FFmpeg Devices Documentation

Table of Contents

- 1 Description
- 2 Device Options
- 3 Input Devices
 - 3.1 alsa
 - 3.1.1 Options
 - 3.2 avfoundation
 - 3.2.1 Options
 - 3.2.2 Examples
 - 3.3 bktr
 - 3.3.1 Options
 - \circ 3.4 decklink
 - 3.4.1 Options
 - 3.4.2 Examples
 - 3.5 kmsgrab
 - 3.5.1 Options
 - 3.5.2 Examples
 - 3.6 libndi_newtek
 - 3.6.1 Options
 - 3.6.2 Examples
 - \circ 3.7 dshow
 - 3.7.1 Options
 - 3.7.2 Examples
 - 3.8 fbdev
 - 3.8.1 Options
 - 3.9 gdigrab
 - 3.9.1 Options
 - 3.10 iec61883
 - 3.10.1 Options
 - 3.10.2 Examples
 - 3.11 jack
 - 3.11.1 Options
 - 3.12 lavfi
 - 3.12.1 Options
 - 3.12.2 Examples
 - 3.13 libcdio
 - 3.13.1 Options
 - 0 3.14 libdc1394
 - 3.15 openal
 - 3.15.1 Options

- 3.15.2 Examples
- 3.16 oss
 - 3.16.1 Options
- 3.17 pulse
 - 3.17.1 Options
 - 3.17.2 Examples
- 3.18 sndio
 - 3.18.1 Options
- 3.19 video4linux2, v4l2
 - 3.19.1 Options
- 3.20 vfwcap
 - 3.20.1 Options
- 3.21 x11grab
 - 3.21.1 Options
- 4 Output Devices
 - 0 4.1 alsa
 - 4.1.1 Examples
 - 4.2 caca
 - 4.2.1 Options
 - 4.2.2 Examples
 - 4.3 decklink
 - 4.3.1 Options
 - 4.3.2 Examples
 - 4.4 libndi_newtek
 - 4.4.1 Options
 - 4.4.2 Examples
 - 4.5 fbdev
 - 4.5.1 Options
 - 4.5.2 Examples
 - 4.6 opengl
 - 4.6.1 Options
 - 4.6.2 Examples
 - 0 4.7 oss
 - \circ 4.8 pulse
 - 4.8.1 Options
 - 4.8.2 Examples
 - 4.9 sdl
 - 4.9.1 Options
 - 4.9.2 Interactive commands
 - 4.9.3 Examples
 - 4.10 sndio
 - 0 4.11 xv
 - 4.11.1 Options

- 4.11.2 Examples
- 5 See Also
- 6 Authors

1 Description# TOC

This document describes the input and output devices provided by the libavdevice library.

2 Device Options# TOC

The libavdevice library provides the same interface as libavformat. Namely, an input device is considered like a demuxer, and an output device like a muxer, and the interface and generic device options are the same provided by libavformat (see the ffmpeg-formats manual).

In addition each input or output device may support so-called private options, which are specific for that component.

Options may be set by specifying *-option value* in the FFmpeg tools, or by setting the value explicitly in the device AVFormatContext options or using the libavutil/opt.h API for programmatic use.

3 Input Devices# TOC

Input devices are configured elements in FFmpeg which enable accessing the data coming from a multimedia device attached to your system.

When you configure your FFmpeg build, all the supported input devices are enabled by default. You can list all available ones using the configure option "–list-indevs".

You can disable all the input devices using the configure option "-disable-indevs", and selectively enable an input device using the option "-enable-indev=*INDEV*", or you can disable a particular input device using the option "-disable-indev=*INDEV*".

The option "-devices" of the ff* tools will display the list of supported input devices.

A description of the currently available input devices follows.

3.1 alsa# TOC

ALSA (Advanced Linux Sound Architecture) input device.

To enable this input device during configuration you need libasound installed on your system.

This device allows capturing from an ALSA device. The name of the device to capture has to be an ALSA card identifier.

An ALSA identifier has the syntax:

```
hw:CARD[,DEV[,SUBDEV]]
```

where the DEV and SUBDEV components are optional.

The three arguments (in order: *CARD*, *DEV*, *SUBDEV*) specify card number or identifier, device number and subdevice number (-1 means any).

To see the list of cards currently recognized by your system check the files /proc/asound/cards and /proc/asound/devices.

For example to capture with ffmpeg from an ALSA device with card id 0, you may run the command:

ffmpeg -f alsa -i hw:0 alsaout.wav

For more information see: http://www.alsa-project.org/alsa-doc/alsa-lib/pcm.html

3.1.1 Options# TOC

```
sample_rate
```

Set the sample rate in Hz. Default is 48000.

channels

Set the number of channels. Default is 2.

3.2 avfoundation# TOC

AVFoundation input device.

AVFoundation is the currently recommended framework by Apple for streamgrabbing on OSX \geq 10.7 as well as on iOS.

The input filename has to be given in the following syntax:

```
-i "[[VIDEO]:[AUDIO]]"
```

The first entry selects the video input while the latter selects the audio input. The stream has to be specified by the device name or the device index as shown by the device list. Alternatively, the video and/or audio input device can be chosen by index using the -video_device_index <INDEX> and/or -audio_device_index <INDEX>, overriding any device name or index given in the input filename.

All available devices can be enumerated by using -list_devices true, listing all device names and corresponding indices.

There are two device name aliases:

default

Select the AVFoundation default device of the corresponding type.

none

Do not record the corresponding media type. This is equivalent to specifying an empty device name or index.

3.2.1 Options# TOC

AVFoundation supports the following options:

```
-list_devices <TRUE | FALSE>
```

If set to true, a list of all available input devices is given showing all device names and indices.

```
-video_device_index <INDEX>
```

Specify the video device by its index. Overrides anything given in the input filename.

-audio_device_index <INDEX>

Specify the audio device by its index. Overrides anything given in the input filename.

-pixel_format <FORMAT>

Request the video device to use a specific pixel format. If the specified format is not supported, a list of available formats is given and the first one in this list is used instead. Available pixel formats are: monob, rgb555be, rgb555le, rgb565be, rgb565le, rgb565le, rgb24, bgr24, 0rgb, bgr0, 0bgr, rgb0, bgr48be, uyvy422, yuva444p, yuva444p16le, yuv444p, yuv422p16, yuv422p10, yuv444p10, yuv420p, nv12, yuyv422, gray

-framerate

Set the grabbing frame rate. Default is ntsc, corresponding to a frame rate of 30000/1001.

