.


Limitations

- Results depend on playback order
- No lookahead
- Cannot normalize silence meaningfully



Adjusts audio loudness by applying a **single, fixed gain** calculated from a prior analysis pass.

Unlike **Normalize: One Pass**, this filter does not adapt during playback; it applies a constant gain determined in advance.



Normalize: Two Pass is an **offline loudness normalization filter**.

What this filter does

- Analyzes the entire clip to measure its loudness
- Computes a single gain value needed to reach a target loudness
- Applies that gain uniformly across the whole clip during playback and export
- Does not change gain over time

This produces predictable, consistent results.

Parameters

Target Loudness (-50.0 to -10.0 LUFS)

Defines the desired integrated loudness for the clip.


- **Lower values (e.g. -30 to -23 LUFS)**
Quieter output, more headroom.
- **Higher values (e.g. -16 to -12 LUFS)**
Louder output, common for online content.

This value is used only to compute the normalization gain.

Analyze

Runs the analysis pass over the entire clip.


- Measures the clip's integrated loudness
- Calculates the gain required to reach the target
- Must be run before normalization can occur



Changing the clip content or **Target Loudness** requires re-running the analysis.

Analysis results

Displayed after analysis completes.



Detected Loudness (LUFS)

Shows the measured integrated loudness of the clip before normalization.

Normalization Gain (dB)

Shows the fixed gain that will be applied to the clip.

- Positive values increase loudness

- Negative values reduce loudness

This gain is applied uniformly to the entire clip.

How normalization is applied

- **First pass:** analyzes the clip and computes a gain value
- **Second pass:** applies that gain during playback and export

The filter applies a flat gain adjustment; it is not dynamic.

Auditory characteristics

- Consistent loudness across the entire clip
- No pumping or breathing
- Dynamics within the clip remain unchanged
- Quiet and loud sections keep their relative balance

Recommended use cases

- Normalizing dialogue clips to a consistent loudness
- Matching loudness between different recordings
- Preparing audio for delivery standards
- Batch normalization of multiple clips

What this filter is not

- Not dynamic loudness control
- Not a **Compressor** or **Limiter**
- Does not protect against clipping
- Does not adapt to changes within the clip

Usage notes and tips

- Always click **Analyze** after changing settings or clip content.
- Combine with a **Limiter** if peak protection is required.
- For clips with large level variation, consider Normalize: One Pass instead.
- Use the displayed gain value to understand how much adjustment is being applied.


Limitations

- Gain is fixed for the entire clip
- Cannot react to internal loudness variations
- Analysis must be repeated after edits



Removes a **very narrow range of frequencies** centered around a selected frequency while leaving the rest of the spectrum largely unchanged.

This filter is designed to eliminate **specific, unwanted tones** such as hum, buzz, or ringing.



Notch is a **frequency-domain filter** evaluated continuously over time.

Parameters

Center frequency (5 - 21 600 Hz)

Sets the **exact frequency** that will be removed.

- **Low values (≈50 - 120 Hz)**
Useful for removing electrical hum or low-frequency rumble.
- **Mid values (≈300 - 3 000 Hz)**
Target problematic resonances in voices or instruments.
- **High values (>3 000 Hz)**
Remove high-pitched whines or ringing.

Note:

This is the frequency you want to suppress.

Bandwidth (5 - 21 600 Hz)

Controls **how wide** the removed frequency range is around the center frequency.

- **Low values**
Very narrow notch; removes a specific tone.
- **High values**
Wider removal; affects more surrounding frequencies.

Important note:

Although expressed in Hz, this controls width, not position.

Rolloff rate (1 - 10)

Controls how steeply frequencies transition into and out of the notch.

- **Low values (1 - 3)**
Gentle slopes; smoother, less obvious filtering.
- **High values (8 - 10)**
Steep slopes; precise removal with minimal side effects.

Higher rolloff values make the notch more focused.

Dry → Wet (0.0 - 100.0%)

Controls the mix between the original signal and the notched signal.

- **0.0%**
Original audio only.
- **50.0%**
Partial removal of the target frequency.
- **100.0%**
Full notch effect.

Allows subtle correction when full removal sounds unnatural.

