

Shotcut User Guide

Table of Contents

(December, 2024)

Glossary / Terminology

Launching Shotcut

Quick Start Guide

Project Management

Adding Media

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

Keyboard Shortcuts

File > Save and Save As

File > Backup and Save

File > Open MLT XML as Clip

File > Export > Frame

View > Notes

View > History

View > Resources

View > Subtitles

Help > Actions and Shortcuts

Settings > Video Mode

Settings > Audio Channels

Settings > Preview Scaling

Settings > Proxy Editing

Settings > Player > Realtime (frame dropping)

Settings > Player > Interpolation

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
Timeline > Output
Timeline > Ripple All Tracks
Timeline > Nudge
Timeline > Markers
Timeline > Record Audio / Voice-over
Timeline > Align To Reference Track
Timeline > Apply Copied Filters
Subtitles > Speech To Text
Export Basics
Export failed - Find solutions
Advanced Export
Export > Resample
Embedding metadata into Exports
Export at the Command Line
Filter Sets
Properties > Extract Sub-clip
Open Other > Color
About Keyframes
Simple Keyframes
Audio Loudness Scope
Audio Peak Meter Scope
Audio Surround Scope
Audio Vector Scope
Video Scopes
Video Waveform Scope
Video Zoom Scope
360: Equirectangular Mask Video Filter
360: Equirectangular to Rectilinear 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

Audio Dance Visualization

Audio Waveform Visualization

Blend Mode Video Filter

Chroma Key: Advanced Video Filter

Chroma Key: Simple Video Filter

Corner Pin Video Filter

Crop: Rectangle Video Filter

Crop: Source Video Filter

Deband Video Filter

Fisheye Video Filter

GPS Text Video Filter

Lut (3D) Video Filter

Mask: Apply Video Filter

Mask: Chroma Key Video Filter

Motion Tracker Video Filter

Reduce Noise: Wavelet Video Filter

Reframe Output Video Filter

Sharpen Video Filter

Size, Position & Rotate Video Filter

Subtitle Burn-In Video Filter

Text: Simple Video Filter

White Balance Video Filter

Downmix Audio Filter

Noise Gate Audio Filter

Pitch Audio Filter

Stereo Enhancer 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.

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.

Safe area

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

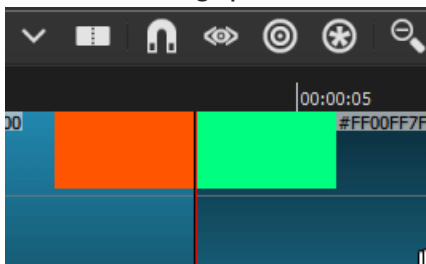
Scrubbing is seeking by clicking some object and dragging it. Typically this is the play head in the player or timeline. But it can also refer to simply rewind and fast forward playback through media.

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.

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

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 named .mlt to named yyyy-mm-ddThh-mm-ss.mlt while keeping named .mlt 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.

Modes	Description
Playlist	Arrange clips in order, apply filters if needed, and then export. No need for the Timeline.
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.

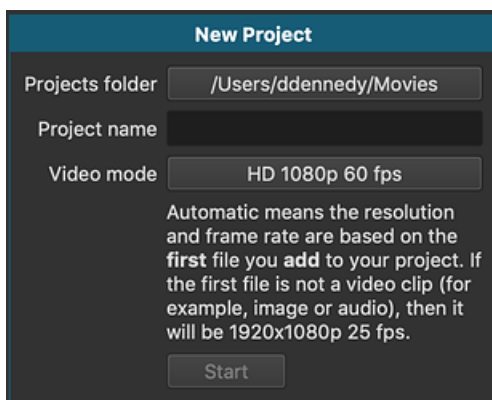
3. Timeline

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**

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 copy or move the files into there automatically at this time.

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.

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 the “**i**” key.

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 the “**o**” key.

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.




Example

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 Filters

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.



How to Reduce Memory Usage

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 can not correctly edit HDR clips directly. However, HDR clips can be converted to a standard format 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:

- ITU-R BT.2020
- SMPTE ST2084 (PQ)
- SMPTE ST428





Converting HDR Clips

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

1. Click on the “Convert...” button in the Properties panel. The Convert button will open a dialog.
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.



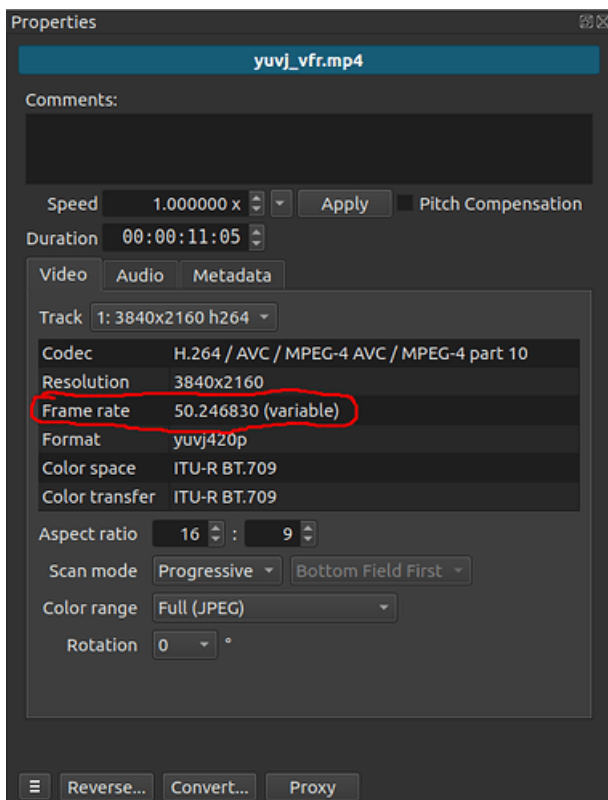
Introduction

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 **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.

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](#).


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

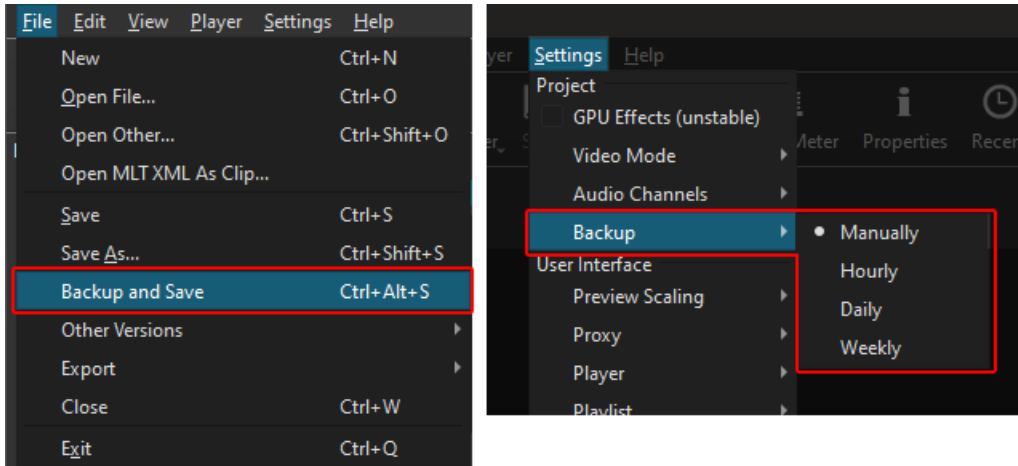
This is only available if you have already saved a project (“Untitled” does not appear in the window title bar). It first gets the last date and time the project file was modified. Then, it duplicates the current project file with a new name that ends with the date and time of last modification. Finally,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

This menu item was added in version 23.09.

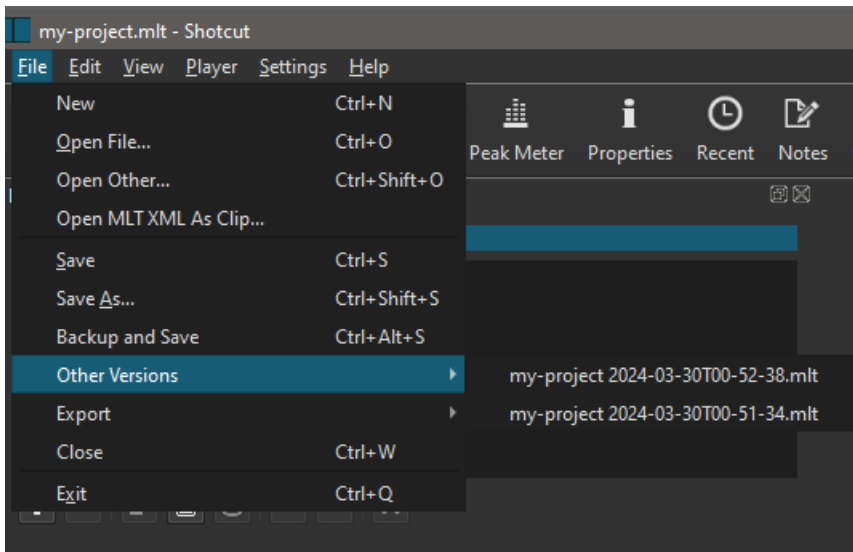
File > Backup and Save

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.

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.
-

If you are a trusted user you can edit this page to improve Shotcut's [documentation!](#)

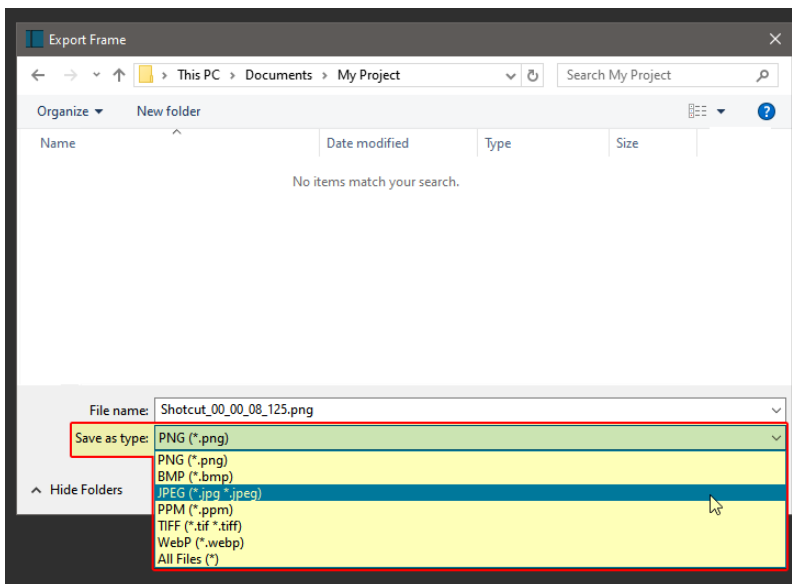
Export Frame is a feature that lets you save the frame currently displayed in the Shotcut viewer as a still image. The available formats are: *PNG*, *BMP*, *JPEG*, *PPM*, *TIFF*, or *WebP*.

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

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

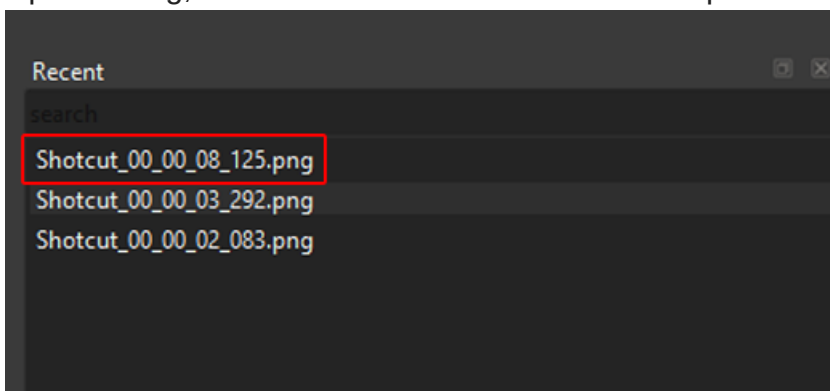
- Using the `Ctrl` + `Shift` + `E` keyboard shortcut.
- Going to **File > Export** and choosing **Frame...**

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:

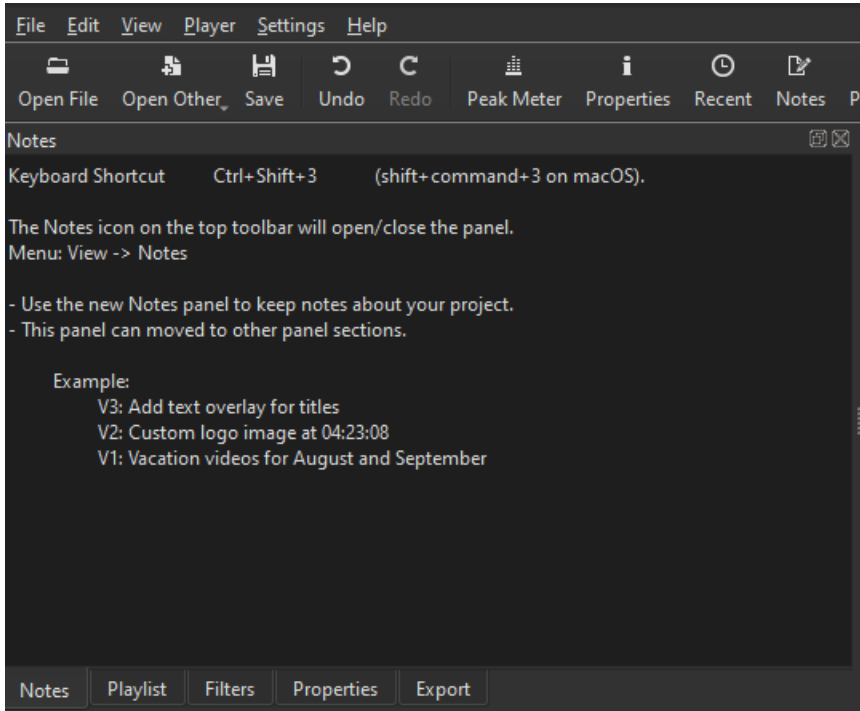
- Export the frame you want to use as a thumbnail.
- In the Recent panel, double-click it to open it in the **Source** viewer.
- Apply a Text filter to add a title.
- Use **Export > Frame** again to export the new thumbnail.

