



Rocky Enterprise Linux 9.2 Manual Pages on command 'pulse-daemon.conf.5'

C:\>man pulse-daemon.conf.5

pulse-daemon.conf(5) File Formats Manual pulse-daemon.conf(5)

NAME

pulse-daemon.conf - PulseAudio daemon configuration file

SYNOPSIS

~/.config/pulse/daemon.conf

~/.config/pulse/daemon.conf.d/*.conf

/etc/pulse/daemon.conf

/etc/pulse/daemon.conf.d/*.conf

DESCRIPTION

The PulseAudio sound server reads configuration directives from a configuration file on startup. If the per-user file `~/.config/pulse/daemon.conf` exists, it is used, otherwise the system configuration file `/etc/pulse/daemon.conf` is used. In addition to those main files, configuration directives can also be put in files under directories `~/.config/pulse/daemon.conf.d/` and `/etc/pulse/daemon.conf.d/`. Those files have to have the `.conf` file name extension, but otherwise the file names can be chosen freely. The files under `daemon.conf.d` are processed in alphabetical order. In case the same option is set in multiple files, the last file to set an option overrides earlier files. The main `daemon.conf` file is processed first, so options set in files under `daemon.conf.d` override the main file.

Please note that the server also reads a configuration script on startup. See `default.pa(5)`.

The configuration file is a simple collection of variable declarations. If the con?

figuration file parser encounters either ; or # it ignores the rest of the line un?

til its end.

For the settings that take a boolean argument the values true, yes, on and 1 are equivalent, resp. false, no, off, 0.

GENERAL DIRECTIVES

daemonize= Daemonize after startup. Takes a boolean value, defaults to no. The --daemonize command line option takes precedence.

fail= Fail to start up if any of the directives in the configuration script de?

fault.pa fail. Takes a boolean argument, defaults to yes. The --fail command line option takes precedence.

allow-module-loading= Allow/disallow module loading after startup. This is a security feature that if disabled makes sure that no further modules may be loaded into the PulseAudio server after startup completed. It is recommended to disable this when system-instance is enabled. Please note that certain features like automatic hot-plug support will not work if this option is enabled. Takes a boolean argument, defaults to yes. The --disallow-module-loading command line option takes precedence.

allow-exit= Allow/disallow exit on user request. Defaults to yes.

resample-method= The resampling algorithm to use. Use one of src-sinc-best-quality, src-sinc-medium-quality, src-sinc-fastest, src-zero-order-hold, src-linear, trivial, speex-float-N, speex-fixed-N, ffmpeg, soxr-mq, soxr-hq, soxr-vhq. See the documentation of libsamplerate and speex for explanations of the different src- and speex- methods, respectively. The method trivial is the most basic algorithm implemented. If you're tight on CPU consider using this. On the other hand it has the worst quality of them all. The Speex resamplers take an integer quality setting in the range 0..10 (bad...good). They exist in two flavours: fixed and float. The former uses fixed point numbers, the latter relies on floating point numbers. On most desktop CPUs the float point resampler is a lot faster, and it also offers slightly better quality. The soxr-family methods are based on libsoxr, a resampler library from the SoX sound processing utility. The mq variant has the best performance of the three. The hq is more expensive and, according to SoX developers, is considered the best choice for audio of up to 16 bits per sample. The vhq variant has more precision than hq and is more suitable for larger samples. The Soxr resamplers gen?

erally offer better quality at less CPU compared to other resamplers, such as speex. The downside is that they can add a significant delay to the output (usually up to around 20 ms, in rare cases more). See the output of `dump-resample-methods` for a complete list of all available resamplers. Defaults to `speex-float-1`. The `--resample-method` command line option takes precedence. Note that some modules overwrite or allow overwriting of the resampler to use.

`avoid-resampling=` If set, try to configure the device to avoid resampling. This only works on devices which support reconfiguring their rate, and when no other streams are already playing or capturing audio. The device will also not be configured to a rate less than the default and alternate sample rates.