Keyframes

All four sliders are keyframeable.

Keyframes allow the notch filter to change over time, enabling:

- Sweeping the notch frequency to track a moving hum or whistle
- Gradually reducing or restoring a problematic tone
- Applying the notch only during specific sections of a clip
- Creating automated resonance suppression


Behavior note:

Because the filter operates in the frequency domain, keyframed changes are

best evaluated during playback rather than on a single frame.

Channel selection

Determines which audio channels the filter is applied to. Buttons are toggle switches.



Available buttons depend on the project's audio channel configuration: **Settings > Audio Channels**

- **L** — Front Left
- **R** — Front Right
- **C** — Center
- **LF** — Low-Frequency Effects (LFE)
- **Ls** — Surround Left
- **Rs** — Surround Right

Note:

Independent filtering per channel requires adding separate instances of the filter.

Auditory characteristics

- Removes a specific tone or resonance
- Minimal impact on overall sound when set narrowly
- Wide settings can make audio sound hollow
- Overuse may reduce naturalness

Recommended use cases

- Removing electrical hum (50/60 Hz)
- Eliminating microphone feedback
- Reducing ringing or resonant tones
- Cleaning up dialogue recordings
- Fixing tonal artifacts without affecting the full mix

Example settings

- **50 Hz hum removal**
 - Center frequency: 50 Hz
 - Bandwidth: narrow
 - Rolloff rate: high
 - Dry → Wet: 100%

- **Voice resonance reduction**
 - Center frequency: 500 - 1 000 Hz
 - Bandwidth: narrow to medium
 - Rolloff rate: medium
 - Dry → Wet: 50 - 100%
- **High-pitched whine**
 - Center frequency: 4 - 8 kHz
 - Bandwidth: narrow
 - Rolloff rate: high

Usage notes and tips


- Sweep the **Center frequency** while listening to locate the problem tone.
- Use the narrowest **Bandwidth** that solves the issue.
- Prefer high **Rolloff rate** for surgical removal.
- Use **Dry** → **Wet** to soften the effect if needed.

Limitations

- Not suitable for broad tonal shaping
- Cannot remove multiple tones at once
- Extreme settings can cause audible artifacts

Adjusts the **stereo position** of audio by distributing signal level between the left and right channels.

This filter is used to place sound toward one side of the stereo field or to animate movement across it.



Pan is a **stereo level-distribution filter**. It does not change tone, timing, or pitch.

Parameters

Channel

Selects which channel is treated as the source for the panning operation.

- **Left**
Uses the left channel as the source signal.
- **Right**
Uses the right channel as the source signal.

The selected channel is redistributed between left and right outputs according to the pan slider.

Important behavior note:

The Pan filter does not mix both channels together. It takes the selected channel and spreads it across the stereo field.

Left → Right (0.00 - 1.00)

Controls where the selected channel is placed in the stereo field.

- **0.00**
Output is fully left.

- **0.50**
Centered (equal level in left and right).
- **1.00**
Output is fully right.

Intermediate values place the sound proportionally between left and right.

Keyframes

The **Left** → **Right** slider is keyframeable.

This allows:

- Smooth panning movement across the stereo field
- Sound traveling from left to right (or the reverse)
- Position changes over time within a clip

Keyframed panning is evaluated continuously during playback.

Auditory characteristics

- Changes perceived position of sound
- Loudness remains similar, but distribution changes
- Extreme positions isolate sound to one speaker
- Center position places sound equally in both speakers

Recommended use cases

- Positioning dialogue or effects in stereo space
- Correcting recordings where sound is biased to one side
- Creating motion effects (e.g. passing objects)
- Separating overlapping sounds in a mix
- Creative stereo placement

Example usages

- **Center a left-only recording**
 - **Channel:** Left
 - **Left** → **Right:** 0.50
- **Move a sound from left to right**
 - **Channel:** Left (or Right, depending on source)
 - Animate **Left** → **Right** from 0.00 to 1.00
- **Place an effect slightly off-center**
 - **Channel:** Left

- **Left** → **Right**: 0.35 or 0.65

Usage notes and tips

- Choose the channel that actually contains the sound you want to pan.
- For subtle placement, avoid extreme values.
- For stereo recordings where both channels matter, consider other mixing tools instead of Pan.
- Always monitor on stereo speakers or headphones for accurate placement.