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 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.

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.

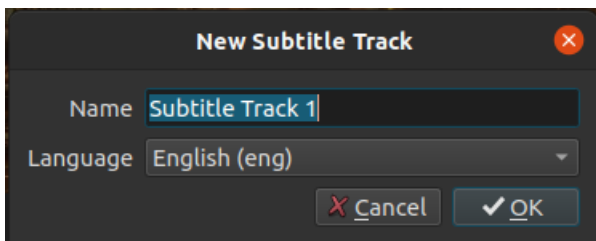

Clicking **Convert** shows the same **Convert to Edit-friendly** dialog that you can reach from the **Properties** panel.

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.



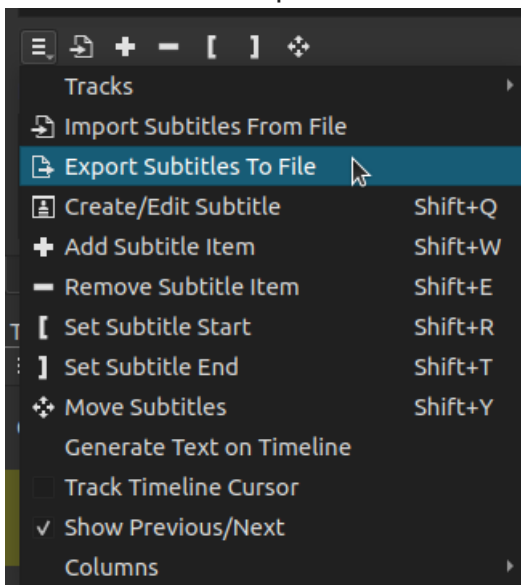
Importing Subtitles

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.

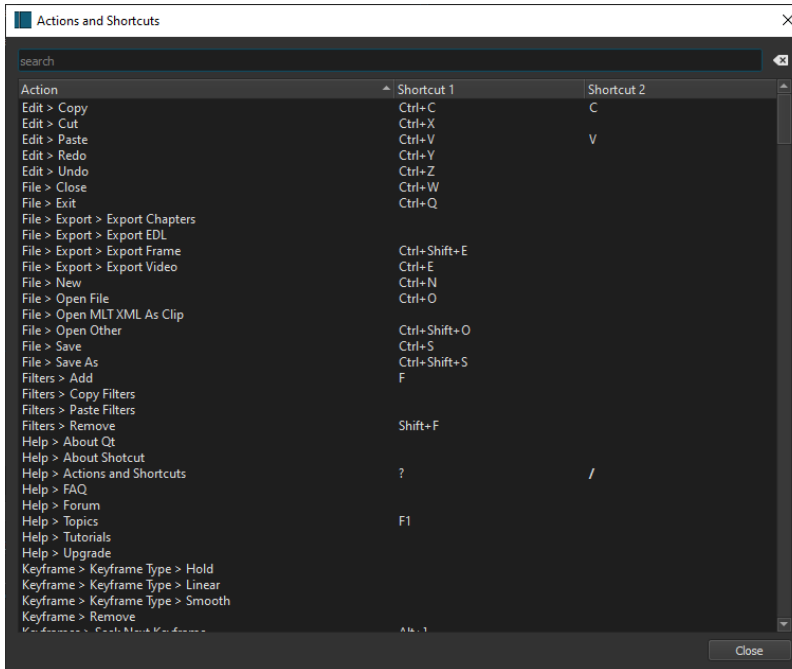


Embedding Subtitles

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.

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.

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.

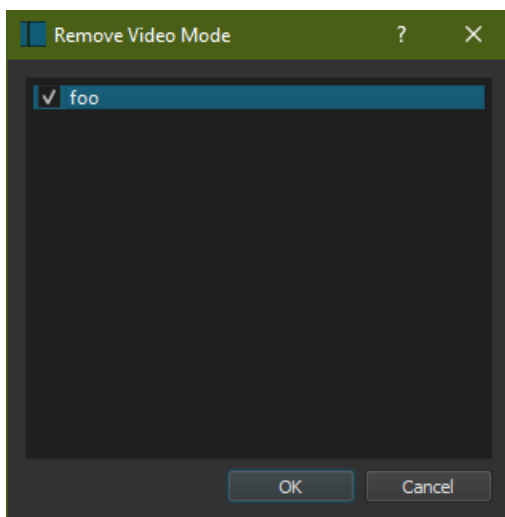
The maximum resolution is 8192 x 8192. Only even numbers are accepted because most video codecs only support even values.

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.




Settings > Audio Channels

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

Settings > Preview Scaling

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 [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.

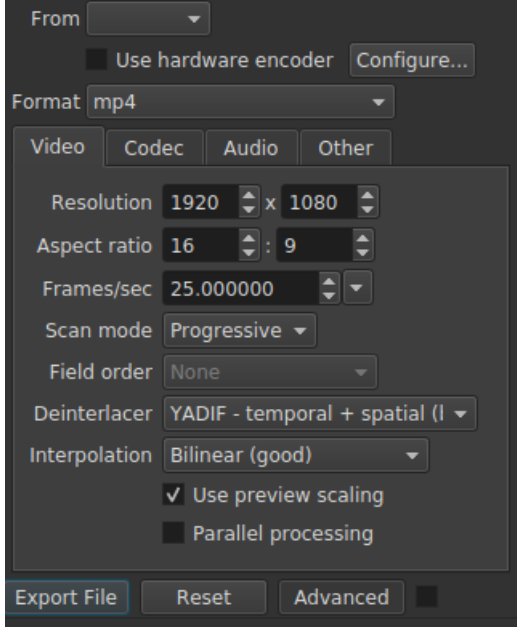
The submenu lists only the vertical resolution. The horizontal resolution depends on the aspect ratio of your project.

- **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.



Settings > Player > Realtime (frame dropping)

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

Interpolation is the computation of values based on neighboring values. It easiest to think of this as the quality level of scaling - or changing the size of an image - but it also affects things like positioning and rotation. It is also called sub-pixel rendering because it gives the impression of positioning between actual pixels.

This settings appears in the **Player** category of **Settings** because this only affects preview. There is a separate interpolation option in **Export > Advanced > Video**.



These names are based on algorithms. You can search for them on Wikipedia to get started on learning more about theme. Shotcut provides hints about the quality and speed trade-offs. **Nearest Neighbor** actually means no interpolation, no sub-pixel rendering.

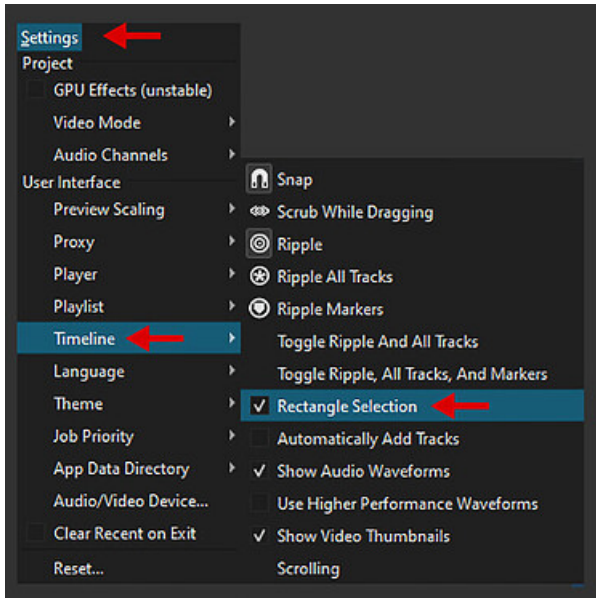
This setting does not affect text or SVG anti-aliasing. It also does not apply to HTML.

Settings > Timeline > Rectangle Selection

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.

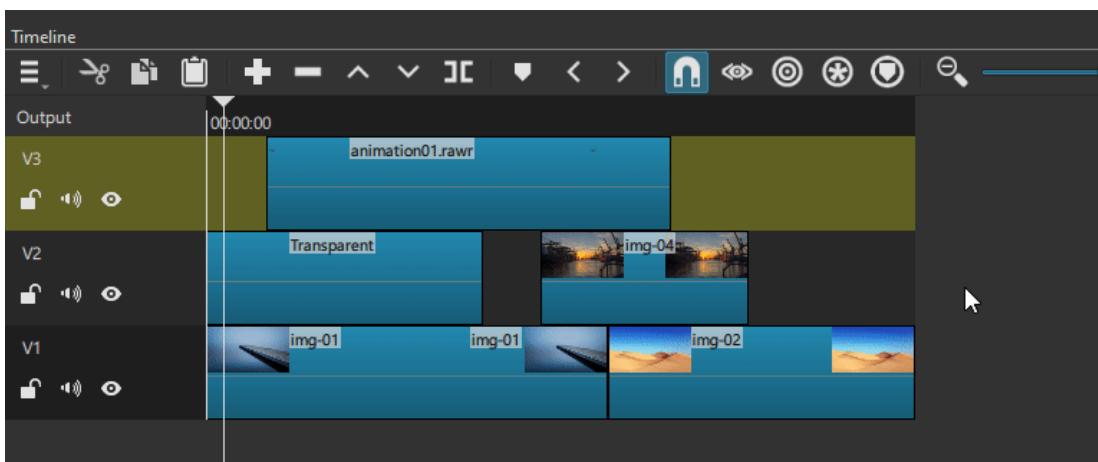
To disable or enable it, go to **Settings > Timeline**



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)

DEMO



(click anywhere on the image to play the animated GIF)

Shortcuts:

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



Settings > Timeline > Automatically Add Tracks

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

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.

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](#).




Settings > Clear Recent on Exit

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.

About the Playlist



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:

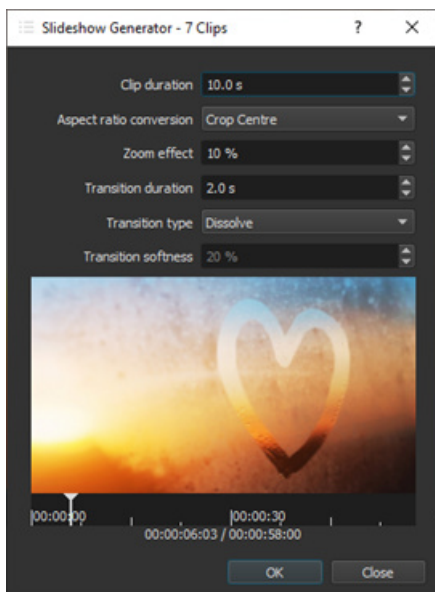
- fill the frame,
- make a smooth zoom animation for still images, and
- add a transition between each item.

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)




Some Use Cases and Suggestions

- Make a video loop for an information screen.
- Show off holiday or travel photos.
- You can now 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 ($\text{Length of track} / \text{number of slides}$) 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.
- Use music files and set **Clip duration** very high to make a music bed track that cross-fades between each song. The resulting clip duration is automatically limited to the clip duration and **Clip duration** acts as the maximum duration.

Source vs. Project Player

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).


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.