`enable-remixing=` If disabled never upmix or downmix channels to different channel maps. Instead, do a simple name-based matching only. Defaults to yes. There is no known valid use case for setting this option to no, therefore, this option is deprecated and may be removed in a future version of PulseAudio.

`remixing-use-all-sink-channels=` If enabled, use all sink channels when remixing. Otherwise, remix to the minimal set of sink channels needed to reproduce all of the source channels. (This has no effect on LFE remixing.) Defaults to yes.

`enable-lfe-remixing=` This is a way to set `remixing-produce-lfe` and `remixing-consume-lfe` to the same value at once. This option only exists for backward compatibility and may be removed in a future version of PulseAudio.

`remixing-produce-lfe=` If enabled, and the sink input does not have the LFE channel, synthesize the output LFE channel as a (lowpass-filtered, if `lfe-crossover-freq` is not 0) average of all input channels. Also, when `lfe-crossover-freq` is not 0, filter out low frequencies from other channels while producing a synthetic LFE output. If disabled, the output LFE channel will only get a signal when an input LFE channel is available as well. Defaults to no.

`remixing-consume-lfe=` If enabled, and the sink does not have an LFE channel, redirect the input LFE channel (if any) to other channels. If disabled, the input LFE channel will remain unused unless the sink has the LFE channel as well. Defaults to no.

`lfe-crossover-freq=` The crossover frequency (in Hz) for the LFE filter. Set it to 0 to disable the LFE filter. Defaults to 0.

`use-pid-file=` Create a PID file in the runtime directory (`$XDG_RUNTIME`

TIME_DIR/pulse/pid). If this is enabled you may use commands like --kill or --check. If you are planning to start more than one PulseAudio process per user, you better disable this option since it effectively disables multiple instances.

Takes a boolean argument, defaults to yes. The --use-pid-file command line option takes precedence.

cpu-limit= If disabled do not install the CPU load limiter, even on platforms where it is supported. This option is useful when debugging/profiling PulseAudio to disable disturbing SIGXCPU signals. Takes a boolean argument, defaults to no. The --no-cpu-limit command line argument takes precedence.

system-instance= Run the daemon as system-wide instance, requires root privileges. Takes a boolean argument, defaults to no. The --system command line argument takes precedence.

local-server-type= Please don't use this option if you don't have to! This option is currently only useful when you want D-Bus clients to use a remote server. This option may be removed in future versions. If you only want to run PulseAudio in the system mode, use the system-instance option. This option takes one of user, system or none as the argument. This is essentially a duplicate for the system-instance option. The difference is the none option, which is useful when you want to use a remote server with D-Bus clients. If both this and system-instance are defined, this option takes precedence. Defaults to whatever the system-instance is set.

enable-shm= Enable data transfer via POSIX or memfd shared memory. Takes a boolean argument, defaults to yes. The --disable-shm command line argument takes precedence.

enable-memfd=. Enable memfd shared memory. Takes a boolean argument, defaults to yes.

shm-size-bytes= Sets the shared memory segment size for the daemon, in bytes. If left unspecified or is set to 0 it will default to some system-specific default, usually 64 MiB. Please note that usually there is no need to change this value, unless you are running an OS kernel that does not do memory overcommit.

lock-memory= Locks the entire PulseAudio process into memory. While this might increase drop-out safety when used in conjunction with real-time scheduling this takes away a lot of memory from other processes and might hence considerably slow down your system. Defaults to no.

flat-volumes= Enable 'flat' volumes, i.e. where possible let the sink volume equal the maximum of the volumes of the inputs connected to it. Takes a boolean argument, defaults to no.

rescue-streams= Enable rescuing of streams if the used sink or source becomes unavailable. Takes a boolean argument. If set to yes, pulseaudio will try to move the streams from a sink or source that becomes unavailable to the default sink or source. If set to no, streams will be killed if the corresponding sink or source disappears. Defaults to yes.