Limitations

- Operates only on one selected channel at a time
- Does not preserve stereo width when applied to a single selected channel
- Not suitable for surround positioning (Front/Rear)

See also: [Balance Audio Filter](#)



The **Pitch** filter lets you change how high or low the sound is with respect to frequency (not loudness).




One may use this effect to disguise a voice or create a silly sounding voice. It is very useful to use in conjunction with the **Speed** option in **Properties**, which inherently changes the pitch. Thus, this filter can be used to correct the pitch. You can simply copy the speed value from Properties to the **Speed Compensation** parameter in the Pitch filter!

The **Octave Shift** parameter is a different way to represent the change in musical form. There are 6 whole steps in music between each octave. Thus, each whole step corresponds to increments or decrements of 0.1666. For example, to change the pitch from A to B, set the Octave shift to 0.166666. There are twelve half-steps (semitones) in an octave. For example, to change the pitch from A to Bb, set the Octave Shift to 0.083333 (1/12).

Simulates the way sound reflects and decays in a physical space, adding a sense of **depth, distance, and environment**.

Reverb recreates both the **early reflections** (first bounces off nearby surfaces) and the reverb tail (the longer decay that follows).



Reverb is a **time-based spatial effect**. All parameters are keyframeable, allowing the perceived space to change over time.

Parameters

Room size (1 - 300 m)

Defines the virtual size of the space in which the sound exists.

- **Small values ($\approx 1 - 10$ m)**
Very small rooms; tight, close reflections.
- **Medium values ($\approx 10 - 50$ m)**
Typical indoor spaces such as rooms or halls.
- **Large values (> 50 m)**
Large halls or cavernous spaces.

Note:

Larger room size = reflections feel farther away and more spacious.

Reverb time (0.1 - 30.0 s)

Controls how long the reverb takes to decay after the sound stops.

- **Short times ($\approx 0.1 - 0.5$ s)**
Subtle ambience; barely noticeable.
- **Medium times ($\approx 0.8 - 3$ s)**
Natural room or hall reverberation.
- **Long times (> 5 s)**
Large, echoing spaces; dramatic or stylized effects.

Important distinction:

Room size affects *space perception*; reverb time affects *decay duration*. They are related but not the same.

Damping (0.0 - 100.0%)

Controls how quickly high frequencies decay in the reverb tail.

- **Low values**
Bright, metallic reverb; highs persist.
- **High values**
Darker, warmer reverb; highs fade faster.

Examples:

Higher damping is often preferred for dialogue to avoid harshness.

Input bandwidth (0.0 - 100.0%)

Limits how much of the incoming frequency range is fed into the reverb.

- **Low values**
Narrower frequency range; darker reverb.
- **High values**
Full-range reverb input.

Note:

This controls how “full” or “thin” the reverb sounds.

Dry signal level (-70.0 to 0.0 dB)

Sets the level of the original, unprocessed sound.

- **0.0 dB**
Original signal at full level.

- **Lower values**
Original signal reduced, making reverb more prominent.

Early reflection level (-70.0 to 0.0 dB)

Controls the level of early reflections.

- Early reflections define perceived room size and distance
- Too much can make sound feel close but boxy
- Too little can reduce realism

Examples:

Helpful for placing voices slightly “into” a room without heavy tail.

Tail level (-70.0 to 0.0 dB)

Controls the level of the **reverb tail** (late reflections).

- Higher values increase ambience and spaciousness
- Lower values keep sound clearer and closer

Behavior note:

The tail is where long decay and “wash” come from.

Keyframes

All reverb parameters are keyframeable.

This allows:

- Gradually moving a sound into or out of a space
- Increasing reverb for transitions or flashbacks
- Changing perceived room size within a clip

Reverb changes are best evaluated during playback.