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

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)

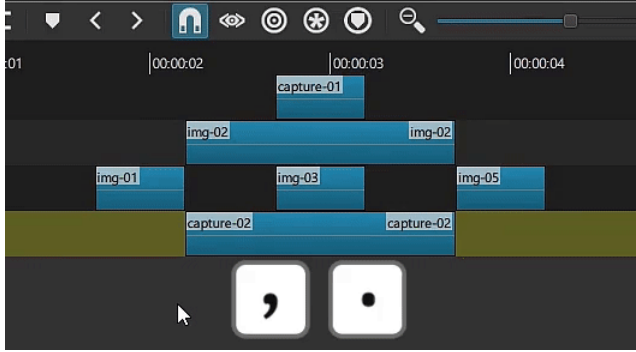
Limitations

- 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)

Timeline > Markers

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.



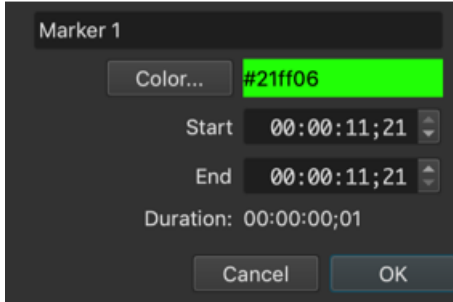
- Markers appear in the time ruler/track/row at the top of the **Timeline**.
- Drag a marker to move it.
- There are several ways to seek to a marker: Click directly on a marker, press **<** or **>**, or click the seeking buttons on the timeline toolbar.




- 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.



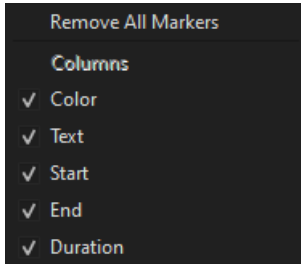
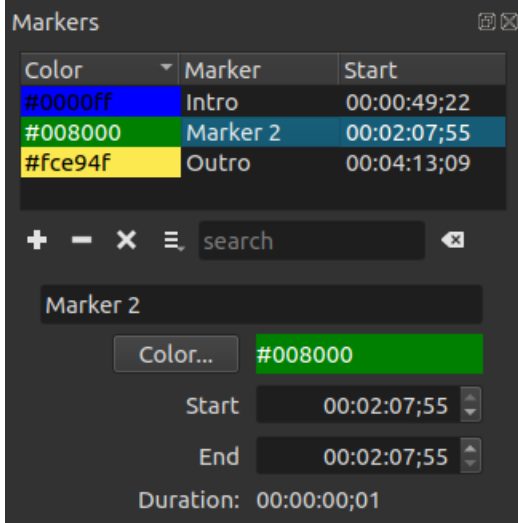
- A color change is remembered for all new markers until a new color is chosen.
- Right-click a marker to open its context menu to **Edit**, **Delete**, **Choose Color...**, or **Choose Recent Color**.




- 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.



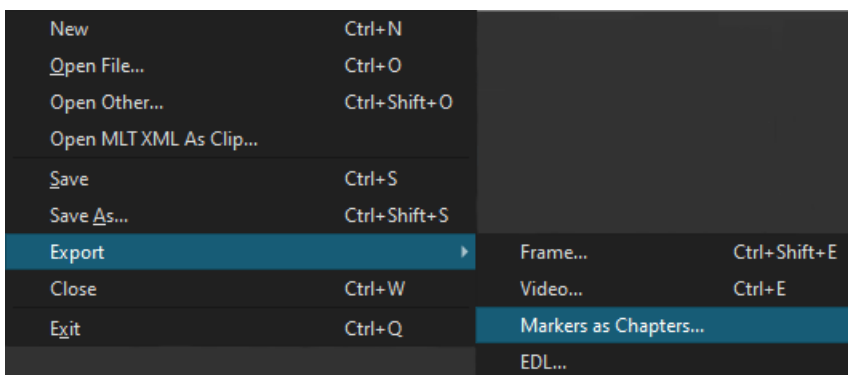
- **View > Markers** or press **Ctrl** + **Shift** + **6** (**shift** + **command** + **6** on macOS) to open the Markers panel where all markers are listed and can be categorized by several columns (**Color**, **Marker**, **Start**, **End**, and **Duration**). Clicking on a marker in the list will seek 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.



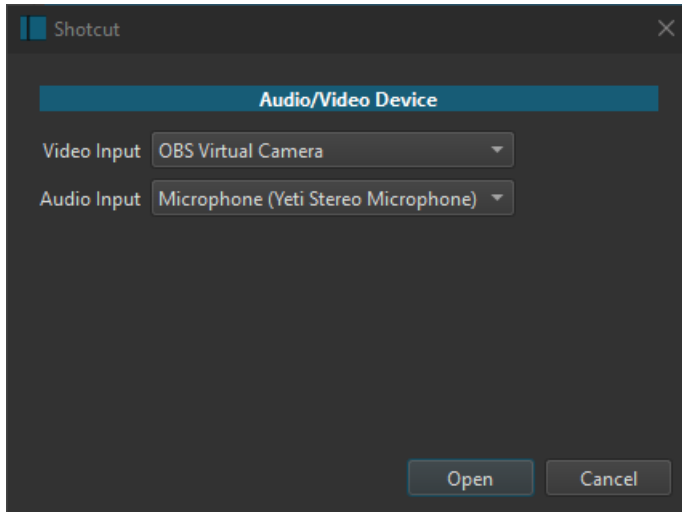
- **File > Export > Markers as Chapters...** outputs a .txt file in the format used to make chapters for YouTube. Only timeline markers with a duration of 1 frame (default) are exported as chapters.



- 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




Click on Open

2. Click on the Timeline
3. Click on the Microphone symbol on the Timeline



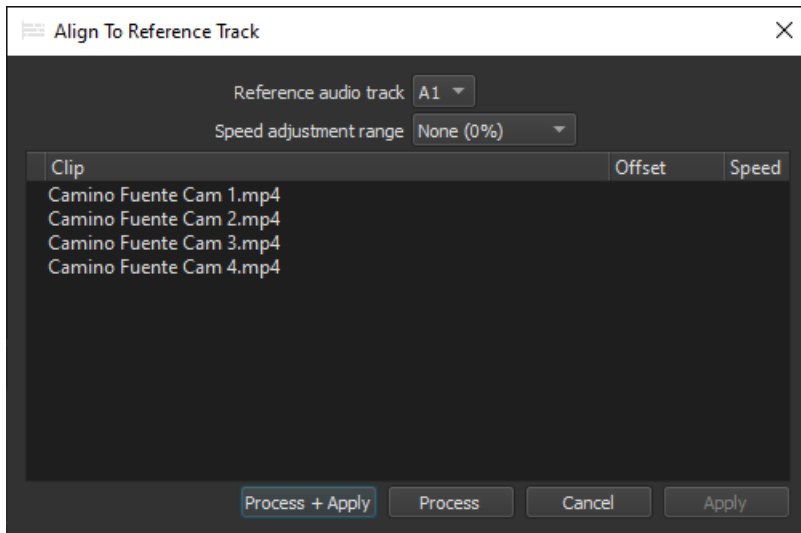
4. Name the Voice Over audio file.
Once you click "Save" for the file name, it starts recording immediately.



5. When done, click the same Microphone symbol. This stops recording and saves to the file name you specified.

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

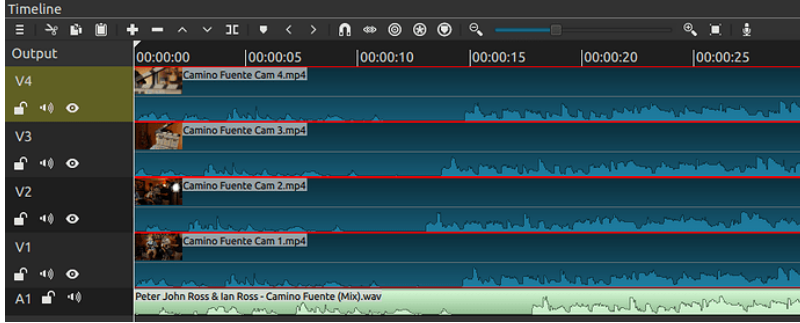
- This tool can only be used to align clips with audio that is similar to the audio in clips on another track. Both the clip to align and the reference track must have audio that is similar so that the algorithm can detect the alignment.
- The tool does not use timecodes for alignment

Usage

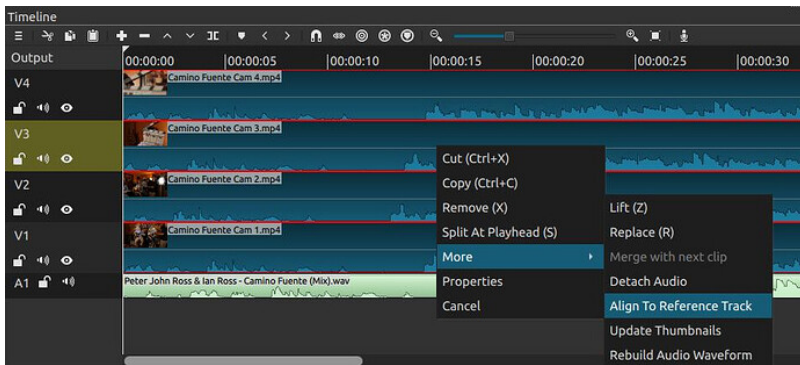
Place clips with similar audio on separate tracks. For example, put all the clips from each camera source on its own track.

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.

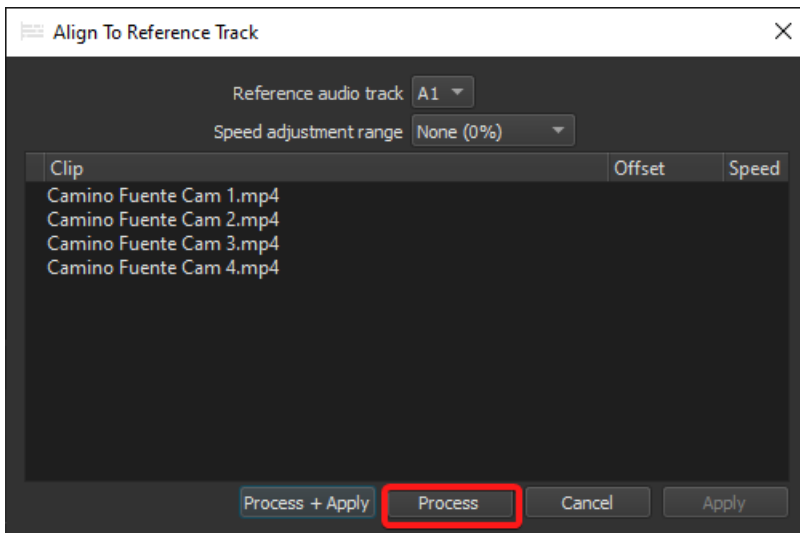
Select the clips to be aligned (do not select any clips on the reference track)



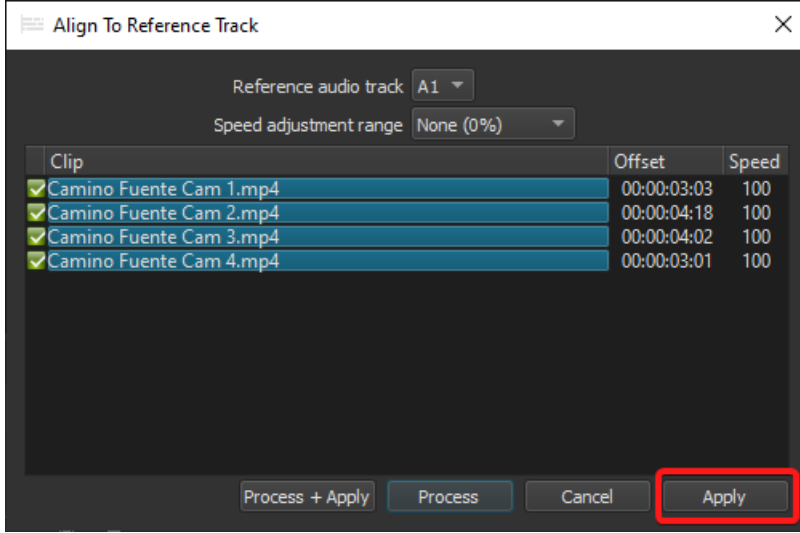
Right-click on the selected clips and choose from the menu: More->Align To Reference Track