-video_size

Set the video frame size.

```
-capture_cursor
```

Capture the mouse pointer. Default is 0.

-capture_mouse_clicks

Capture the screen mouse clicks. Default is 0.

3.2.2 Examples# TOC

• Print the list of AVFoundation supported devices and exit:

```
$ ffmpeg -f avfoundation -list_devices true -i ""
```

- Record video from video device 0 and audio from audio device 0 into out.avi:
 - \$ ffmpeg -f avfoundation -i "0:0" out.avi
- Record video from video device 2 and audio from audio device 1 into out.avi:

```
$ ffmpeg -f avfoundation -video_device_index 2 -i ":1" out.avi
```

• Record video from the system default video device using the pixel format bgr0 and do not record any audio into out.avi:

```
$ ffmpeg -f avfoundation -pixel_format bgr0 -i "default:none" out.avi
```

3.3 bktr# TOC

BSD video input device.

3.3.1 Options# TOC

framerate

Set the frame rate.

video_size

Set the video frame size. Default is vga.

standard

Available values are:

'pal' 'ntsc' 'secam' 'paln' 'palm' 'ntscj'

3.4 decklink# TOC

The decklink input device provides capture capabilities for Blackmagic DeckLink devices.

To enable this input device, you need the Blackmagic DeckLink SDK and you need to configure with the appropriate --extra-cflags and --extra-ldflags. On Windows, you need to run the IDL files through widl.

DeckLink is very picky about the formats it supports. Pixel format of the input can be set with raw_format. Framerate and video size must be determined for your device with -list_formats 1. Audio sample rate is always 48 kHz and the number of channels can be 2, 8 or 16. Note that all audio channels are bundled in one single audio track.

3.4.1 Options# TOC

list_devices

If set to true, print a list of devices and exit. Defaults to false.

list_formats

If set to true, print a list of supported formats and exit. Defaults to false.

format_code <FourCC>

This sets the input video format to the format given by the FourCC. To see the supported values of your device(s) use list_formats. Note that there is a FourCC 'pal ' that can also be used as pal (3 letters). Default behavior is autodetection of the input video format, if the hardware supports it.

bm_v210

This is a deprecated option, you can use raw_format instead. If set to '1', video is captured in 10 bit v210 instead of uyvy422. Not all Blackmagic devices support this option.

raw_format

Set the pixel format of the captured video. Available values are:

```
'uyvy422'
'yuv422p10'
'argb'
'bgra'
'rgb10'
teletext_lines
```

If set to nonzero, an additional teletext stream will be captured from the vertical ancillary data. Both SD PAL (576i) and HD (1080i or 1080p) sources are supported. In case of HD sources, OP47 packets are decoded.

This option is a bitmask of the SD PAL VBI lines captured, specifically lines 6 to 22, and lines 318 to 335. Line 6 is the LSB in the mask. Selected lines which do not contain teletext information will be ignored. You can use the special all constant to select all possible lines, or standard to skip lines 6, 318 and 319, which are not compatible with all receivers.

For SD sources, ffmpeg needs to be compiled with --enable-libzvbi. For HD sources, on older (pre-4K) DeckLink card models you have to capture in 10 bit mode.

channels

Defines number of audio channels to capture. Must be '2', '8' or '16'. Defaults to '2'.

duplex_mode

Sets the decklink device duplex mode. Must be 'unset', 'half' or 'full'. Defaults to 'unset'.

video_input

Sets the video input source. Must be 'unset', 'sdi', 'hdmi', 'optical_sdi', 'component', 'composite' or 's_video'. Defaults to 'unset'.

audio_input

Sets the audio input source. Must be 'unset', 'embedded', 'aes_ebu', 'analog', 'analog_xlr', 'analog_rca' or 'microphone'. Defaults to 'unset'.

video_pts

Sets the video packet timestamp source. Must be 'video', 'audio', 'reference' or 'wallclock'. Defaults to 'video'.

audio_pts

Sets the audio packet timestamp source. Must be 'video', 'audio', 'reference' or 'wallclock'. Defaults to 'audio'.

draw_bars

If set to 'true', color bars are drawn in the event of a signal loss. Defaults to 'true'.

queue_size

Sets maximum input buffer size in bytes. If the buffering reaches this value, incoming frames will be dropped. Defaults to '1073741824'.

audio_depth

Sets the audio sample bit depth. Must be '16' or '32'. Defaults to '16'.

3.4.2 Examples# TOC

• List input devices:

ffmpeg -f decklink -list_devices 1 -i dummy

• List supported formats:

ffmpeg -f decklink -list_formats 1 -i 'Intensity Pro'

• Capture video clip at 1080i50:

ffmpeg -format_code Hi50 -f decklink -i 'Intensity Pro' -c:a copy -c:v copy output.avi

• Capture video clip at 1080i50 10 bit:

ffmpeg -bm_v210 1 -format_code Hi50 -f decklink -i 'UltraStudio Mini Recorder' -c:a copy -c:v copy output.avi

• Capture video clip at 1080i50 with 16 audio channels:

ffmpeg -channels 16 -format_code Hi50 -f decklink -i 'UltraStudio Mini Recorder' -c:a copy -c:v copy output.avi

3.5 kmsgrab# TOC

KMS video input device.

Captures the KMS scanout framebuffer associated with a specified CRTC or plane as a DRM object that can be passed to other hardware functions.

Requires either DRM master or CAP_SYS_ADMIN to run.

If you don't understand what all of that means, you probably don't want this. Look at x11grab instead.

3.5.1 Options# TOC

device

DRM device to capture on. Defaults to /dev/dri/card0.

format

Pixel format of the framebuffer. Defaults to bgr0.

format_modifier

Format modifier to signal on output frames. This is necessary to import correctly into some APIs, but can't be autodetected. See the libdrm documentation for possible values.

crtc_id

KMS CRTC ID to define the capture source. The first active plane on the given CRTC will be used.

plane_id

KMS plane ID to define the capture source. Defaults to the first active plane found if neither crtc_id nor plane_id are specified.

framerate

Framerate to capture at. This is not synchronised to any page flipping or framebuffer changes - it just defines the interval at which the framebuffer is sampled. Sampling faster than the framebuffer update rate will generate independent frames with the same content. Defaults to 30.