SCHEDULING

high-priority= Renice the daemon after startup to become a high-priority process. This is a good idea if you experience drop-outs during playback. However, this is a certain security issue, since it works when called SUID root only, or RLIMIT_NICE is used. root is dropped immediately after gaining the nice level on startup, thus it is presumably safe. See pulseaudio(1) for more information. Takes a boolean argument, defaults to yes. The --high-priority command line option takes precedence.

realtime-scheduling= Try to acquire SCHED_FIFO scheduling for the IO threads. The same security concerns as mentioned above apply. However, if PA enters an endless loop, realtime scheduling causes a system lockup. Thus, realtime scheduling should only be enabled on trusted machines for now. Please note that only the IO threads of PulseAudio are made real-time. The controlling thread is left a normally scheduled thread. Thus enabling the high-priority option is orthogonal. See pulseaudio(1) for more information. Takes a boolean argument, defaults to yes. The --realtime command line option takes precedence.

realtime-priority= The realtime priority to acquire, if realtime-scheduling is enabled. Note: JACK uses 10 by default, 9 for clients. Thus it is recommended to choose the PulseAudio real-time priorities lower. Some PulseAudio threads might choose a priority a little lower or higher than the specified value. Defaults to 5.

nice-level= The nice level to acquire for the daemon, if high-priority is enabled. Note: on some distributions X11 uses -10 by default. Defaults to -11.

IDLE TIMES

exit-idle-time= Terminate the daemon after the last client quit and this time in seconds passed. Use a negative value to disable this feature. Defaults to 20. The --exit-idle-time command line option takes precedence.

When PulseAudio runs in the per-user mode and detects a login session, then any positive value will be reset to 0 so that PulseAudio will terminate immediately on logout. A positive value therefore has effect only in environments where there's no support for login session tracking. A negative value can still be used to disable any automatic exit.

When PulseAudio runs in the system mode, automatic exit is always disabled, so this option does nothing.

`scache-idle-time=` Unload autoloaded sample cache entries after being idle for this time in seconds. Defaults to 20. The `--scache-idle-time` command line option takes precedence.

PATHS

`dl-search-path=` The path where to look for dynamic shared objects (DSOs/plugins). You may specify more than one path separated by colons. The default path depends on compile time settings. The `--dl-search-path` command line option takes precedence.

`default-script-file=` The default configuration script file to load. Specify an empty string for not loading a default script file. The default behaviour is to load `~/.config/pulse/default.pa`, and if that file does not exist fall back to the system wide installed version `/etc/pulse/default.pa`. If run in system-wide mode the file `/etc/pulse/system.pa` is used instead. If `-n` is passed on the command line or `default-script-file=` is disabled the default configuration script is ignored.

`load-default-script-file=` Load the default configuration script file as specified in `default-script-file=`. Defaults to yes.

LOGGING

`log-target=` The default log target. Use either `stderr`, `syslog`, `journal` (optional), `auto`, `file:PATH` or `newfile:PATH`. On traditional systems `auto` is equivalent to `syslog`. On `systemd`-enabled systems, `auto` is equivalent to `journal`, in case `daemonize` is enabled, and to `stderr` otherwise. If set to `file:PATH`, logging is directed to the file indicated by `PATH`. `newfile:PATH` is otherwise the same as `file:PATH`, but existing files are never overwritten. If the specified file already exists, a suffix is added to the file name to avoid overwriting. Defaults to `auto`. The `--log-target` command line option takes precedence.

`log-level=` Log level, one of `debug`, `info`, `notice`, `warning`, `error`. Log messages with a lower log level than specified here are not logged. Defaults to `notice`. The

--log-level command line option takes precedence. The -v command line option might alter this setting.

log-meta= With each logged message log the code location the message was generated from. Defaults to no.

log-time= With each logged message log the relative time since startup. Defaults to no.

log-backtrace= When greater than 0, with each logged message log a code stack trace up the specified number of stack frames. Defaults to 0.