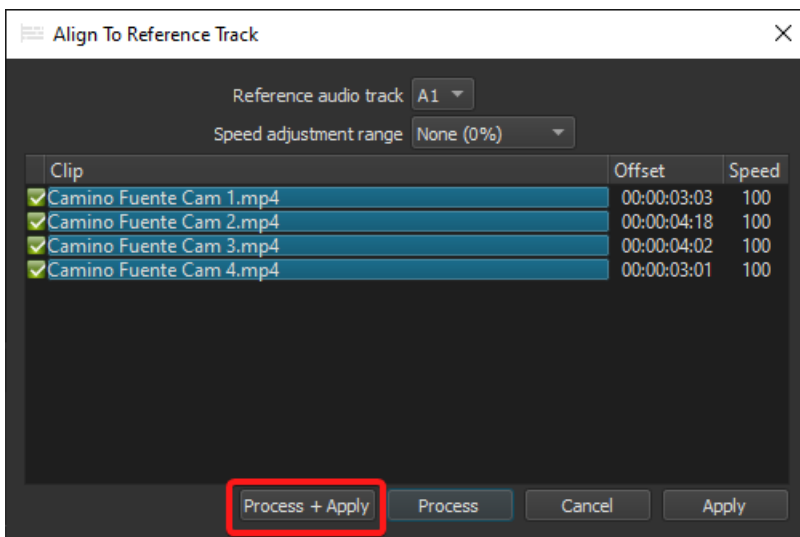
In the dialog, select the reference track. Then click “Process”



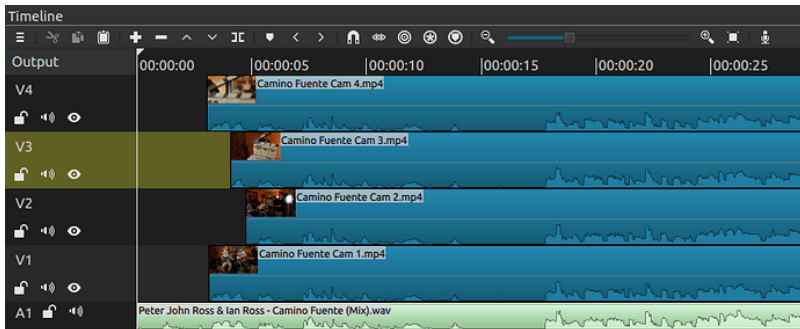
If the results look good, click “Apply”.



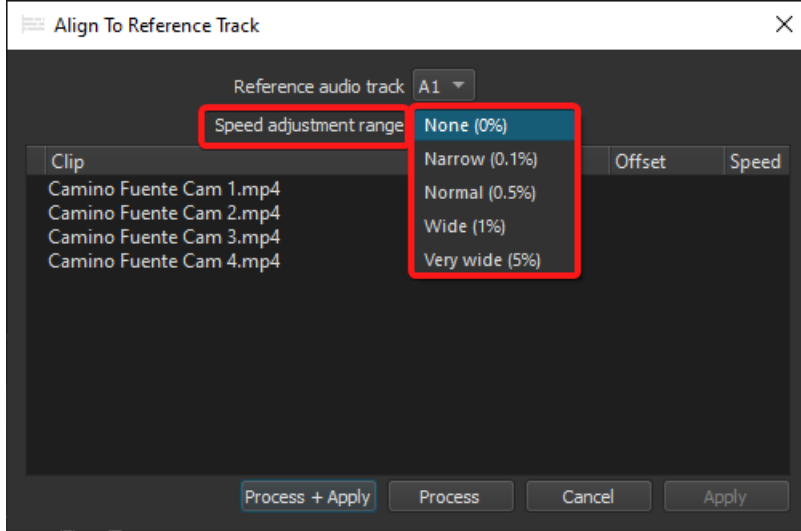
Alternatively, both “Process” and “Apply” can be run as one action by clicking “Process + Apply”.



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.



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.

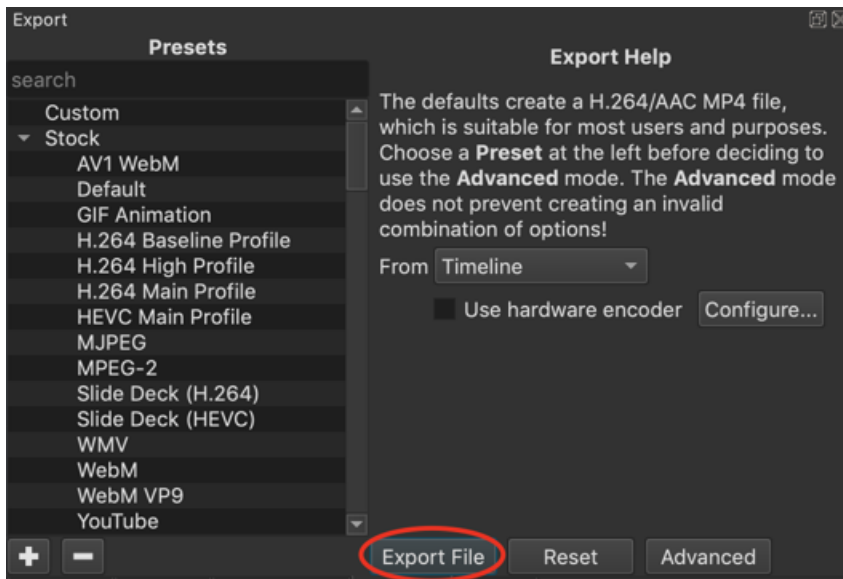
Known Quirks:

- 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.

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! You can export a single clip that has been trimmed and filtered in the **Source** player. You can also add that clip to the **Playlist**, add more clips, and export the playlist as a single, sequential file or a separate file for each item.
- **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



How to compress a video | Reduce video file size on iPhone, Mac, and more |...

Learn how to compress video on iPhone, Android, Mac, and with popular video compressors online. Reduce file size and keep your video high-quality.


Est. reading time: 9 minutes

Export failed - Find solutions

Your export failed for one or more reasons. There is always a reason why it failed. Below are the posts by others describing their faults and solutions. This tutorial is not a definite tutorial as you may have a unique situation not described in here.


Time Code

After failure to export it may display a time code, make a note of the time displayed and look on your timeline to see what's there that could be causing the problem. It could be a clip, audio or video file, transition.




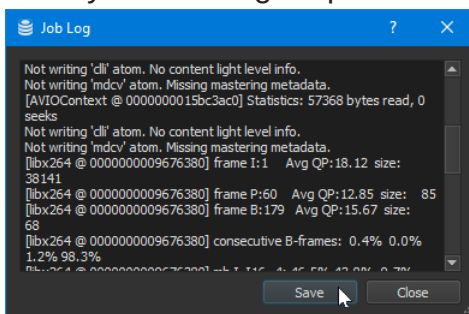
Export Job Log

The major source you're going to look at is the Export job log. When the job failed it will be in the job's window. At times it may be requested to share "View XML" as well. Right click on the job then View Log.



To share with this forum save the Job Log by clicking Save, and name the file.

Then you can drag/drop the file into the forum or upload by the Upload icon. 



You may not know what to look for, but look at the bottom of the log for clues to what is going on. There may be **Exit codes** at the bottom of the report.

You can always search the forum with that exit code to find a solution, or suggestions to how to resolve the export failure.

I compiled a list of posts with problems, suggestions, and some solutions. I only went back a few years and even few issues have been resolved with the [latest release](#) of Shotcut. Click the arrows the expand the list for each code.

Windows Controlled Folder Access

If you are on **Windows**, one thing to look for in the job log is a line near the end like

```
[consumer avformat] Could not open ...
```

This is most likely due to **Windows Settings > Windows Security > Controlled folder access**. This is not a bug in Shotcut. You can read more about it on Microsoft's Web site [here](#) and [here](#).

Antivirus

👁 If you're using an anti-virus/malware set the program to approve **melt.exe** and **ffmpeg.exe**.

-
- ▶ Failed with exit code -1073741819
 - ▶ Exit Code 1
 - ▶ Exit Code 3
 - ▶ Exit Code 11

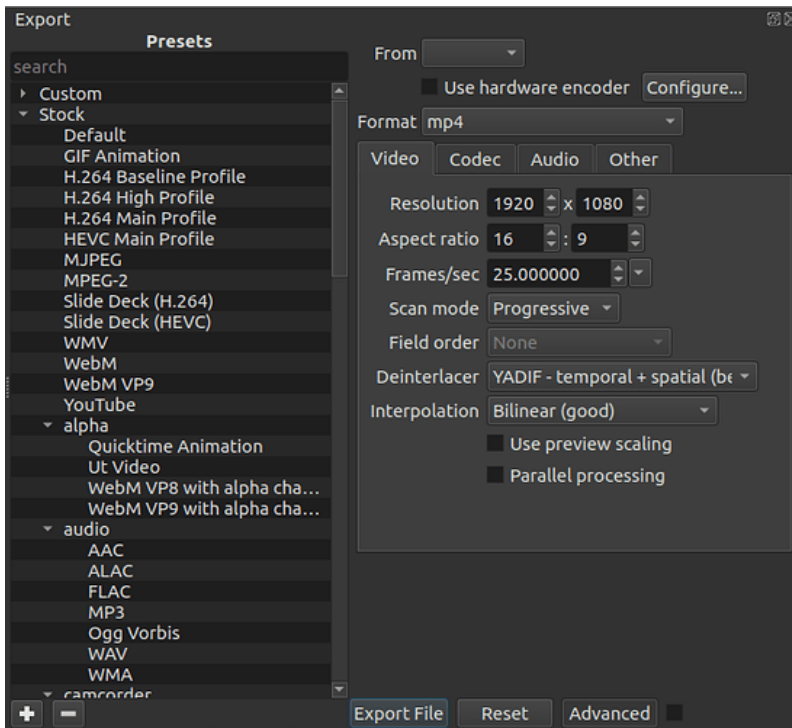
If you can't find a solution that works for you, fill out a new post and describe what's going on.

Note: If there are more relevant export fail posts to add I can certainly edit this tutorial to include them.

Insufficient Memory

A common cause of an export failure is running out of system memory. While an export is in process, use the task manager (system monitor, etc) to observe the memory utilization. If the memory is exhausted, the export may fail. See [How to Reduce Memory Usage](#)

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**.



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

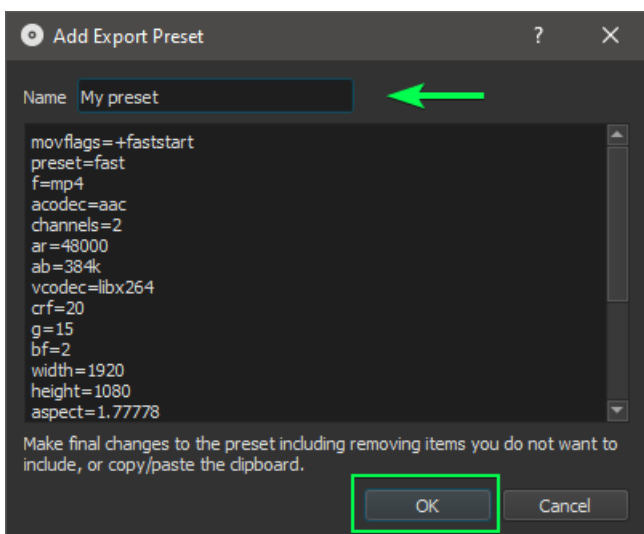
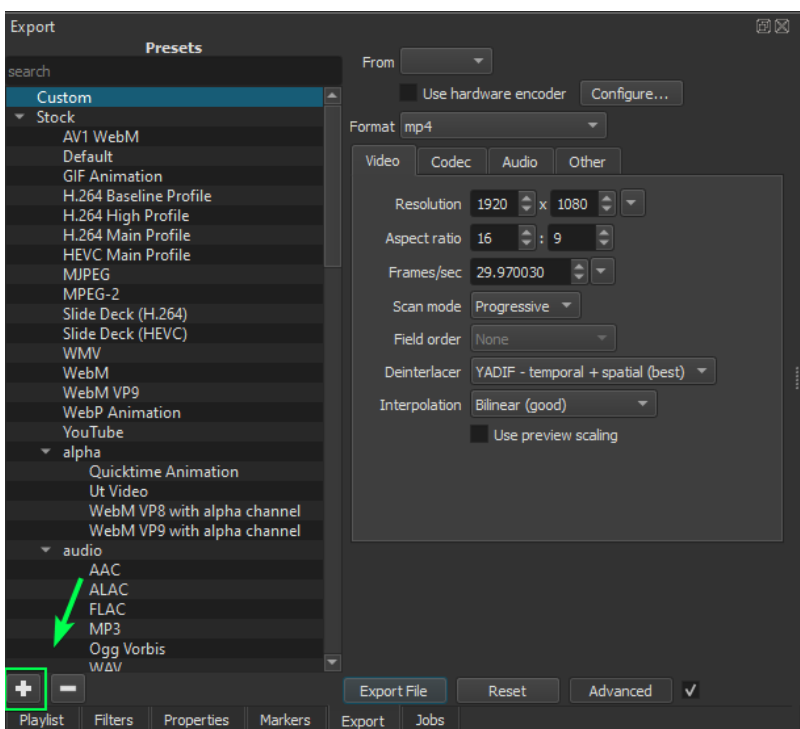
many things you cannot copy from ffmpeg examples are possible here. Only options from FFmpeg that affect the encoder or muxer.