3.5.2 Examples# TOC

• Capture from the first active plane, download the result to normal frames and encode. This will only work if the framebuffer is both linear and mappable - if not, the result may be scrambled or fail to download.

ffmpeg -f kmsgrab -i - -vf 'hwdownload,format=bgr0' output.mp4

• Capture from CRTC ID 42 at 60fps, map the result to VAAPI, convert to NV12 and encode as H.264.

ffmpeg -crtc_id 42 -framerate 60 -f kmsgrab -i - -vf 'hwmap=derive_device=vaapi,scale_vaapi=w=1920:h=1080:format=nv12' -c:v h264_vaapi output.mp4

3.6 libndi_newtek# TOC

The libndi_newtek input device provides capture capabilities for using NDI (Network Device Interface, standard created by NewTek).

Input filename is a NDI source name that could be found by sending -find_sources 1 to command line - it has no specific syntax but human-readable formatted.

To enable this input device, you need the NDI SDK and you need to configure with the appropriate --extra-cflags and --extra-ldflags.

3.6.1 Options# TOC

find_sources

If set to true, print a list of found/available NDI sources and exit. Defaults to false.

wait_sources

Override time to wait until the number of online sources have changed. Defaults to 0.5.

allow_video_fields

When this flag is false, all video that you receive will be progressive. Defaults to true.

3.6.2 Examples# TOC

• List input devices:

ffmpeg -f libndi_newtek -find_sources 1 -i dummy

• Restream to NDI:

ffmpeg -f libndi_newtek -i "DEV-5.INTERNAL.M1STEREO.TV (NDI_SOURCE_NAME_1)" -f libndi_newtek -y NDI_SOURCE_NAME_2

3.7 dshow# TOC

Windows DirectShow input device.

DirectShow support is enabled when FFmpeg is built with the mingw-w64 project. Currently only audio and video devices are supported.

Multiple devices may be opened as separate inputs, but they may also be opened on the same input, which should improve synchronism between them.

The input name should be in the format:

```
TYPE=NAME [: TYPE=NAME]
```

where TYPE can be either audio or video, and NAME is the device's name or alternative name..

3.7.1 Options# TOC

If no options are specified, the device's defaults are used. If the device does not support the requested options, it will fail to open.

video_size

Set the video size in the captured video.

```
framerate
```

Set the frame rate in the captured video.

```
sample_rate
```

Set the sample rate (in Hz) of the captured audio.

sample_size

Set the sample size (in bits) of the captured audio.

channels

Set the number of channels in the captured audio.

list_devices

If set to true, print a list of devices and exit.

list_options

If set to true, print a list of selected device's options and exit.

video_device_number

Set video device number for devices with the same name (starts at 0, defaults to 0).

audio_device_number

Set audio device number for devices with the same name (starts at 0, defaults to 0).

pixel_format

Select pixel format to be used by DirectShow. This may only be set when the video codec is not set or set to rawvideo.

audio_buffer_size

Set audio device buffer size in milliseconds (which can directly impact latency, depending on the device). Defaults to using the audio device's default buffer size (typically some multiple of 500ms). Setting this value too low can degrade performance. See also http://msdn.microsoft.com/en-us/library/windows/desktop/dd377582(v=vs.85).aspx

video_pin_name

Select video capture pin to use by name or alternative name.

audio_pin_name

Select audio capture pin to use by name or alternative name.

crossbar_video_input_pin_number

Select video input pin number for crossbar device. This will be routed to the crossbar device's Video Decoder output pin. Note that changing this value can affect future invocations (sets a new default) until system reboot occurs.

```
crossbar_audio_input_pin_number
```

Select audio input pin number for crossbar device. This will be routed to the crossbar device's Audio Decoder output pin. Note that changing this value can affect future invocations (sets a new default) until system reboot occurs.

show_video_device_dialog

If set to true, before capture starts, popup a display dialog to the end user, allowing them to change video filter properties and configurations manually. Note that for crossbar devices, adjusting values in this dialog may be needed at times to toggle between PAL (25 fps) and NTSC (29.97) input frame rates, sizes, interlacing, etc. Changing these values can enable different scan rates/frame rates and avoiding green bars at the bottom, flickering scan lines, etc. Note that with some devices, changing these properties can also affect future invocations (sets new defaults) until system reboot occurs.

show_audio_device_dialog

If set to true, before capture starts, popup a display dialog to the end user, allowing them to change audio filter properties and configurations manually.

show_video_crossbar_connection_dialog

If set to true, before capture starts, popup a display dialog to the end user, allowing them to manually modify crossbar pin routings, when it opens a video device.

```
show_audio_crossbar_connection_dialog
```

If set to true, before capture starts, popup a display dialog to the end user, allowing them to manually modify crossbar pin routings, when it opens an audio device.

```
show_analog_tv_tuner_dialog
```

If set to true, before capture starts, popup a display dialog to the end user, allowing them to manually modify TV channels and frequencies.

show_analog_tv_tuner_audio_dialog

If set to true, before capture starts, popup a display dialog to the end user, allowing them to manually modify TV audio (like mono vs. stereo, Language A,B or C).

```
audio_device_load
```

Load an audio capture filter device from file instead of searching it by name. It may load additional parameters too, if the filter supports the serialization of its properties to. To use this an audio capture source has to be specified, but it can be anything even fake one.

audio_device_save

Save the currently used audio capture filter device and its parameters (if the filter supports it) to a file. If a file with the same name exists it will be overwritten.

```
video_device_load
```

Load a video capture filter device from file instead of searching it by name. It may load additional parameters too, if the filter supports the serialization of its properties to. To use this a video capture source has to be specified, but it can be anything even fake one.

```
video_device_save
```

Save the currently used video capture filter device and its parameters (if the filter supports it) to a file. If a file with the same name exists it will be overwritten.

3.7.2 Examples# TOC

- Print the list of DirectShow supported devices and exit:
 - \$ ffmpeg -list_devices true -f dshow -i dummy
- Open video device *Camera*:
 - \$ ffmpeg -f dshow -i video="Camera"
- Open second video device with name *Camera*:

\$ ffmpeg -f dshow -video_device_number 1 -i video="Camera"

- Open video device *Camera* and audio device *Microphone*:
 - \$ ffmpeg -f dshow -i video="Camera":audio="Microphone"
- Print the list of supported options in selected device and exit:
 - \$ ffmpeg -list_options true -f dshow -i video="Camera"
- Specify pin names to capture by name or alternative name, specify alternative device name:
- Configure a crossbar device, specifying crossbar pins, allow user to adjust video capture properties at startup:

3.8 fbdev# TOC

Linux framebuffer input device.