Auditory characteristics

- Adds depth and spatial context
- Makes sounds feel closer or farther away
- Excessive reverb reduces clarity
- Long tails can mask speech and detail

Recommended use cases

- Adding natural ambience to dialogue
- Placing sounds in a believable environment
- Enhancing music or ambient soundscapes
- Creating distance or memory/flashback effects

- Stylized or dramatic spatial effects

Example settings

- **Subtle dialogue ambience**
 - Room size: 10 - 20 m
 - Reverb time: 0.5 - 1.2 s
 - Damping: 60 - 80%
 - Early reflections: -12 to -18 dB
 - Tail level: -20 to -30 dB
- **Large hall / dramatic space**
 - Room size: 50 - 100 m
 - Reverb time: 4 - 8 s
 - Damping: 30 - 50%
 - Tail level: -10 to -15 dB
- **Distant or dreamlike sound**
 - Dry signal level: reduced
 - Tail level: increased
 - Long reverb time

Usage notes and tips


- Start with small amounts; reverb builds quickly.
- For dialogue, keep the dry signal dominant.
- Adjust early reflections before increasing tail level.
- Too much reverb can sound muddy or unnatural.

Limitations

- Not a physical room simulator
- No per-surface control
- Can reduce intelligibility if overused
- Extreme settings may cause ringing or wash

This filter is used to make an interesting stereo sound from a mono source - more interesting than simply duplicating the channel, which is the default when not using this filter. Duplicating a mono channel to stereo still sounds like mono.

Stereo Enhancer was added in version 22.03.30.




The Stereo Enhancer filter is based on the FFmpeg [haas](#) filter.

The Source parameter has four choices.

- Left: The source only has mono sound coming from the left channel.
- Right: The source only has mono sound coming from the right channel.
- Middle (L+R): The source already has some stereo effects that you wish to enhance or you do not know if the mono sound is on the left or the right channel.
- Side (L-R): The source has a split audio track where separate sets of sounds are limited to each of the two channels. For example, the left audio channel has vocals only and the right audio channel has music and sound effects only.

Exchanges the audio content of **two selected channels** within the same clip. This filter is used to correct channel placement errors or reassign content when channels are wired or mapped incorrectly.



Swap Channels is a **channel-routing utility filter**. It does not change loudness, tone, timing, or dynamics.

Parameters

Swap

Selects the first channel to be exchanged.

The available options depend on the project's **Audio Channels** setting.

With

Selects the second channel to be exchanged.

When the filter is active, the audio content of the two selected channels is swapped.

Channel availability

The channels listed depend on the project configuration in **Settings > Audio Channels**.

1 (Mono)

- **Swap:** inactive
- **With:** inactive

With only one channel, there is nothing to exchange.

2 (Stereo)

- **Swap / With:**
 - Left
 - Right

Allows swapping the left and right stereo channels.

4 (Quad / Ambisonics)

- **Swap / With:**
 - Front left
 - Front right
 - Left surround
 - Right surround

Any two channels among the four can be exchanged.

6 (5.1 Surround)

- **Swap / With:**
 - Front left
 - Front right
 - Center
 - Low frequency
 - Left surround
 - Right surround

Provides full flexibility to correct or rearrange channel placement.

What this filter does

- Exchanges audio between two channels
- Preserves original levels, timing, and frequency content
- Operates only while the filter is active
- Does not permanently modify the source audio

Disabling or removing the filter restores the original channel layout.

Auditory characteristics

- Sound appears to move from one speaker to another
- No change in loudness or tone
- No added resonance, echo, or artifacts
- Effect is purely positional

Recommended use cases

- Fixing reversed left/right stereo recordings
- Correcting misrouted surround channels
- Swapping dialogue accidentally recorded in the wrong speaker
- Repairing imported files with incorrect channel order
- Adjusting channel layout before further processing

Example usages

- **Stereo channels reversed**
 - **Swap:** Left
 - **With:** Right
 - **Result:** Correct stereo orientation
- **Dialogue in wrong surround channel**
 - **Swap:** Center
 - **With:** Front left (or right, depending on error)
- **Misplaced bass content**
 - **Swap:** Low frequency
 - **With:** another channel (for correction purposes)

Usage notes and tips

- Use **Swap Channels** when audio content is in the wrong channel, not when it needs blending.
- For duplicating or copying audio, use **Copy Channel** instead.
- For silencing a channel, use **Mute**.
- Place this filter early in the audio filter chain for routing corrections.