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



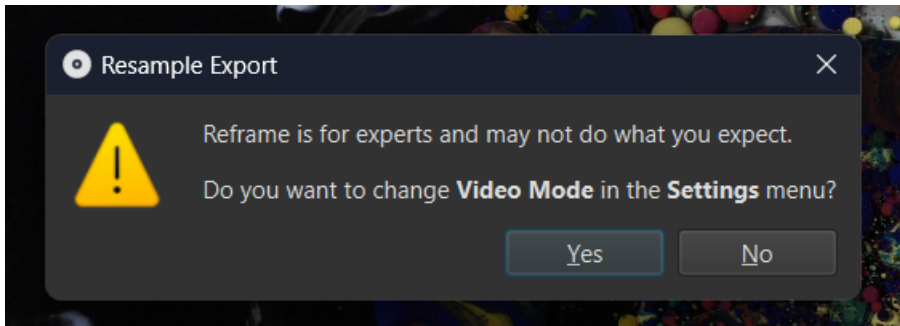
Resample is a button in **Export > Advanced > Video** that *enables* the **Resolution**, **Aspect ratio**, and **Frames/sec** fields.

This was added in version 24.09, but the concept of resample in export always existed in Shotcut. It is named “Resample” because it usually occurs after each frame of the Timeline or Playlist is *rendered*. Export = render + encode (uncompressed is formatting and writing).

Some sample uses for this feature:

- Video Mode resolution = source (can make editing interaction faster), AND you want to down-sample to make a smaller file.
- You composed this project at resolution X but I need Y, AND you have some stuff sensitive to size and position, AND you do not want to take the time to adjust these by changing Video Mode. Best quality is not a concern.
- You made this project at aspect ratio X but you want Y, AND you really do want black bars.
- You made this project at frame rate X but you want Y, AND $Y < X$.
- All your sources are \leq Video Mode resolution, AND you are not using text or vector graphics, AND you want to upscale export in order to game YouTube’s encoding (or avoid some other upscaler such as in my TV).

In version 23.09 clicking this shows a dialog:



(“Reframe” in that dialog was a typo and should have been “Resample.”)

Click **No** to acknowledge that you really do not want to change the Video Mode at this time and resample the export.

In the next version 24.10 this button and dialog are removed and instead an inline warning message may appear that does not block anything:

Video

Codec

Audio

Other

Reframe

Resolution 1920 x 1080

Aspect ratio 17 : 9



Aspect ratio does not match project Video Mode, which causes black bars.