The Linux framebuffer is a graphic hardware-independent abstraction layer to show graphics on a computer monitor, typically on the console. It is accessed through a file device node, usually /dev/fb0.

For more detailed information read the file Documentation/fb/framebuffer.txt included in the Linux source tree.

See also http://linux-fbdev.sourceforge.net/, and fbset(1).

To record from the framebuffer device /dev/fb0 with ffmpeg:

ffmpeg -f fbdev -framerate 10 -i /dev/fb0 out.avi

You can take a single screenshot image with the command:

ffmpeg -f fbdev -framerate 1 -i /dev/fb0 -frames:v 1 screenshot.jpeg

3.8.1 Options# TOC

framerate

Set the frame rate. Default is 25.

3.9 gdigrab# TOC

Win32 GDI-based screen capture device.

This device allows you to capture a region of the display on Windows.

There are two options for the input filename:

desktop

or

title=window_title

The first option will capture the entire desktop, or a fixed region of the desktop. The second option will instead capture the contents of a single window, regardless of its position on the screen.

For example, to grab the entire desktop using ffmpeg:

ffmpeg -f gdigrab -framerate 6 -i desktop out.mpg

Grab a 640x480 region at position 10, 20:

ffmpeg -f gdigrab -framerate 6 -offset_x 10 -offset_y 20 -video_size vga -i desktop out.mpg

Grab the contents of the window named "Calculator"

ffmpeg -f gdigrab -framerate 6 -i title=Calculator out.mpg

3.9.1 Options# TOC

draw_mouse

Specify whether to draw the mouse pointer. Use the value 0 to not draw the pointer. Default value is 1.

framerate

Set the grabbing frame rate. Default value is ntsc, corresponding to a frame rate of 30000/1001.

show_region

Show grabbed region on screen.

If *show_region* is specified with 1, then the grabbing region will be indicated on screen. With this option, it is easy to know what is being grabbed if only a portion of the screen is grabbed.

Note that *show_region* is incompatible with grabbing the contents of a single window.

For example:

```
ffmpeg -f gdigrab -show_region 1 -framerate 6 -video_size cif -offset_x 10 -offset_y 20 -i desktop out.mpg
```

video_size

Set the video frame size. The default is to capture the full screen if desktop is selected, or the full window size if title=window_title is selected.

offset_x

When capturing a region with *video_size*, set the distance from the left edge of the screen or desktop.

Note that the offset calculation is from the top left corner of the primary monitor on Windows. If you have a monitor positioned to the left of your primary monitor, you will need to use a negative $offset_x$ value to move the region to that monitor.

offset_y

When capturing a region with *video_size*, set the distance from the top edge of the screen or desktop.

Note that the offset calculation is from the top left corner of the primary monitor on Windows. If you have a monitor positioned above your primary monitor, you will need to use a negative *offset_y* value to move the region to that monitor.

3.10 iec61883# TOC

FireWire DV/HDV input device using libiec61883.

To enable this input device, you need libiec61883, libraw1394 and libavc1394 installed on your system. Use the configure option --enable-libiec61883 to compile with the device enabled.

The iec61883 capture device supports capturing from a video device connected via IEEE1394 (FireWire), using libiec61883 and the new Linux FireWire stack (juju). This is the default DV/HDV input method in Linux Kernel 2.6.37 and later, since the old FireWire stack was removed.

Specify the FireWire port to be used as input file, or "auto" to choose the first port connected.

3.10.1 Options# TOC

dvtype

Override autodetection of DV/HDV. This should only be used if auto detection does not work, or if usage of a different device type should be prohibited. Treating a DV device as HDV (or vice versa) will not work and result in undefined behavior. The values auto, dv and hdv are supported.

dvbuffer

Set maximum size of buffer for incoming data, in frames. For DV, this is an exact value. For HDV, it is not frame exact, since HDV does not have a fixed frame size.

dvguid

Select the capture device by specifying its GUID. Capturing will only be performed from the specified device and fails if no device with the given GUID is found. This is useful to select the input if multiple devices are connected at the same time. Look at /sys/bus/firewire/devices to find out the GUIDs.

3.10.2 Examples# TOC

• Grab and show the input of a FireWire DV/HDV device.

ffplay -f iec61883 -i auto

• Grab and record the input of a FireWire DV/HDV device, using a packet buffer of 100000 packets if the source is HDV.

ffmpeg -f iec61883 -i auto -hdvbuffer 100000 out.mpg

3.11 jack# TOC

JACK input device.

To enable this input device during configuration you need libjack installed on your system.

A JACK input device creates one or more JACK writable clients, one for each audio channel, with name *client_name*:input_N, where *client_name* is the name provided by the application, and N is a number which identifies the channel. Each writable client will send the acquired data to the FFmpeg input device.

Once you have created one or more JACK readable clients, you need to connect them to one or more JACK writable clients.

To connect or disconnect JACK clients you can use the jack_connect and jack_disconnect programs, or do it through a graphical interface, for example with qjackctl.

To list the JACK clients and their properties you can invoke the command jack_lsp.

Follows an example which shows how to capture a JACK readable client with ffmpeg.

```
# Create a JACK writable client with name "ffmpeg".
$ ffmpeg -f jack -i ffmpeg -y out.wav
# Start the sample jack_metro readable client.
$ jack_metro -b 120 -d 0.2 -f 4000
# List the current JACK clients.
$ jack_lsp -c
system:capture_1
system:capture_2
system:playback_1
system:playback_2
ffmpeg:input_1
metro:120_bpm
# Connect metro to the ffmpeg writable client.
$ jack_connect metro:120_bpm ffmpeg:input_1
```

For more information read: http://jackaudio.org/

3.11.1 Options# TOC

channels

Set the number of channels. Default is 2.

3.12 lavfi# TOC

Libavfilter input virtual device.

This input device reads data from the open output pads of a libavfilter filtergraph.

For each filtergraph open output, the input device will create a corresponding stream which is mapped to the generated output. Currently only video data is supported. The filtergraph is specified through the option graph.

3.12.1 Options# TOC

graph

Specify the filtergraph to use as input. Each video open output must be labelled by a unique string of the form "outN", where N is a number starting from 0 corresponding to the mapped input stream generated by the device. The first unlabelled output is automatically assigned to the "out0" label, but all the others need to be specified explicitly.