RESOURCE LIMITS

See `getrlimit(2)` for more information. Set to -1 if PulseAudio shall not touch the resource limit. Not all resource limits are available on all operating systems.

rlimit-as Defaults to -1.

rlimit-rss Defaults to -1.

rlimit-core Defaults to -1.

rlimit-data Defaults to -1.

rlimit-fsize Defaults to -1.

rlimit-nofile Defaults to 256.

rlimit-stack Defaults to -1.

rlimit-nproc Defaults to -1.

rlimit-locks Defaults to -1.

rlimit-sigpending Defaults to -1.

rlimit-msgqueue Defaults to -1.

rlimit-memlock Defaults to 16 KiB. Please note that the JACK client libraries may require more locked memory.

rlimit-nice Defaults to 31. Please make sure that the default nice level as configured with nice-level fits in this resource limit, if high-priority is enabled.

rlimit-rtprio Defaults to 9. Please make sure that the default real-time priority level as configured with realtime-priority= fits in this resource limit, if real-time-scheduling is enabled. The JACK client libraries require a real-time priority of 9 by default.

rlimit-rttime Defaults to 1000000.

DEFAULT DEVICE SETTINGS

Most drivers try to open the audio device with these settings and then fall back to

lower settings. The default settings are CD quality: 16bit native endian, 2 channels, 44100 Hz sampling.

default-sample-format= The default sampling format. Specify one of u8, s16le, s16be, s24le, s24be, s24-32le, s24-32be, s32le, s32be, float32le, float32be, ulaw, alaw. Depending on the endianness of the CPU the formats s16ne, s16re, s24ne, s24re, s24-32ne, s24-32re, s32ne, s32re, float32ne, float32re (for native, resp. reverse endian) are available as aliases.

default-sample-rate= The default sample frequency.

default-sample-channels The default number of channels.

default-channel-map The default channel map.

alternate-sample-rate The alternate sample frequency. Sinks and sources will use either the default-sample-rate value or this alternate value, typically 44.1 or 48kHz. Switching between default and alternate values is enabled only when the sinks/sources are suspended. This option is ignored in passthrough mode where the stream rate will be used. If set to the same value as the default sample rate, this feature is disabled.

DEFAULT FRAGMENT SETTINGS

Some hardware drivers require the hardware playback buffer to be subdivided into several fragments. It is possible to change these buffer metrics for machines with high scheduling latencies. Not all possible values that may be configured here are available in all hardware. The driver will find the nearest setting supported. Modern drivers that support timer-based scheduling ignore these options.

default-fragments= The default number of fragments. Defaults to 4.

default-fragment-size-msec=The duration of a single fragment. Defaults to 25ms (i.e. the total buffer is thus 100ms long).

DEFAULT DEFERRED VOLUME SETTINGS

With the flat volume feature enabled, the sink HW volume is set to the same level as the highest volume input stream. Any other streams (with lower volumes) have the appropriate adjustment applied in SW to bring them to the correct overall level. Sadly hardware mixer changes cannot be timed accurately and thus this change of volumes can sometimes cause the resulting output sound to be momentarily too loud or too soft. So to ensure SW and HW volumes are applied concurrently without any glitches, their application needs to be synchronized. The sink implementation needs

to support deferred volumes. The following parameters can be used to refine the process.

`enable-deferred-volume`= Enable deferred volume for the sinks that support it. This feature is enabled by default.

`deferred-volume-safety-margin-usec`= The amount of time (in usec) by which the HW volume increases are delayed and HW volume decreases are advanced. Defaults to 8000 usec.

`deferred-volume-extra-delay-usec`= The amount of time (in usec) by which HW volume changes are delayed. Negative values are also allowed. Defaults to 0.

AUTHORS

The PulseAudio Developers <pulseaudio-discuss (at) lists (dot) freedesktop (dot) org>; PulseAudio is available from <http://pulseaudio.org/>

SEE ALSO

`pulse-client.conf(5)`, `default.pa(5)`, `pulseaudio(1)`, `pacmd(1)`

Manuals

User

`pulse-daemon.conf(5)`