Frames/sec 29.970030



Embedding metadata into Exports

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:
a lang=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 movf1
```

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`

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

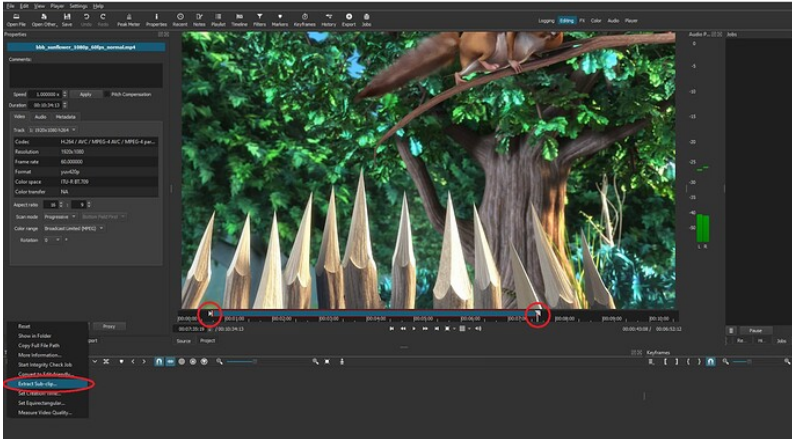
- Download the ZIP file containing the filter set and un-zip it.

- 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**



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 a way 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

Open Other > Color


Open Other > 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.


About Keyframes

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:-

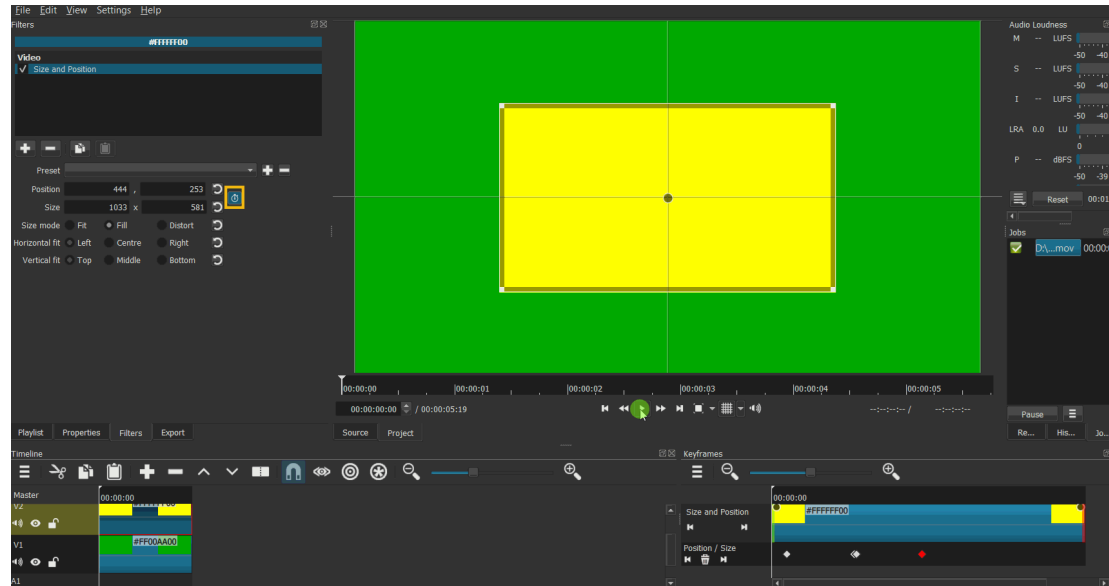




Image source from [Zoom Image and then Pan - #2](#) by samth

Simple Keyframes

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.


Audio Loudness Scope

The Audio Loudness scope provides a real-time calculation of the program loudness for the preview player.

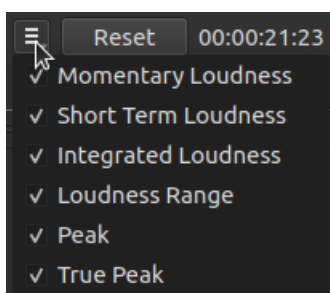
The meter displays 6 metrics as defined in [EBU-R128](#).

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.

Audio Peak Meter Scope

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.



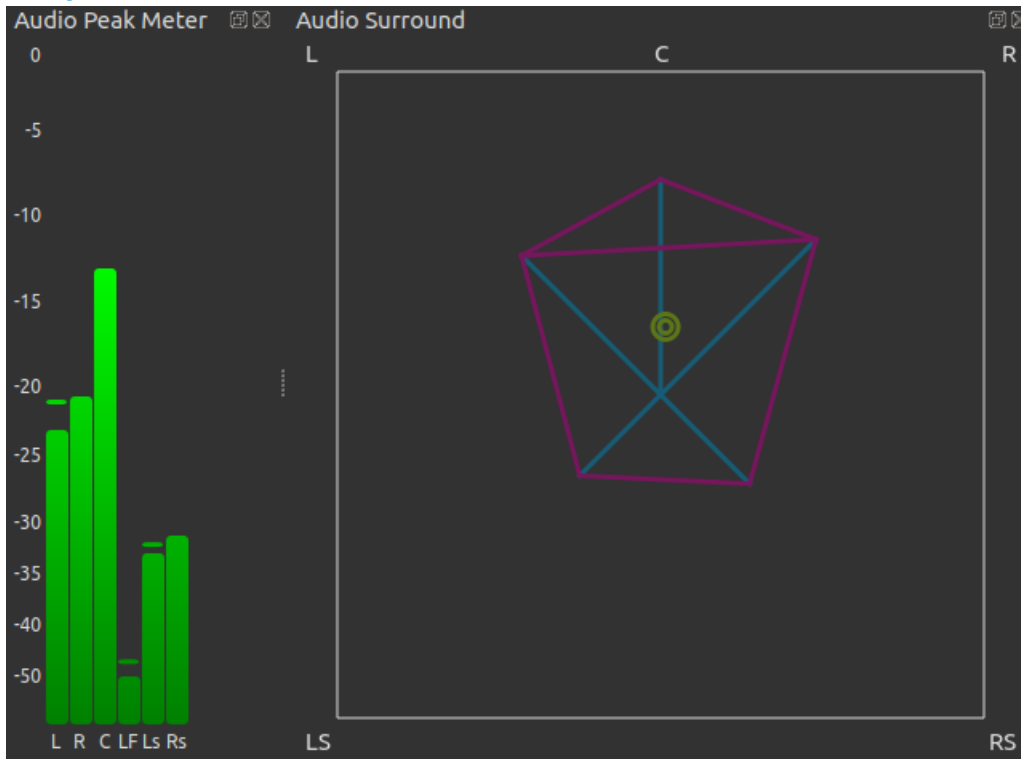
The scope also shows a “recent peak” indicator as a small line above the level bar.

Audio Surround Scope

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.

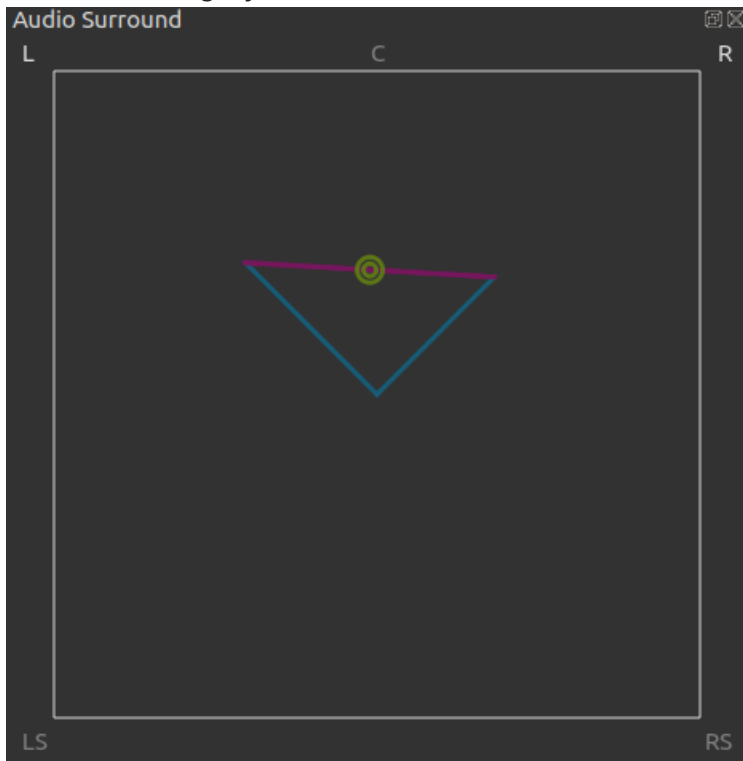
The blue lines extending from the center of the display visualize the signal strength for each channel. In the illustration above, you can see the relationship between the level for each channel in the Audio Peak Meter and the Audio Surround Scope.

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.




Audio Vector Scope

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:



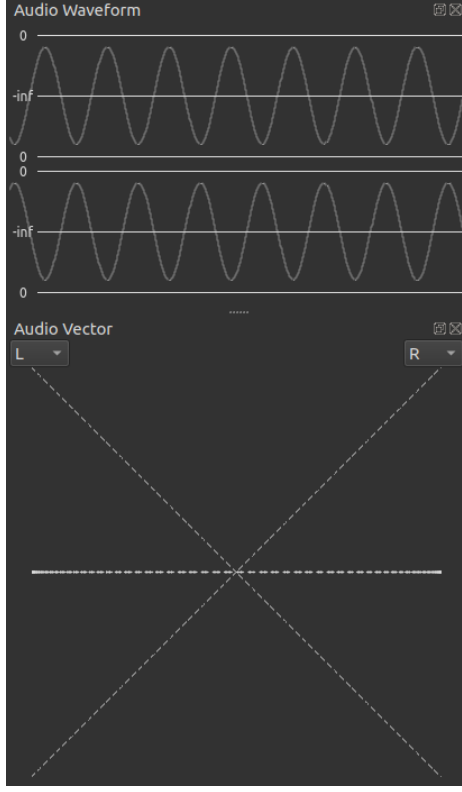
To show the phase relationship between different channels, choose the channels from the drop down boxes:



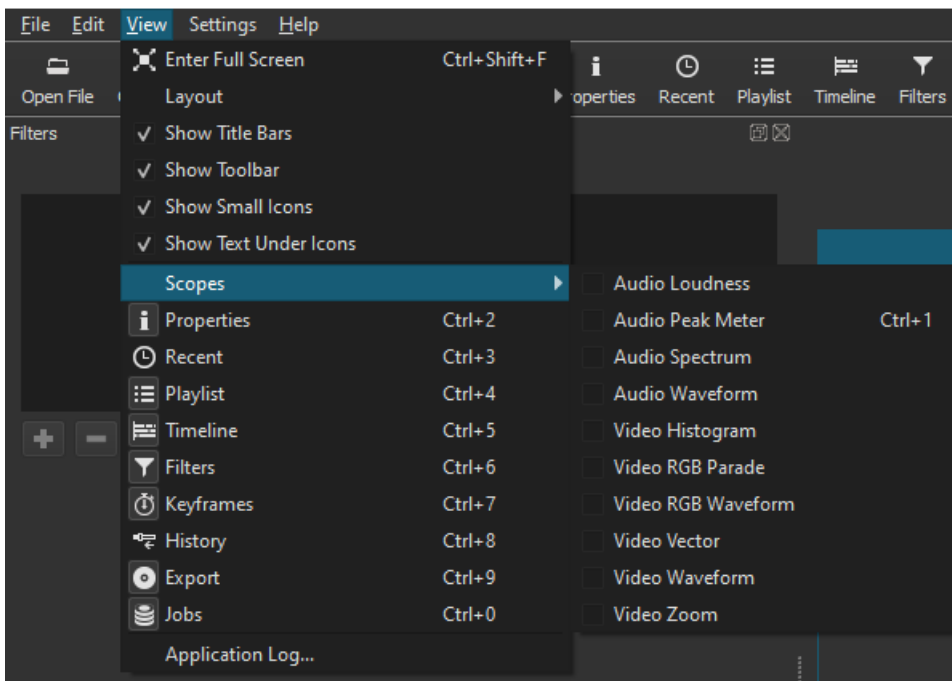
If the vector scope shows a vertical line without any horizontal separation between the points, that means the two channels are perfectly in phase. Typically caused by a mono signal that is split to left and right:



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.



Shotcut contains various video scopes to help with color adjustment and picture validation.



The video scopes represent the picture from the preview player. The preview player always displays as YUV 4:2:0. As a result, the representation in the scope does not necessarily represent the original source material. Nor does it necessarily represent what will be in the final export file.

Format Conversions

Shotcut manipulates images by applying various services (filters, transitions, playlist, timeline). Different services operate on different image formats. Shotcut will automatically convert the image format as needed as the image passes through the processing pipeline from one service to the next. Depending on the services that user has applied an image may go through multiple conversions before it reaches the preview player.


Internally, Shotcut supports two image formats:

- Limited range YUV as defined by [BT.709](#) and [BT.601](#)
- Full range sRGB as defined by [IEC 61966-2-1](#)

Note that a full range YUV image will be converted to limited range upon decoding. If a limited range source image has values that are outside of the limited range, those out-of-bounds values will be passed to the preview window (and therefore the scopes) unless they go through a conversion. In order to satisfy the needs of various services (transition, filter, etc.) an image may be automatically converted

between YUV and RGB. These automatic internal conversions will result in out-of-bounds values being rounded down to be within the limited range.

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.

Video Zoom Scope


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.



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%.



360: Equirectangular to Rectilinear Video Filter

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.



360: Hemispherical to Equirectangular Video Filter

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.



360: Rectilinear to Equirectangular Video Filter

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](#)

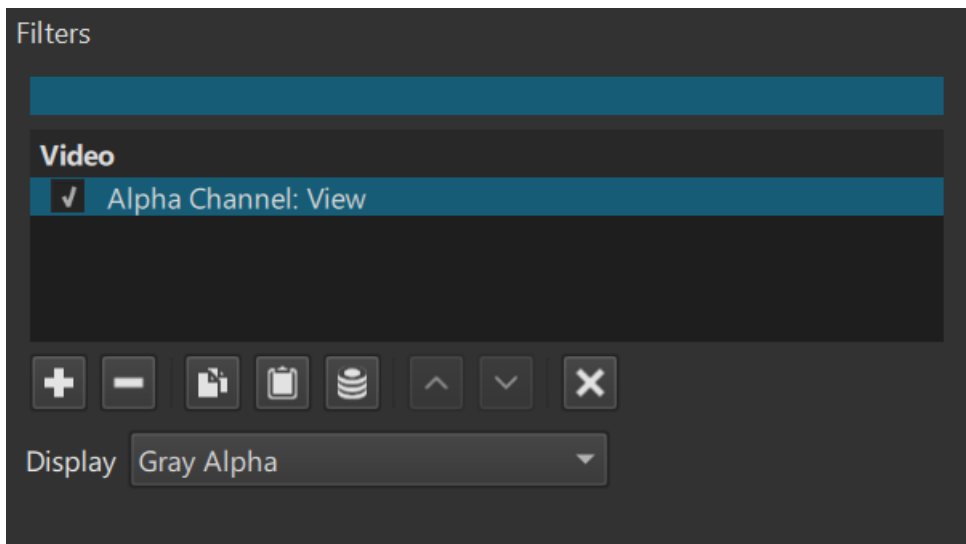
A large collection of free and portable video plugins - dyne/frei0r

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.




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.



Demo with all settings to default:

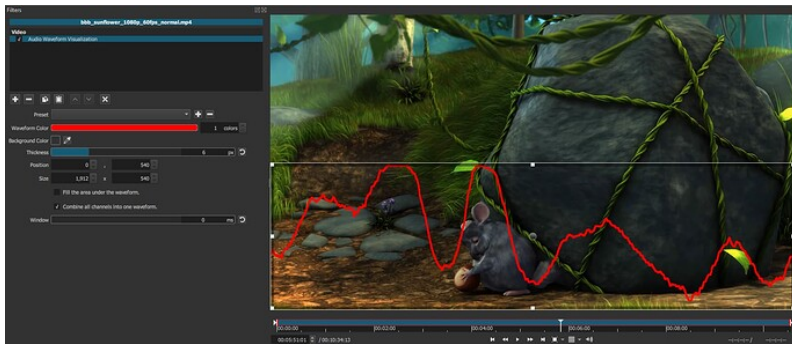
Settings

- **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.

Audio Waveform Visualization


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**.

Blend Mode Video Filter


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 **Over** (simple alpha compositing).



Most of the options are the standard **Porter/Duff** blend modes. Full technical details about the different blend modes are defined in the [Cairo graphics library](#).

Chroma Key: Advanced Video Filter

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.

Chroma Key: Simple Video Filter

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:

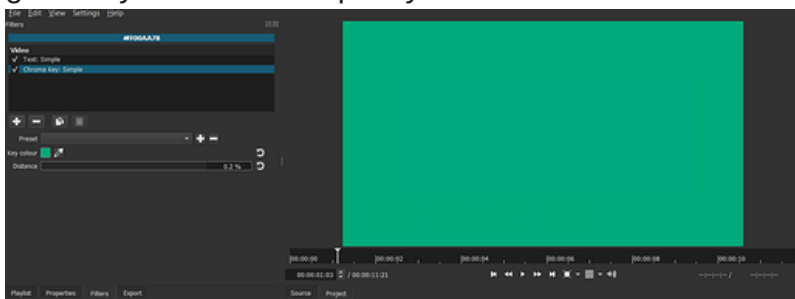
The **Key colour** sets the background colour. The eye dropper icon can be used to activate the colour picker. By clicking now on the background, the colour will be selected as the key colour.

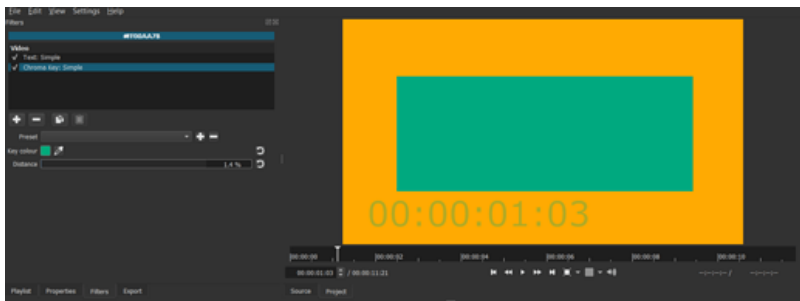

The **Distance** adjusts which shades or colours are treated like the key colour. This is especially important for similar colours that are present in the background as well as in the foreground. If the colours are very similar, only few percentages decide whether both are keyed out or not. This is also 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.

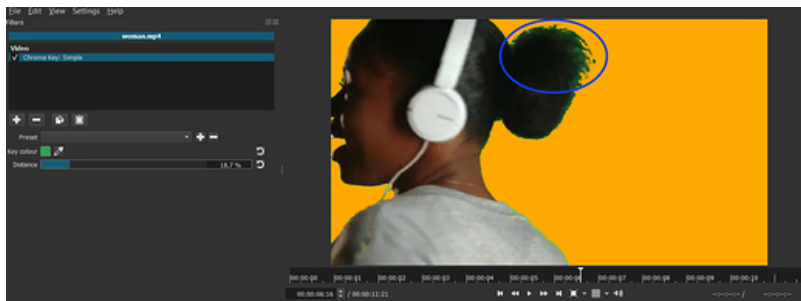




In this example one strand of hair is partially removed beyond a certain percentage.




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.

Corner Pin Video Filter

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

Crop: Rectangle Video Filter

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.

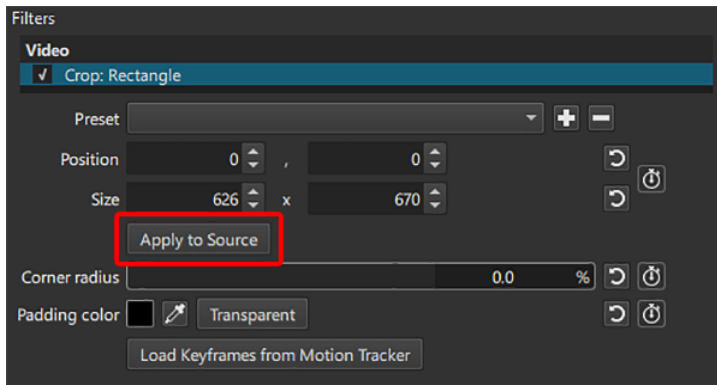


Example #2:

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 Video Filter

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

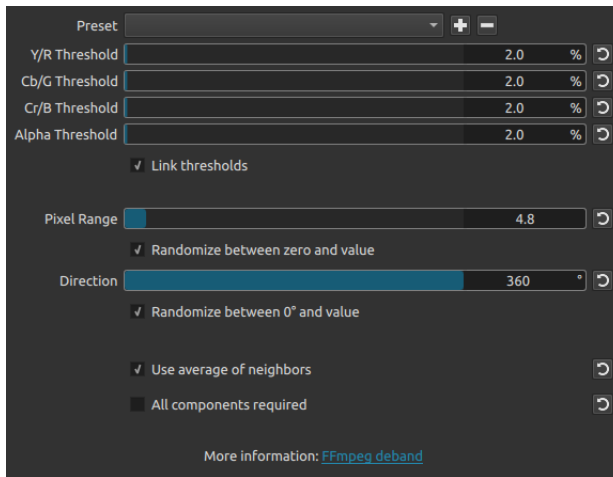
The values are in pixel units in terms of the current Video Mode, not the source resolution that may change depending upon whether a proxy is being used. Therefore, to convert the values from source to video resolution use the equations:

- $\text{Top/Bottom} = \text{source-top} / \text{source-height} * \text{mode-height}$
- $\text{Left/Right} = \text{source-left} / \text{source-width} * \text{mode-width}$

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.

What Is It?

The Deband filter tries to remove banding from an image by replacing streaks of solid colors (bands) with gradients or dithering.



Quick How-To

Add the Deband filter to a clip. Start with a preset based on the following criteria:

- For minimal processing, use the “Full range to Limited range” preset. (This is also the default setting.) This preset is designed to gently remove bands that appear when full-range video is dropped onto Shotcut’s limited-range timeline. Minimal processing is good for 2D graphics and PowerPoint-type material where edges should remain crisp.
- For average processing, use the “FFmpeg default values” preset. This is good for live-action camera footage.
- If a streaky blue sky is involved, heavier processing from the “Blue gradient in open sky” preset may be needed. For a red sunset, try the “Red gradient in open sky” preset.

These are just starting points. Customize as needed.

Depending on the source video, the results of Deband may seem very subtle. If the output is viewed on a four-inch smartphone screen, this is certainly true. But on a screen that is “medium-sized” or larger, the Deband filter can create profound improvements.

Also, the most distracting thing about bands is their tendency to wiggle like caffeinated worms during video playback. Although Deband results may look subtle on a still image, using Deband to stop the worm-wiggling in a video makes a significant improvement in perceived realism and the viewer’s ability to stay focused.

Common Uses

Banding is a common side effect of the following processes:

- Full to limited range conversion (including image files)
- Low-bitrate codecs smearing gradients to reduce file size
- Compression macroblocking is essentially a series of short bands
- Posterization or clipping introduced by extreme post-processing filters

“Full range” refers to a video signal (RGB or YCbCr) that is encoded in the 0-255 range. “Full range” is also called JPEG or PC range.


“Limited range” refers to the 16-235 range, which leaves room for sync signals, ringing artefacts, and the Gibbs phenomenon. “Limited range” is also called MPEG, TV, legal, broadcast, or studio range.

Computer screen recordings are often RGB 0-255, and some video cameras save their files as YCbCr 0-255. Since the Shotcut timeline is YCbCr 16-235, a conversion must take place for these sources. Shotcut invokes `swscale` in a non-dithering mode that often introduces banding (overlap) as colors are squished into the reduced range. The Deband filter tries to detect these bands and convert them back into smooth gradients, typically by adding dithering.

Note that many image file formats (like JPEG, GIF, PNG) are encoded with 0-255 values. They will be converted to 16-235 when placed on the timeline, which makes them prone to banding. The Deband filter therefore has potential to make many still images look better than they would without any processing.

Demonstration

The following image has gradients that go from RGB 0 to RGB 255:




When Shotcut loads this image, banding can be seen in the Video Zoom scope as multiple lines having the same color value. Although the Y (luma) value changes from 16-235 without skipping values, the display conversion back to RGB 0-255 does skip values.

For instance, going from Y 61 to Y 62 correlates to RGB 52 to RGB 54. The skipped RGB 53 causes RGB 54 to look unexpectedly brighter than its neighbors. This discontinuity ruins the illusion of a smooth gradient.

The Video Waveform scope also shows uneven value distribution rather than a solid vertical bar. A solid bar would mean equal representation for every value from 0 to 100 IRE.

Below, we see what the image looks like after applying the Deband filter with the default settings:



Unfortunately, the forum software has shrunk the image so much that the dithering is hard to see. Here is an exaggerated view (enlarged and emphasized with a Curves filter):



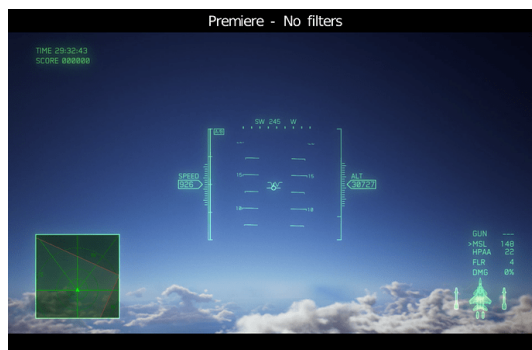
The strategy behind the default settings is to assume that bands will be very narrow (at most three pixels wide) as a result of gradients being squished during range conversion. This explains the small default pixel range. We also expect banding to happen only when values are very similar to each other (where overlap happens when squished), therefore the thresholds are small as well. These are the gentlest and least invasive Deband settings. For slightly more processing, use the “FFmpeg default values” preset.

Some scenes, like open blue sky, require more aggressive processing from the Deband filter. @IceFox2421 contributed the following screenshot (used with permission) which features a sky that was splotchy in the game itself, then was captured in 0-255 range, then was brought into Shotcut, which introduced even more banding with the timeline’s full-to-limited conversion. The screenshot also shows what the footage looks like in Adobe Premiere, which benefits from a better full-to-limited conversion process. Lastly, the screenshot goes back to Shotcut using the Deband filter, and shows very competitive results in the final image quality.

On some browsers, this will animate (click the image to enlarge first):



If not, here are the frames for comparison:



For open sky, the strategy is to use a higher threshold for blue and a lower threshold for all the other colors. Bear in mind that the Shotcut timeline is YCbCr unless overridden with `mlt_image_format=rgb24`, meaning blue is manipulated with the Cb of YCbCr rather than the B of RGB. Lastly, we use a large pixel range to scan for slowly-changing gradients, as the sky does not quickly change from one tone to the next.

How Does It Work?

The Shotcut Deband filter uses `avfilter.deband` ([FFmpeg](#) | [MLT Framework](#)) for processing. The general idea is to take a pixel and look at its neighbors within Pixel Range. If they're the same general color (within Threshold range), then assume this area should have been a gradient (or an in-between value) rather than a band. The pixel will be replaced with the average value of its neighbors. If the Randomize option is selected to skew that average, then any bands will in essence become dithered. If dithering introduces too much "snow" or impulse noise into the image, try adding a Sharpen filter with the Amount around or below 30% (which will actually create blur despite the name) and with Size around 12%. Sometimes this can blur small patches of dots out of existence without harming the overall image. In extreme cases, consider a denoise filter to smooth the dot patches out (but not to the point of recreating bands, of course).

Limitations

A more proper way to convert full range video to limited range is to provide a scaler with access to the original 0-255 data, then perform a dithered conversion using error diffusion in linear color space. `swscale` and `zscale` can be manually invoked in such a way, but this is not directly available in Shotcut. (See the [Alternatives](#) section below for details.) Since Shotcut is currently 8-bit, dithering is not invoked in `swscale`. Likewise, since Shotcut does not use `zscale`, conversions are not done in linear space.

Since the Deband filter is seeing data that has already been compressed to the 16-235 range by `swscale`, Deband does not have the luxury of examining the original 0-255 video data. This is an automatic disadvantage in terms of technical accuracy. But, given proper parameters on a scene-by-scene basis, even a discerning audience would struggle to see the difference.

Alternatives

Video can be manually pre-processed into limited range before importing into Shotcut. This gives `swscale` or `zscale` access to the original 0-255 values.

Let's create a sample 0-255 input video for this demonstration by using the test image provided at the top of this document:

```
ffmpeg -loop true -i FullRangeGradients.png -t
00:00:04.000 -r 30000/1001 -vsync cfr -pix_fmt gbrp -c:v
utvideo -an FullRangeGradients.mkv
```

Below are two ways to convert the sample input video into a dithered limited-range intermediate.

Convert with `swscale` passing through 10-bit so dithering can be invoked when returning to 8-bit:

(see [this note about dithering](#) on FFmpeg Trac)

```
ffmpeg -i FullRangeGradients.mkv -filter:v
scale=out_color_matrix=bt709:out_range=full:flags=neighbor+accurate_rnd+full_chroma_inp+full
-pix_fmt yuv422p -color_range mpeg -colorspace bt709 -
color_primaries bt709 -color_trc bt709 -c:v libx264 -qp 0
-g 8 -bf 0 -preset ultrafast -movflags
+faststart+write_colr -an -y FullRangeGradients-
swscale.mp4
```

Convert with `zscale` using dithering:

```
ffmpeg -i FullRangeGradients.mkv -filter:v
zscale=matrix=709:range=limited:filter=bicubic:dither=error_diffusion,format=yuv422p
-pix_fmt yuv422p -color_range mpeg -colorspace bt709 -
color_primaries bt709 -color_trc bt709 -c:v libx264 -qp 0
```

```
-g 8 -bf 0 -preset ultrafast -movflags  
+faststart+write_colr -an -y FullRangeGradients-zscale.mp4
```

Stack all three videos on separate tracks in Shotcut and toggle track visibility to notice that the swscale and zscale versions have smoother gradients than the source video.

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

StereogrphaicPanoramic

- 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 [Preview Scaling](#) will have a massive impact allowing realtime playback even for 4K sources.
- See also [Lens Correction](#) which supports keyframes unlike this 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.

0:00



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.

```
Text: Speed: #gps_speed# km/h  
Altitude: #gps_elev# m  
HR: #gps_hr# bpm
```

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 offsetting the GPS track. You can use the `#file_datetime_now#` keyword to see it live in the output.

```
Video start time: 2021-02-15 13:20:47  
GPS start time: 2021-02-15 06:50:46
```


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.

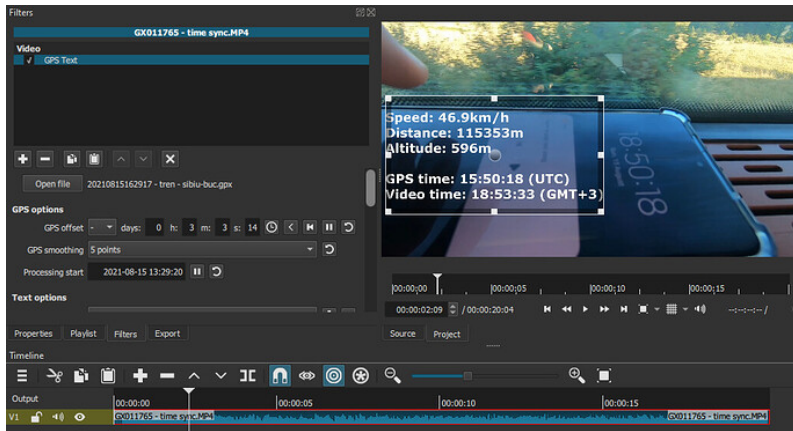
Sync button 4: Sync to now

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).



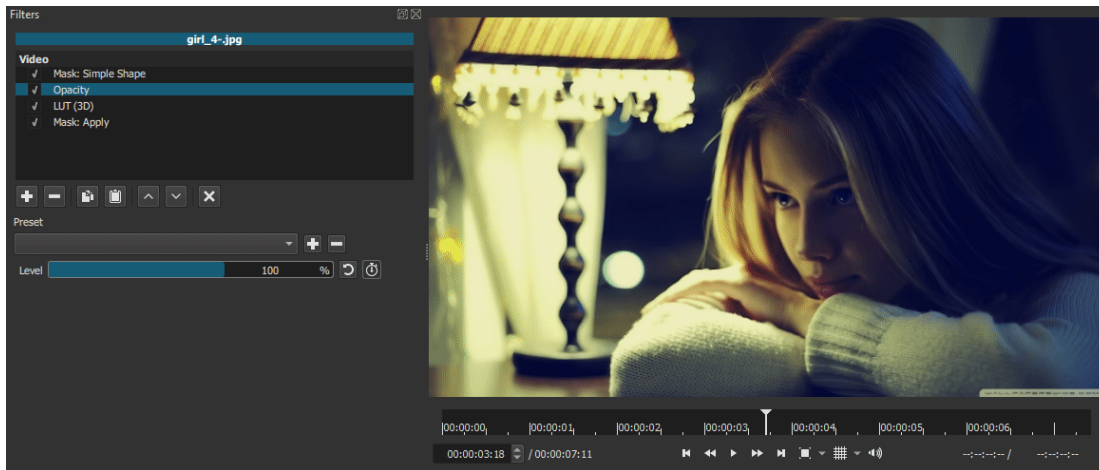
More info

[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).

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**




Mask: Apply Video Filter

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 goes after a sequence of other filters.

There are no parameters or UI for this filter. Also, it can only be used once per clip or track.

To really make this work you first add a [Mask: Simple Shape](#), [Mask: From File](#), or [Mask: Chroma Key](#) video filter to select the portion of the image you want to affect. Then, add one or more filters that apply some effect to that area. Finally, add **Mask: Apply**, 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:



You can also add video filters that further affect the alpha channel created by the initial mask creation filter. These go between the mask creation filter and **Mask: Apply**. For example:

- Add **Alpha: Adjust** to invert the mask and affect everything except the chosen color range or outside of the mask.
- Add **Mask: Simple Shape** along with **Operation: Minimum** to further limit the mask created by chroma key to a geometric-shaped (e.g. rectangle or circle) area of the video. Or use **Operation: Subtract** to exclude a certain area.
- Add the **Alpha Channel: View** video filter before **Mask: Apply** (or disable it) to better visualize what is going to be affected. Just do not forget to remove or disable **Alpha Channel: View** in order to see the final combined result.


You can include video filters before a mask creation filter to affect the entire image before the mask filters. And you can add other video filters after **Mask: Apply** that also applies to the entire image.

See also:


- [Mask: Chroma Key Video Filter](#)
- [Mask: From File](#)
- [Mask: Simple Shape](#)

Mask: Chroma Key Video Filter

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](#)



0:00



This is required to use motion tracking. 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.

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 NLMMeans 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!




Reframe Output Video Filter

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.

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.

The black bars you see in this embedded player are not part of the video!

Here it is on [YouTube Shorts](#)



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).

Size, Position & Rotate Video Filter

Location: Filters - Video - Size & Position.



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
- 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

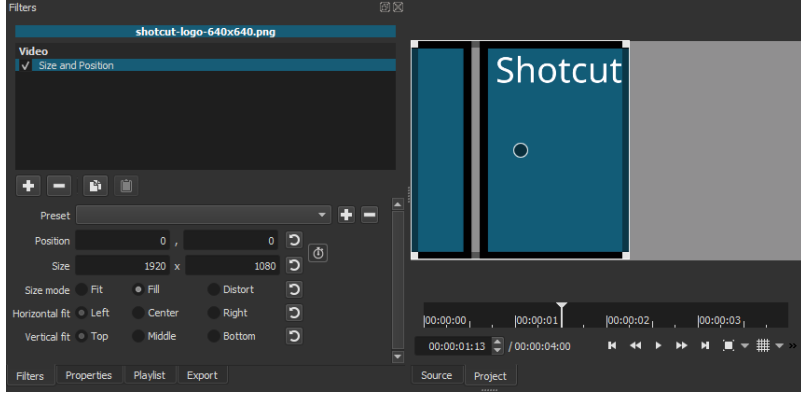
Vertical fit:

- 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.



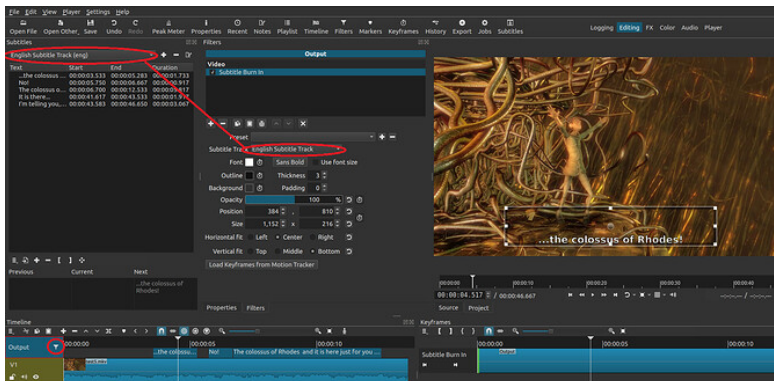
Subtitle Burn-In Video Filter

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 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 [strftime 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_norma
  <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</prop
  <property name="meta.media.0.stream.sample_aspect_ratio">
  <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</prop
  <property name="meta.media.0.codec.sample_aspect_ratio">.
  <property name="meta.media.0.codec.colorspace">709</prop
  <property name="meta.media.0.codec.name">h264</property>
  <property name="meta.media.0.codec.long_name">H.264 / AVC
  <property name="meta.media.0.codec.bit_rate">4001453</prop
```