The suffix "+subcc" can be appended to the output label to create an extra stream with the closed captions packets attached to that output (experimental; only for EIA-608 / CEA-708 for now). The subcc streams are created after all the normal streams, in the order of the corresponding stream. For example, if there is "out19+subcc", "out7+subcc" and up to "out42", the stream #43 is subcc for stream #7 and stream #44 is subcc for stream #19.

If not specified defaults to the filename specified for the input device.

graph_file

Set the filename of the filtergraph to be read and sent to the other filters. Syntax of the filtergraph is the same as the one specified by the option *graph*.

dumpgraph

Dump graph to stderr.

3.12.2 Examples# TOC

• Create a color video stream and play it back with ffplay:

ffplay -f lavfi -graph "color=c=pink [out0]" dummy

• As the previous example, but use filename for specifying the graph description, and omit the "out0" label:

ffplay -f lavfi color=c=pink

• Create three different video test filtered sources and play them:

ffplay -f lavfi -graph "testsrc [out0]; testsrc,hflip [out1]; testsrc,negate [out2]" test3

• Read an audio stream from a file using the amovie source and play it back with ffplay: ffplay -f lavfi "amovie=test.wav"

• Read an audio stream and a video stream and play it back with ffplay:

ffplay -f lavfi "movie=test.avi[out0];amovie=test.wav[out1]"

• Dump decoded frames to images and closed captions to a file (experimental):

ffmpeg -f lavfi -i "movie=test.ts[out0+subcc]" -map v frame%08d.png -map s -c copy -f rawvideo subcc.bin

3.13 libcdio# TOC

Audio-CD input device based on libcdio.

To enable this input device during configuration you need libcdio installed on your system. It requires the configure option --enable-libcdio.

This device allows playing and grabbing from an Audio-CD.

For example to copy with ffmpeg the entire Audio-CD in /dev/sr0, you may run the command:

ffmpeg -f libcdio -i /dev/sr0 cd.wav

3.13.1 Options# TOC

speed

Set drive reading speed. Default value is 0.

The speed is specified CD-ROM speed units. The speed is set through the libcdio cdio_cddap_speed_set function. On many CD-ROM drives, specifying a value too large will result in using the fastest speed.

paranoia_mode

Set paranoia recovery mode flags. It accepts one of the following values:

```
'disable'
'verify'
'overlap'
'neverskip'
'full'
```

Default value is 'disable'.

For more information about the available recovery modes, consult the paranoia project documentation.

3.14 libdc1394# TOC

IIDC1394 input device, based on libdc1394 and libraw1394.

Requires the configure option --enable-libdc1394.

3.15 openal# TOC

The OpenAL input device provides audio capture on all systems with a working OpenAL 1.1 implementation.

To enable this input device during configuration, you need OpenAL headers and libraries installed on your system, and need to configure FFmpeg with --enable-openal.

OpenAL headers and libraries should be provided as part of your OpenAL implementation, or as an additional download (an SDK). Depending on your installation you may need to specify additional flags via the --extra-cflags and --extra-ldflags for allowing the build system to locate the OpenAL headers and libraries.

An incomplete list of OpenAL implementations follows:

Creative

The official Windows implementation, providing hardware acceleration with supported devices and software fallback. See http://openal.org/.

OpenAL Soft

Portable, open source (LGPL) software implementation. Includes backends for the most common sound APIs on the Windows, Linux, Solaris, and BSD operating systems. See http://kcat.strangesoft.net/openal.html.

Apple

OpenAL is part of Core Audio, the official Mac OS X Audio interface. See http://developer.apple.com/technologies/mac/audio-and-video.html

This device allows one to capture from an audio input device handled through OpenAL.

You need to specify the name of the device to capture in the provided filename. If the empty string is provided, the device will automatically select the default device. You can get the list of the supported devices by using the option *list_devices*.

3.15.1 Options# TOC

channels

Set the number of channels in the captured audio. Only the values 1 (monaural) and 2 (stereo) are currently supported. Defaults to 2.

sample_size

Set the sample size (in bits) of the captured audio. Only the values 8 and 16 are currently supported. Defaults to 16.

sample_rate

Set the sample rate (in Hz) of the captured audio. Defaults to 44.1k.

list_devices

If set to true, print a list of devices and exit. Defaults to false.

3.15.2 Examples# TOC

Print the list of OpenAL supported devices and exit:

\$ ffmpeg -list_devices true -f openal -i dummy out.ogg

Capture from the OpenAL device DR-BT101 via PulseAudio:

\$ ffmpeg -f openal -i 'DR-BT101 via PulseAudio' out.ogg

Capture from the default device (note the empty string " as filename):

\$ ffmpeg -f openal -i '' out.ogg

Capture from two devices simultaneously, writing to two different files, within the same ffmpeg command:

\$ ffmpeg -f openal -i 'DR-BT101 via PulseAudio' out1.ogg -f openal -i 'ALSA Default' out2.ogg

Note: not all OpenAL implementations support multiple simultaneous capture - try the latest OpenAL Soft if the above does not work.

3.16 oss# TOC

Open Sound System input device.

The filename to provide to the input device is the device node representing the OSS input device, and is usually set to /dev/dsp.

For example to grab from /dev/dsp using ffmpeg use the command:

ffmpeg -f oss -i /dev/dsp /tmp/oss.wav

For more information about OSS see: http://manuals.opensound.com/usersguide/dsp.html

3.16.1 Options# TOC

```
sample_rate
```

Set the sample rate in Hz. Default is 48000.

channels

Set the number of channels. Default is 2.

3.17 pulse# TOC

PulseAudio input device.

To enable this output device you need to configure FFmpeg with --enable-libpulse.

The filename to provide to the input device is a source device or the string "default"

To list the PulseAudio source devices and their properties you can invoke the command pactl list sources.

More information about PulseAudio can be found on http://www.pulseaudio.org.

3.17.1 Options# TOC

server

Connect to a specific PulseAudio server, specified by an IP address. Default server is used when not provided.

name

Specify the application name PulseAudio will use when showing active clients, by default it is the LIBAVFORMAT_IDENT string.

```
stream_name
```

Specify the stream name PulseAudio will use when showing active streams, by default it is "record".

sample_rate

Specify the samplerate in Hz, by default 48kHz is used.

channels

Specify the channels in use, by default 2 (stereo) is set.

frame_size

Specify the number of bytes per frame, by default it is set to 1024.

fragment_size

Specify the minimal buffering fragment in PulseAudio, it will affect the audio latency. By default it is unset.

wallclock

Set the initial PTS using the current time. Default is 1.