Limitations

- No mixing or blending
- No keyframing
- No time-based behavior
- Only two channels can be swapped per filter instance, but you can use more than one of the same filter



Track Auto Fade Audio Filter

This filter address a problem where clicks and pops may be heard in export when switching between clips due to discontinuities in the audio stream. A recommended technique was to add 2-frame **Fade Audio Out** and **Fade Audio In** filters on clips.

Track Auto Fade automates that. This filter can only be placed on a track by clicking the track header (in the left column of the Timeline) to select it.

This filter first appeared in version 23.07.29.



Track Seam Audio Filter

This filter address a problem where clicks and pops may be heard in export when switching between clips due to discontinuities in the audio stream. A recommended technique was to add 2-frame **Fade Audio Out** and **Fade Audio In** filters on clips. **Track Seam** works on a similar level but rather attempts to do a smoothing between outgoing and incoming samples around edit points. This filter can only be placed on a track by clicking the track header (in the left column of the Timeline) to select it.


This filter first appeared in version 23.07.29.

For technical reasons this appears in the **Time** category of the **Filters** chooser, but it is an audio filter.

This first appeared in version 23.09 with the name **Declick** and the name in version 24.01 was changed to **Declick Audio** to make it more obvious.

This filter is intended to remove click, crackle, and pop artifacts in the source audio - not from splitting clips or switching between clips on the timeline. For those situations, see the **Track Auto Fade Audio** and **Track Seam** (recommended) audio filters.

Here are the parameters:




The screenshot shows the parameter interface for the Declick Audio Filter. It features a dark-themed UI with four rows of controls. The first row is for 'Preset', showing a dropdown menu, a '+' button, and a '-' button. The second row is for 'Window', with a blue progress bar, a text field containing '50', the unit 'ms', and a reset button. The third row is for 'Threshold', with a text field containing '2' and a reset button. The fourth row is for 'Burst Fusion', with a blue progress bar, a text field containing '20', the unit '%', and a reset button.

Parameter	Value	Unit
Preset	[Dropdown]	
Window	50	ms
Threshold	2	
Burst Fusion	20	%

- **Window** sets the amount of audio (number of samples in milliseconds) that will be processed at-a-time.
- **Threshold** is the strength in % of impulse noise to be removed. The lower value, the more samples will be detected as impulsive noise.
- **Burst fusion** is a percentage of the window - if any two samples detected as noise are spaced less than this value, any sample between those two samples will be also detected as noise.


This filter is similar to **Time Remap** but can be easier to use.



This filter first appeared in version 23.05.14.


Speed: Forward Only Time Filter

This filter is similar to **Time Remap** but can be easier to use. It does not usually require converting the clip first unlike **Time Remap** and **Speed: Forward & Reverse**.



This filter first appeared in version 23.05.14.

The Time Remap filter is the first filter that allows altering the timing of the underlying source clip. It maps clip time to output time.



See also

[Speed: Forward Only Time Filter](#)

This filter is similar to Time Remap but can be easier to use. It does not usually require converting the clip first unlike Time Remap and Speed: Forward & Reverse. [\[Screenshot Speed: Forward Only filter\]](#) This filter first appeared in version 23.05.14.

[Speed: Forward & Reverse Time Filter](#)

This filter is similar to Time Remap but can be easier to use. [\[Screenshot Speed: Forward Only filter\]](#) This filter first appeared in version 23.05.14.

Limitations

- It can only be added onto clips with NO B-frames. If you try to add it to such a clip, it shows the **Convert to Edit-friendly** dialog.
- A proxy clip can mask the above check for B-frames and then export may be very slow or give bad results. However, it may also be a very convenient workaround for the check and dialog mentioned above if you are patient and confident to troubleshoot and fix problem areas upon export.
- This filter cannot be disabled.
- This filter does not support simple keyframes or filter trimming.
- This filter can not be applied to tracks or the timeline output.
- The audio waveform does not reflect any of the changes.