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 \#.



White Balance Video Filter

The White Balance filter works well by following these steps:

Restore “Neutral color” and “Color temperature” to their defaults (6500 degrees).

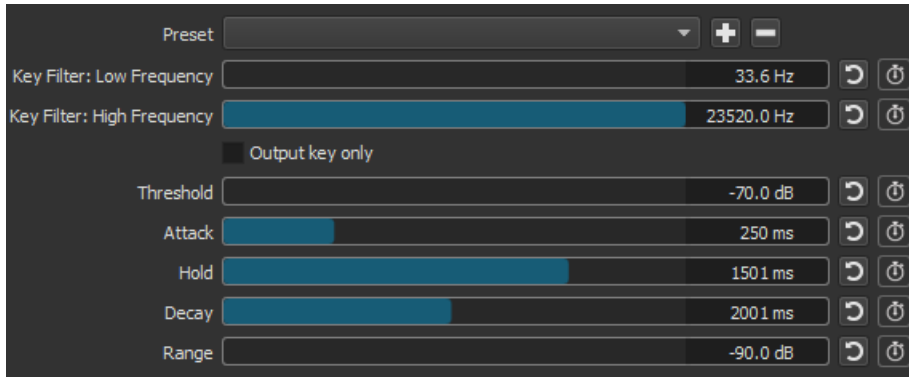
Click on the eyedropper and position the crosshairs over an object which is supposed to be white, e.g. a white card, white shirt, etc. and click.

The Downmix audio filter converts multichannel source audio to mono audio. This filter is simple in its operation. It adds all input channels together and sends the same signal out to all output channels. It is not aware of center, surround sound, or frequency channels. The most common use case is to mix stereo sources to mono to create two channel mono audio.

Noise Gate Audio Filter

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.

The **Pitch** filter lets you change how high or low the sound is with respect to frequency (not loudness).



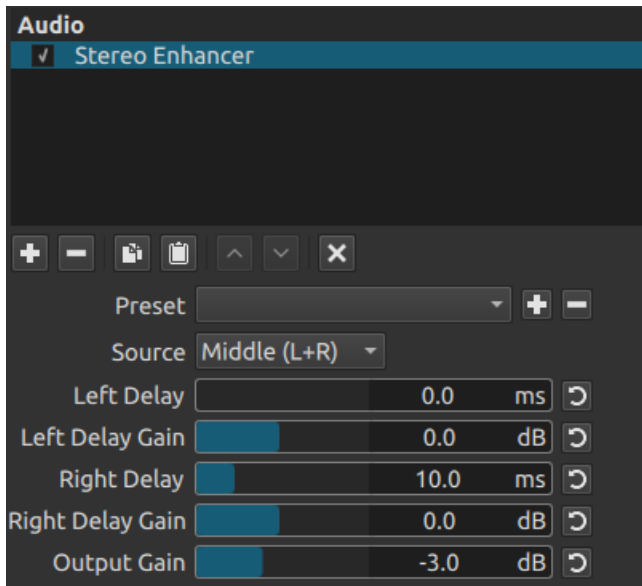
One may use this effect to disguise a voice or create a silly sounding voice. It is very useful to use in conjunction with the **Speed** option in **Properties**, which inherently changes the pitch. Thus, this filter can be used to correct the pitch. You can simply copy the speed value from Properties 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).

Stereo Enhancer Audio Filter

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.



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.

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.


Declick Audio Filter

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.

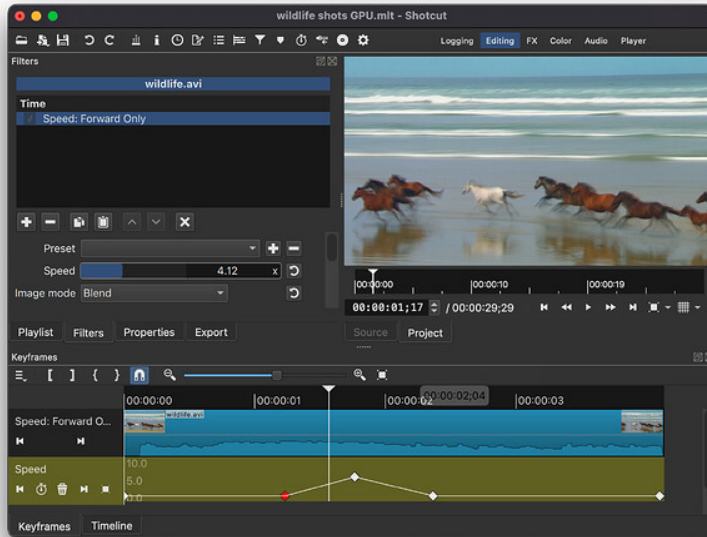
Here are the parameters:



- **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.

Speed: Forward & Reverse Time Filter

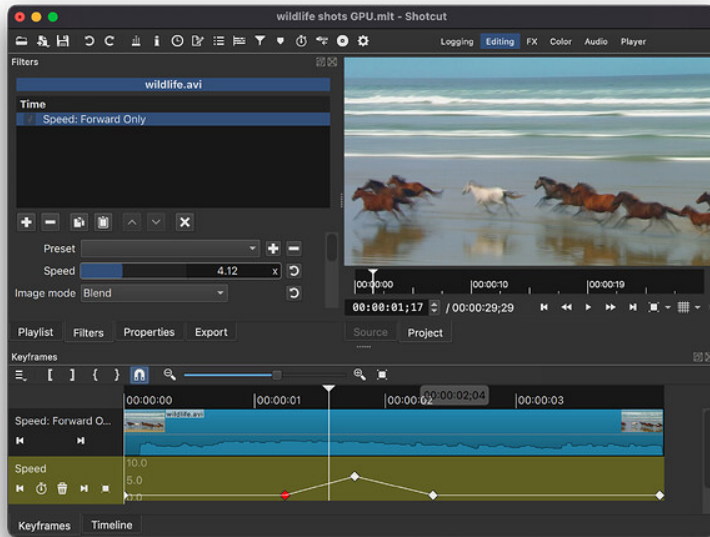
This filter is similar to **Time Remap** but can be easier to use.



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

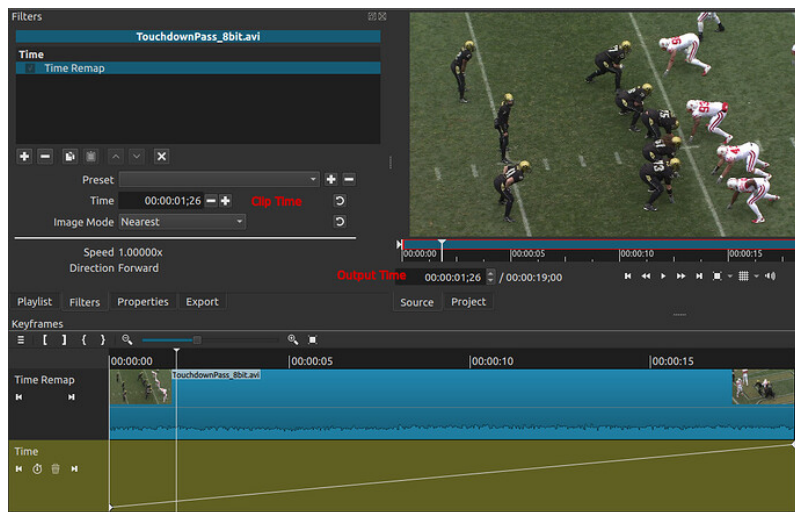
TL;DR

Filters > Time > Time Remap > Keyframes

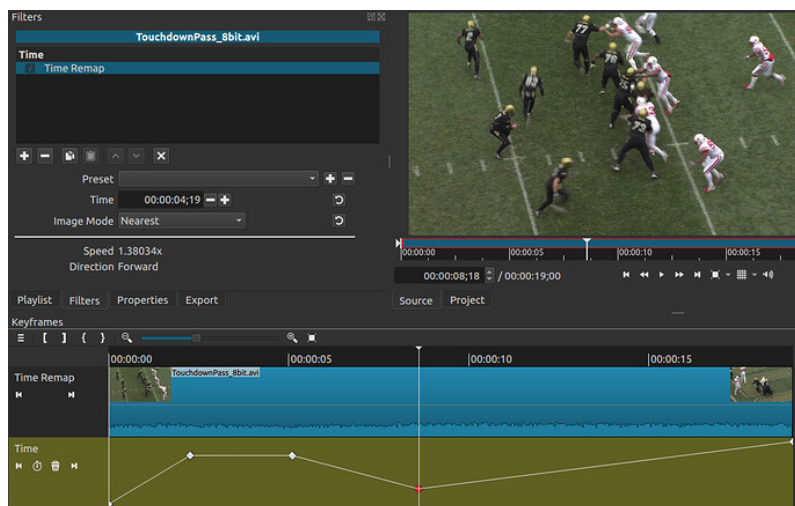
Hold **Ctrl** to drag a keyframe vertical only or **Alt** to drag horizontal only except on macOS these are **⌘** to drag vertical only or **⌥** to drag horizontal only.

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 and 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.

Preset Locations: Filters

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_stereo360	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
speedForwardReverseSpeed	Speed: Forward & Reverse	Time
speedForward	Speed: Forward Only	Time
timeremap	Time Remap	Time