3.17.2 Examples# TOC

Record a stream from default device:

ffmpeg -f pulse -i default /tmp/pulse.wav

3.18 sndio# TOC

sndio input device.

To enable this input device during configuration you need libsndio installed on your system.

The filename to provide to the input device is the device node representing the sndio input device, and is usually set to /dev/audio0.

For example to grab from /dev/audio0 using ffmpeg use the command:

ffmpeg -f sndio -i /dev/audio0 /tmp/oss.wav

3.18.1 Options# TOC

sample_rate

Set the sample rate in Hz. Default is 48000.

channels

Set the number of channels. Default is 2.

3.19 video4linux2, v4l2# TOC

Video4Linux2 input video device.

"v4l2" can be used as alias for "video4linux2".

If FFmpeg is built with v4l-utils support (by using the --enable-libv4l2 configure option), it is possible to use it with the -use_libv4l2 input device option.

The name of the device to grab is a file device node, usually Linux systems tend to automatically create such nodes when the device (e.g. an USB webcam) is plugged into the system, and has a name of the kind /dev/videoN, where N is a number associated to the device.

Video4Linux2 devices usually support a limited set of *widthxheight* sizes and frame rates. You can check which are supported using -list_formats all for Video4Linux2 devices. Some devices, like TV cards, support one or more standards. It is possible to list all the supported standards using -list_standards all.

The time base for the timestamps is 1 microsecond. Depending on the kernel version and configuration, the timestamps may be derived from the real time clock (origin at the Unix Epoch) or the monotonic clock (origin usually at boot time, unaffected by NTP or manual changes to the clock). The -timestamps abs or -ts abs option can be used to force conversion into the real time clock.

Some usage examples of the video4linux2 device with ffmpeg and ffplay:

• List supported formats for a video4linux2 device:

ffplay -f video4linux2 -list_formats all /dev/video0

• Grab and show the input of a video4linux2 device:

ffplay -f video4linux2 -framerate 30 -video_size hd720 /dev/video0

• Grab and record the input of a video4linux2 device, leave the frame rate and size as previously set:

ffmpeg -f video4linux2 -input_format mjpeg -i /dev/video0 out.mpeg

For more information about Video4Linux, check http://linuxtv.org/.

3.19.1 Options# TOC

standard

Set the standard. Must be the name of a supported standard. To get a list of the supported standards, use the list_standards option.

channel

Set the input channel number. Default to -1, which means using the previously selected channel.

video_size

Set the video frame size. The argument must be a string in the form *WIDTHxHEIGHT* or a valid size abbreviation.

pixel_format

Select the pixel format (only valid for raw video input).

```
input_format
```

Set the preferred pixel format (for raw video) or a codec name. This option allows one to select the input format, when several are available.

framerate

Set the preferred video frame rate.

list_formats

List available formats (supported pixel formats, codecs, and frame sizes) and exit.

Available values are:

'all'

Show all available (compressed and non-compressed) formats.

'raw'

Show only raw video (non-compressed) formats.

'compressed'

Show only compressed formats.

list_standards

List supported standards and exit.

Available values are:

'all'

Show all supported standards.

timestamps, ts

Set type of timestamps for grabbed frames.

Available values are:

'default'

Use timestamps from the kernel.

'abs'

Use absolute timestamps (wall clock).

'mono2abs'

Force conversion from monotonic to absolute timestamps.

Default value is default.

use_libv412

Use libv4l2 (v4l-utils) conversion functions. Default is 0.

3.20 vfwcap# TOC

VfW (Video for Windows) capture input device.

The filename passed as input is the capture driver number, ranging from 0 to 9. You may use "list" as filename to print a list of drivers. Any other filename will be interpreted as device number 0.

3.20.1 Options# TOC

video_size

Set the video frame size.

framerate

Set the grabbing frame rate. Default value is ntsc, corresponding to a frame rate of 30000/1001.

3.21 x11grab# TOC

X11 video input device.

To enable this input device during configuration you need libxcb installed on your system. It will be automatically detected during configuration.

This device allows one to capture a region of an X11 display.

The filename passed as input has the syntax:

[hostname]:display_number.screen_number[+x_offset,y_offset]

hostname:display_number.screen_number specifies the X11 display name of the screen to grab from. *hostname* can be omitted, and defaults to "localhost". The environment variable DISPLAY contains the default display name.

 x_{offset} and y_{offset} specify the offsets of the grabbed area with respect to the top-left border of the X11 screen. They default to 0.

Check the X11 documentation (e.g. man X) for more detailed information.

Use the xdpyinfo program for getting basic information about the properties of your X11 display (e.g. grep for "name" or "dimensions").

For example to grab from : 0.0 using ffmpeg:

ffmpeg -f x11grab -framerate 25 -video_size cif -i :0.0 out.mpg

Grab at position 10, 20:

ffmpeg -f x11grab -framerate 25 -video_size cif -i :0.0+10,20 out.mpg

3.21.1 Options# TOC

draw_mouse

Specify whether to draw the mouse pointer. A value of 0 specifies not to draw the pointer. Default value is 1.

follow_mouse

Make the grabbed area follow the mouse. The argument can be centered or a number of pixels *PIXELS*.

When it is specified with "centered", the grabbing region follows the mouse pointer and keeps the pointer at the center of region; otherwise, the region follows only when the mouse pointer reaches within *PIXELS* (greater than zero) to the edge of region.

For example:

ffmpeg -f x11grab -follow_mouse centered -framerate 25 -video_size cif -i :0.0 out.mpg

To follow only when the mouse pointer reaches within 100 pixels to edge:

```
ffmpeg -f x11grab -follow_mouse 100 -framerate 25 -video_size cif -i :0.0 out.mpg
```

framerate

Set the grabbing frame rate. Default value is ntsc, corresponding to a frame rate of 30000/1001.

show_region

Show grabbed region on screen.

If *show_region* is specified with 1, then the grabbing region will be indicated on screen. With this option, it is easy to know what is being grabbed if only a portion of the screen is grabbed.

region_border

Set the region border thickness if -show_region 1 is used. Range is 1 to 128 and default is 3 (XCB-based x11grab only).