Usage


TL;DR

Filters > Time > Time Remap > Keyframes


Hold **Ctrl** to drag a keyframe vertical only or **Alt** to drag horizontal only except on macOS these are **⌘** to drag vertical only or **⌥** to drag horizontal only.

More detail

The default setting for the **Time Remap** filter sets the output time equal to the clip time which results in no change to the clip.



The user can interact with the keyframes to change the time mapping. When the slope of the line is up, time is moving forward. When the line is flat, time is frozen (freeze frame). When the slope of the line is down, time is moving backwards (reverse). A steeper slope moves time faster.




When the speed is faster than 1.0, multiple clip frames can map to the same output time.

What the keyframes mean

When a video clip is added to the timeline, there is *always* a mapping between the clip frames and the timeline frames. Under normal circumstances, this mapping is

created internally by Shotcut, it is invisible, and it is outside of the user's control. This is how Shotcut is able to add video clips with many different frame rates to the timeline without having to preprocess all clips to a standardized rate.


Consider what happens when 48fps GoPro footage is added to a 30fps timeline:



The gray area in the above image represents the frame mapping that Shotcut automatically applies so that “one second” of GoPro time equates to “one second” of 30fps timeline time. In this example, it means skipping over approximately every third frame of the GoPro footage.

The key takeaway is that a process always exists to connect input frames to output frames. The **Time Remap** filter allows the user to take manual control of that process and override the frame mapping.

To visualize how this works, imagine that the source clip is now on the vertical axis rather than being horizontal on the timeline. Here is an example, where the gray area represents the keyframes in the **Time Remap** filter:



For each keyframe, imagine that you're standing on top of that dot looking towards the source clip wall. Whatever source frame is directly in front of you is what you

will see at that point in time on the timeline. In essence, the keyframes create a “lookup table” that allows any source frame to be referenced from any timeline frame, enabling time effects such as slow motion, freeze frame, and even reverse.

Since the source clip is on the “vertical axis”, the length of the source clip on the timeline is now arbitrary. The clip’s length does nothing but allocate time for the **Time Remap** filter to do its work. If a video is being slowed down, then the clip will probably need to have its **Duration** extended so that it’s long enough to accommodate the slower video. Similarly, if all that’s needed is a quick freeze frame, then the clip can be trimmed short.

To visualize how reverse works, consider the red arrows in the image. When exporting Frame 6, the export engine will look at the keyframes to determine which source frame should be selected. In this example, Frame 3 of the source is selected. Comparing timeline Frames 5-7, we get:

```
Timeline 5 -> Source 4  
Timeline 6 -> Source 3  
Timeline 7 -> Source 4
```

From the viewer’s point of view, the source clip went backwards from Frames 5 to 6 because the selected source frames went from 4 to 3. But when playing Frames 6 to 7, time will appear to move forward again because the selected source frames went from 3 to 4. And yes, creating a steep line (up or down) that skips over source frames will look as though the source clip is sped up.

Following this logic, we also see that source frame 4 is selected for timeline frames 3, 4, and 5. When the same source frame is selected for multiple timeline frames in a row like this, the visual effect is a freeze frame.

The Alt and Ctrl modifiers when dragging keyframes are extremely useful for scrubbing to find a specific source frame without modifying the keyframe’s placement on the timeline, or for preserving a specific source frame while changing its placement on the timeline.

Parameters

Time

This is the output time of the current frame.

- and + decrement or increment by one frame at a time. You can click and hold these to repeat.

< Sets the speed between the previous keyframe and the current frame by adding a keyframe at the current time position if needed.

> Sets the speed between the the current frame and the next keyframe by adding a keyframe at the current time position if needed.

Enable pitch compensation


Enable this if you want the audio pitch to remain close to the original and not affected by speed. The quality is degraded especially for speeds less than 0.5 or greater than 2.0.

Image mode

The **Image Mode** allows the user to choose between selecting the nearest frame, or blending frames frames.

If 30fps footage is sped up to 2x, it effectively becomes 60fps footage but with half the duration of the original. In this example, the sped-up clip is twice the frame rate of the timeline. This means two frames of source fit into the duration of one timeline frame. The export engine has to choose what to do:

1. Choose a single source frame that is closest to the timestamp of the timeline (nearest mode).
2. Blend multiple source frames into one mash-up frame with averaging (blend mode).