For example:

```
ffmpeg -f x11grab -show_region 1 -framerate 25 -video_size cif -i :0.0+10,20 out.mpg
```

With *follow_mouse*:

```
ffmpeg -f xllgrab -follow_mouse centered -show_region 1 -framerate 25 -video_size cif -i :0.0 out.mpg
```

video_size

Set the video frame size. Default value is vga.

```
grab_x
grab_y
```

Set the grabbing region coordinates. They are expressed as offset from the top left corner of the X11 window and correspond to the *x_offset* and *y_offset* parameters in the device name. The default value for both options is 0.

4 Output Devices# TOC

Output devices are configured elements in FFmpeg that can write multimedia data to an output device attached to your system.

When you configure your FFmpeg build, all the supported output devices are enabled by default. You can list all available ones using the configure option "–list-outdevs".

You can disable all the output devices using the configure option "-disable-outdevs", and selectively enable an output device using the option "-enable-outdev=*OUTDEV*", or you can disable a particular input device using the option "-disable-outdev=*OUTDEV*".

The option "-devices" of the ff* tools will display the list of enabled output devices.

A description of the currently available output devices follows.

4.1 alsa# TOC

ALSA (Advanced Linux Sound Architecture) output device.

4.1.1 Examples# TOC

• Play a file on default ALSA device:

ffmpeg -i INPUT -f alsa default

• Play a file on soundcard 1, audio device 7:

```
ffmpeg -i INPUT -f alsa hw:1,7
```

4.2 caca# TOC

CACA output device.

This output device allows one to show a video stream in CACA window. Only one CACA window is allowed per application, so you can have only one instance of this output device in an application.

To enable this output device you need to configure FFmpeg with --enable-libcaca. libcaca is a graphics library that outputs text instead of pixels.

For more information about libcaca, check: http://caca.zoy.org/wiki/libcaca

4.2.1 Options# TOC

```
window_title
```

Set the CACA window title, if not specified default to the filename specified for the output device.

window_size

Set the CACA window size, can be a string of the form *widthxheight* or a video size abbreviation. If not specified it defaults to the size of the input video.

driver

Set display driver.

algorithm

Set dithering algorithm. Dithering is necessary because the picture being rendered has usually far more colours than the available palette. The accepted values are listed with <code>-list_dither</code> algorithms.

antialias

Set antialias method. Antialiasing smoothens the rendered image and avoids the commonly seen staircase effect. The accepted values are listed with <code>-list_dither</code> antialiases.

charset

Set which characters are going to be used when rendering text. The accepted values are listed with -list_dither charsets.

color

Set color to be used when rendering text. The accepted values are listed with -list_dither colors.

list_drivers

If set to true, print a list of available drivers and exit.

list_dither

List available dither options related to the argument. The argument must be one of algorithms, antialiases, charsets, colors.

4.2.2 Examples# TOC

• The following command shows the ffmpeg output is an CACA window, forcing its size to 80x25:

ffmpeg -i INPUT -c:v rawvideo -pix_fmt rgb24 -window_size 80x25 -f caca -

• Show the list of available drivers and exit:

ffmpeg -i INPUT -pix_fmt rgb24 -f caca -list_drivers true -

• Show the list of available dither colors and exit:

ffmpeg -i INPUT -pix_fmt rgb24 -f caca -list_dither colors -

4.3 decklink# TOC

The decklink output device provides playback capabilities for Blackmagic DeckLink devices.

To enable this output device, you need the Blackmagic DeckLink SDK and you need to configure with the appropriate --extra-cflags and --extra-ldflags. On Windows, you need to run the IDL files through widl.

DeckLink is very picky about the formats it supports. Pixel format is always uyvy422, framerate, field order and video size must be determined for your device with *list_formats* 1. Audio sample rate is always 48 kHz.

4.3.1 Options# TOC

```
list_devices
```

If set to true, print a list of devices and exit. Defaults to false.

list_formats

If set to true, print a list of supported formats and exit. Defaults to false.

preroll

Amount of time to preroll video in seconds. Defaults to 0.5.

4.3.2 Examples# TOC

• List output devices:

ffmpeg -i test.avi -f decklink -list_devices 1 dummy

• List supported formats:

ffmpeg -i test.avi -f decklink -list_formats 1 'DeckLink Mini Monitor'

• Play video clip:

ffmpeg -i test.avi -f decklink -pix_fmt uyvy422 'DeckLink Mini Monitor'

• Play video clip with non-standard framerate or video size:

ffmpeg -i test.avi -f decklink -pix_fmt uyvy422 -s 720x486 -r 24000/1001 'DeckLink Mini Monitor'

4.4 libndi_newtek# TOC

The libndi_newtek output device provides playback capabilities for using NDI (Network Device Interface, standard created by NewTek).

Output filename is a NDI name.

To enable this output device, you need the NDI SDK and you need to configure with the appropriate --extra-cflags and --extra-ldflags.

NDI uses uyvy422 pixel format natively, but also supports bgra, bgr0, rgba and rgb0.

4.4.1 Options# TOC

reference_level

The audio reference level in dB. This specifies how many dB above the reference level (+4dBU) is the full range of 16 bit audio. Defaults to 0.

clock_video

These specify whether video "clock" themselves. Defaults to false.

clock_audio

These specify whether audio "clock" themselves. Defaults to false.

4.4.2 Examples# TOC

• Play video clip:

ffmpeg -i "udp://@239.1.1.1:10480?fifo_size=10000006overrun_nonfatal=1" -vf "scale=720:576,fps=fps=25,setdar=dar=16/9,format=pix_fmts=uyvy422" -f libndi_newtek NEW_NDI1

4.5 fbdev# TOC

Linux framebuffer output device.

The Linux framebuffer is a graphic hardware-independent abstraction layer to show graphics on a computer monitor, typically on the console. It is accessed through a file device node, usually /dev/fb0.

For more detailed information read the file Documentation/fb/framebuffer.txt included in the Linux source tree.

4.5.1 Options# TOC

```
xoffset
yoffset
```

Set x/y coordinate of top left corner. Default is 0.

4.5.2 Examples# TOC

Play a file on framebuffer device /dev/fb0. Required pixel format depends on current framebuffer settings.

```
ffmpeg -re -i INPUT -c:v rawvideo -pix_fmt bgra -f fbdev /dev/fb0
```

See also http://linux-fbdev.sourceforge.net/, and fbset(1).