Quick tip: “Blend” mode could be useful for timelapse videos where the action needs to be sped up around 8x or more. Instead of seeing people or stars or whatever flickering between unpredictable positions on the screen, Blend mode would average their movements and create connective trails similar to extreme motion blur. This makes it much easier to track where people are moving and gives a much more relaxed and artsy vibe to the video. This also means not having to do a long-exposure photo sequence with your camera then stitch the photos into a video later to get the same effect... instead, it can be done as a single normal video, which can serve the dual purpose of extracting clips you care about in normal speed, but motion blurring the rest of it for a timelapse.

Here’s why it works:

Nearest mode is the same as the frame drop-or-dupe method that the timeline has used for years to handle frame rate mismatches. If someone puts 60fps video on a

30fps timeline, it drops every other frame of the 60fps video to scale it down to 30fps. It selected the “Nearest” frame to the 30fps cadence.

Extending that concept, if someone takes a 10-second clip and goes to Properties > Speed > 5x so that it becomes a 2-second clip, Shotcut accomplishes the speed-up by dropping (ignoring) four frames then displaying the fifth.

If the Time Remap filter is set to Nearest mode, then the usual dupe-or-drop method described above continues to be used. For the exported frame below, I took a 10-second countdown and did a Time Remap to 2 seconds in Nearest mode (same as Properties > Speed > 5x) then exported:



But here is (almost) the same frame when the Time Remap filter is set to Blend mode:



Instead of ignoring four frames and showing only the fifth, the five frames are now averaged together. This causes a gradient in the circular swipe pattern, because the dark bottom-left swipe was on the screen longer than the upper-left, therefore it went darker when averaged.

The faster the speed-up, the more frames that will be blended together.



These are the locations (folders) of the filter presets that are both custom and default filter presets you create and have saved. To find the preset folder for your operating system, refer to the [FAQ](#) or from the main menu choose **Settings > App Data Directory > Show...**, your file manager opens to the location, and open the `presets` folder.

Preset locations as of Shotcut Version [24.02.29](#)

Filter Name Language: English (United States)

n/a = Not Available

Preset Location	Filter Name	Type
bigsh0t_eq_mask	360: Equirectangular Mask	Video
bigsh0t_eq_to_rect	360: Equirectangular to Rectilinear	Video
bigsh0t_eq_to_stereo	360: Equirectangular to Stereographic	Video
bigsh0t_hemi_to_eq	360: Hemispherical to Equirectangular	Video
bigsh0t_rect_to_eq	360: Rectilinear to Equirectangular	Video
bigsh0t_stabilize_360	360: Stabilize	Video
bigsh0t_stabilize_360	360: Transform	Video
n/a	Alpha Channel: Adjust	Video
n/a	Alpha Channel: View	Video
audioDance	Audio Dance Visualization	Video
audioLevelGraph	Audio Level Visualization	Video
audioLightshow	Audio Light Visualization	Video
audioSpectrum	Audio Spectrum Visualization	Video
audioWaveform	Audio Waveform Visualization	Video
n/a	Blend Mode	Video
boxblur	Blur: Box	Video
blur_exponential	Blur: Exponential	Video
blur_gaussian	Blur: Gaussian	Video
blur_lowpass	Blur: Low Pass	Video
blur_pad	Blur: Pad	Video
brightness	Brightness	Video
choppy	Choppy	Video
avfilter.chromahold	Chroma Hold	Video
frei0r.select0r	Chroma Key: Advanced	Video
frei0r.bluescreen0r	Chroma Hold: Simple	Video
lift_gamma_gain	Color Grading	Video
contrast	Contrast	Video
frei0r.c0rners	Corner Pin	Video
n/a	Crop: Circle	Video
cropRectangle	Crop: Rectangle	Video