4.6 opengl# TOC

OpenGL output device.

To enable this output device you need to configure FFmpeg with --enable-opengl.

This output device allows one to render to OpenGL context. Context may be provided by application or default SDL window is created.

When device renders to external context, application must implement handlers for following messages: AV_DEV_TO_APP_CREATE_WINDOW_BUFFER - create OpenGL context on current thread. AV_DEV_TO_APP_PREPARE_WINDOW_BUFFER - make OpenGL context current. AV_DEV_TO_APP_DISPLAY_WINDOW_BUFFER - swap buffers. AV_DEV_TO_APP_DESTROY_WINDOW_BUFFER - destroy OpenGL context. Application is also required to inform a device about current resolution by sending AV_APP_TO_DEV_WINDOW_SIZE message.

4.6.1 Options# TOC

background

Set background color. Black is a default.

no_window

Disables default SDL window when set to non-zero value. Application must provide OpenGL context and both window_size_cb and window_swap_buffers_cb callbacks when set.

```
window_title
```

Set the SDL window title, if not specified default to the filename specified for the output device. Ignored when no_window is set.

```
window_size
```

Set preferred window size, can be a string of the form widthxheight or a video size abbreviation. If not specified it defaults to the size of the input video, downscaled according to the aspect ratio. Mostly usable when no_window is not set.

4.6.2 Examples# TOC

Play a file on SDL window using OpenGL rendering:

```
ffmpeg -i INPUT -f opengl "window title"
```

4.7 oss# TOC

OSS (Open Sound System) output device.

4.8 pulse# TOC

PulseAudio output device.

To enable this output device you need to configure FFmpeg with --enable-libpulse.

More information about PulseAudio can be found on http://www.pulseaudio.org

4.8.1 Options# TOC

server

Connect to a specific PulseAudio server, specified by an IP address. Default server is used when not provided.

name

Specify the application name PulseAudio will use when showing active clients, by default it is the LIBAVFORMAT_IDENT string.

stream_name

Specify the stream name PulseAudio will use when showing active streams, by default it is set to the specified output name.

device

Specify the device to use. Default device is used when not provided. List of output devices can be obtained with command pactl list sinks.

buffer_size
buffer_duration

Control the size and duration of the PulseAudio buffer. A small buffer gives more control, but requires more frequent updates.

buffer_size specifies size in bytes while buffer_duration specifies duration in milliseconds.

When both options are provided then the highest value is used (duration is recalculated to bytes using stream parameters). If they are set to 0 (which is default), the device will use the default PulseAudio duration value. By default PulseAudio set buffer duration to around 2 seconds.

prebuf

Specify pre-buffering size in bytes. The server does not start with playback before at least prebuf bytes are available in the buffer. By default this option is initialized to the same value as buffer_size or buffer_duration (whichever is bigger).

minreq

Specify minimum request size in bytes. The server does not request less than minreq bytes from the client, instead waits until the buffer is free enough to request more bytes at once. It is recommended to not set this option, which will initialize this to a value that is deemed sensible by the server.

4.8.2 Examples# TOC

Play a file on default device on default server:

```
ffmpeg -i INPUT -f pulse "stream name"
```

4.9 sdl# TOC

SDL (Simple DirectMedia Layer) output device.

This output device allows one to show a video stream in an SDL window. Only one SDL window is allowed per application, so you can have only one instance of this output device in an application.

To enable this output device you need libsdl installed on your system when configuring your build.

For more information about SDL, check: http://www.libsdl.org/

4.9.1 Options# TOC

window_title

Set the SDL window title, if not specified default to the filename specified for the output device.

```
icon_title
```

Set the name of the iconified SDL window, if not specified it is set to the same value of window_title.

window_size

Set the SDL window size, can be a string of the form *widthxheight* or a video size abbreviation. If not specified it defaults to the size of the input video, downscaled according to the aspect ratio.

```
window_fullscreen
```

Set fullscreen mode when non-zero value is provided. Default value is zero.

4.9.2 Interactive commands# TOC

The window created by the device can be controlled through the following interactive commands.

q, ESC

Quit the device immediately.

4.9.3 Examples# TOC

The following command shows the ffmpeg output is an SDL window, forcing its size to the qcif format:

ffmpeg -i INPUT -c:v rawvideo -pix_fmt yuv420p -window_size qcif -f sdl "SDL output"

4.10 sndio# TOC

sndio audio output device.

4.11 xv# TOC

XV (XVideo) output device.

This output device allows one to show a video stream in a X Window System window.

4.11.1 Options# TOC

display_name

Specify the hardware display name, which determines the display and communications domain to be used.

The display name or DISPLAY environment variable can be a string in the format *hostname*[:*number*].

hostname specifies the name of the host machine on which the display is physically attached. *number* specifies the number of the display server on that host machine. *screen_number* specifies the screen to be used on that server.

If unspecified, it defaults to the value of the DISPLAY environment variable.

For example, dual-headed:0.1 would specify screen 1 of display 0 on the machine named "dual-headed".

Check the X11 specification for more detailed information about the display name format.

window_id

When set to non-zero value then device doesn't create new window, but uses existing one with provided *window_id*. By default this options is set to zero and device creates its own window.

window_size

Set the created window size, can be a string of the form *widthxheight* or a video size abbreviation. If not specified it defaults to the size of the input video. Ignored when *window_id* is set.

window_x window_y

Set the X and Y window offsets for the created window. They are both set to 0 by default. The values may be ignored by the window manager. Ignored when *window_id* is set.

```
window_title
```

Set the window title, if not specified default to the filename specified for the output device. Ignored when *window_id* is set.

For more information about XVideo see http://www.x.org/.

4.11.2 Examples# TOC

• Decode, display and encode video input with ffmpeg at the same time:

ffmpeg -i INPUT OUTPUT -f xv display

• Decode and display the input video to multiple X11 windows:

ffmpeg -i INPUT -f xv normal -vf negate -f xv negated

5 See Also# TOC

ffmpeg, ffplay, ffprobe, ffserver, libavdevice

6 Authors# TOC

The FFmpeg developers.

For details about the authorship, see the Git history of the project (git://source.ffmpeg.org/ffmpeg), e.g. by typing the command git log in the FFmpeg source directory, or browsing the online repository at http://source.ffmpeg.org.

Maintainers for the specific components are listed in the file MAINTAINERS in the source code tree.

This document was generated using makeinfo.