Preset Location	Filter Name	Type
crop	Crop: Source	Video
deband	Deband	Video
frei0r.distort0r	Distort	Video
dither	Dither	Video
elastic_scale	Elastic Scale	Video
n/a	Fade In Video	Video
n/a	Fade Out Video	Video
frei0r.defish0r	Fisheye	Video
n/a	Flip	Video
frei0r.glitch0r	Glitch	Video
frei0r.glow	Glow	Video
gpsgraphic	GPS Graphic	Video
gpstext	GPS Text	Video
gradient	Gradient	Video
frei0r.cairoimagegrid	Grid	Video
halftone	Halftone	Video
avfilter.hue	Hue/Lightness/Saturation	Video
n/a	Invert Colors	Video
frei0r.keyspillm0pup	Key Spill: Advanced	Video
n/a	Key Spill: Simple	Video
frei0r.lenscorrection	Lens Correction	Video
frei0r.levels	Levels	Video
n/a	LUT (3D)	Video
n/a	Mask: Apply	Video
maskChromaKey	Mask: Chroma Key	Video
n/a	Mask: Draw (Glaximate)	Video
maskFromFile	Mask: From File	Video
maskSimpleShape	Mask: Simple Shape	Video
n/a	Mirror	Video
frei0r.pixeliz0r	Mosaic	Video
opencv.tracker	Motion Tracker	Video
avfilter.random	Nervous	Video
nosync	No Sync	Video
noise_fast	Noise: Fast	Video
noise_keyframes	Noise: Keyframes	Video
dust	Old Film: Dust	Video
grain	Old Film: Grain	Video
oldfilm	Old Film: Projector	Video
lines	Old Film: Scratches	Video
tcolor	Old Film: Technicolor	Video
brightnessOpacity	Opacity	Video
posterize	Posterize	Video
frei0r.hqdn3d	Reduce Noise: HQDN3D	Video

Preset Location	Filter Name	Type
avfilter.smartblur	Reduce Noise: Smart Blur	Video
vaguedenoiser	Reduce Noise: Wavelet	Video
n/a	Reflect	Video
frei0r.rgbsplit0r	RGB Shift	Video
frei0r.saturat0r	Saturation	Video
n/a	Scan Lines	Video
sepia	Sepia Tone	Video
frei0r.sharpness	Sharpen	Video
affineSizePosition	Size, Position & Rotate	Video
charcoal	Sketch	Video
spotRemover	Spot Remover	Video
n/a	Stabilize	Video
richText	Text: Rich	Video
dynamicText	Text: Simple	Video
threshold	Threshold	Video
timer	Timer	Video
avfilter.tmix	Trails	Video
n/a	Unpremultiply Alpha	Video
vertigo	Vertigo	Video
vignette	Vignette	Video
wave	Wave	Video
frei0r.colgate	White Balance	Video

Preset Location	Filter Name	Type
ambisonic-decoder	Ambisonic Decoder	Audio
audioBalance	Balance	Audio
ladspa.1892	Band Pass	Audio
ladspa.1882	Compressor	Audio
n/a	Copy Channel	Audio
ladspa.1192	Delay	Audio
n/a	Downmix	Audio
15BandEq	Equalizer: 15-Band	Audio
3BandEq	Equalizer: 3-Band (Bass & Treble)	Audio
parametricEq	Equalizer: Parametric	Audio
ladspa.1883	Expander	Audio
n/a	Fade In Audio	Audio
n/a	Fade Out Audio	Audio
volume	Gain / Volume	Audio
ladspa.1890	High Pass	Audio
n/a	Invert	Audio
ladspa.1913	Limiter	Audio
ladspa.1891	Low Pass	Audio
n/a	Mute	Audio

Preset Location	Filter Name	Type
ladspa.1410	Noise Gate	Audio
dynamic_loudness	Normalize: One Pass	Audio
n/a	Normalize: Two Pass	Audio
ladspa.1894	Notch	Audio
audioPan	Pan	Audio
rbpitch	Pitch	Audio
ladspa.1216	Reverb	Audio
stereoEnhance	Stereo Enhancer	Audio
n/a	Swap Channels	Audio

Preset Location	Filter Name	Type
audioDeclick	Declick	Time
speedForwardReverse	Speed: Forward & Reverse	Time
speedForward	Speed: Forward Only	Time
timeremap	Time Remap	Time