

NAME

systemd.timer - Timer unit configuration

SYNOPSIS

timer.timer

DESCRIPTION

A unit configuration file whose name ends in ".timer" encodes information about a timer controlled and supervised by systemd, for timer-based activation.

This man page lists the configuration options specific to this unit type. See [systemd.unit\(5\)](#) for the common options of all unit configuration files. The common configuration items are configured in the generic [Unit] and [Install] sections. The timer specific configuration options are configured in the [Timer] section.

For each timer file, a matching unit file must exist, describing the unit to activate when the timer elapses. By default, a service by the same name as the timer (except for the suffix) is activated. Example: a timer file `foo.timer` activates a matching service `foo.service`. The unit to activate may be controlled by *Unit=* (see below).

Note that in case the unit to activate is already active at the time the timer elapses it is not restarted, but simply left running. There is no concept of spawning new service instances in this case. Due to this, services with *RemainAfterExit=* set (which stay around continuously even after the service's main process exited) are usually not suitable for activation via repetitive timers, as they will only be activated once, and then stay around forever.

AUTOMATIC DEPENDENCIES

Timer units automatically gain a *Before=* dependency on the service they are supposed to activate.

Unless *DefaultDependencies=* in the "[Unit]" section is set to **false**, all timer units will implicitly have dependencies of type *Requires=* and *After=* on `sysinit.target`, a dependency of type *Before=* on `timers.target`, as well as *Conflicts=* and *Before=* on `shutdown.target` to ensure that they are stopped cleanly prior to system shutdown. Timer units with at least one *OnCalendar=* directive will have an additional *After=* dependency on `timer-sync.target` to avoid being started before the system clock has been correctly set. Only timer units involved with early boot or late system shutdown should disable the *DefaultDependencies=* option.

OPTIONS

Timer files must include a [Timer] section, which carries information about the timer it defines. The options specific to the [Timer] section of timer units are the following:

OnActiveSec=, *OnBootSec=*, *OnStartupSec=*, *OnUnitActiveSec=*, *OnUnitInactiveSec=*

Defines monotonic timers relative to different starting points: *OnActiveSec=* defines a timer relative to the moment the timer itself is activated. *OnBootSec=* defines a timer relative to when the machine was booted up. *OnStartupSec=* defines a timer relative to when systemd was first started.

OnUnitActiveSec= defines a timer relative to when the unit the timer is activating was last activated.

OnUnitInactiveSec= defines a timer relative to when the unit the timer is activating was last deactivated.

Multiple directives may be combined of the same and of different types. For example, by combining *OnBootSec=* and *OnUnitActiveSec=*, it is possible to define a timer that elapses in regular intervals and activates a specific service each time.

The arguments to the directives are time spans configured in seconds. Example: "`OnBootSec=50`" means 50s after boot-up. The argument may also include time units. Example: "`OnBootSec=5h 30min`" means 5 hours and 30 minutes after boot-up. For details about the syntax of time spans, see [systemd.time\(7\)](#).

If a timer configured with *OnBootSec=* or *OnStartupSec=* is already in the past when the timer unit is activated, it will immediately elapse and the configured unit is started. This is not the case for timers defined in the other directives.

These are monotonic timers, independent of wall-clock time and timezones. If the computer is temporarily suspended, the monotonic clock stops too.

If the empty string is assigned to any of these options, the list of timers is reset, and all prior assignments will have no effect.

Note that timers do not necessarily expire at the precise time configured with these settings, as they are subject to the *AccuracySec=* setting below.

OnCalendar=

Defines realtime (i.e. wallclock) timers with calendar event expressions. See [systemd.time\(7\)](#) for more information on the syntax of calendar event expressions. Otherwise, the semantics are similar to *OnActiveSec=* and related settings.

Note that timers do not necessarily expire at the precise time configured with this setting, as it is subject to the *AccuracySec=* setting below.

AccuracySec=

Specify the accuracy the timer shall elapse with. Defaults to 1min. The timer is scheduled to elapse within a time window starting with the time specified in *OnCalendar=*, *OnActiveSec=*, *OnBootSec=*, *OnStartupSec=*, *OnUnitActiveSec=* or *OnUnitInactiveSec=* and ending the time configured with *AccuracySec=* later. Within this time window, the expiry time will be placed at a host-specific, randomized, but stable position that is synchronized between all local timer units. This is done in order to optimize power consumption to suppress unnecessary CPU wake-ups. To get best accuracy, set this option to 1us. Note that the timer is still subject to the timer slack configured via [systemd-system.conf\(5\)](#)'s *TimerSlackNSec=* setting. See [prctl\(2\)](#) for details. To optimize power consumption, make sure to set this value as high as possible and as low as necessary.

RandomizedDelaySec=

Delay the timer by a randomly selected, evenly distributed amount of time between 0 and the specified time value. Defaults to 0, indicating that no randomized delay shall be applied. Each timer unit will determine this delay randomly each time it is started, and the delay will simply be added on top of the next determined elapsing time. This is useful to stretch dispatching of similarly configured timer events over a certain amount time, to avoid that they all fire at the same time, possibly resulting in resource congestion. Note the relation to *AccuracySec=* above: the latter allows the service manager to coalesce timer events within a specified time range in order to minimize wakeups, the former does the opposite: it stretches timer events over a time range, to make it unlikely that they fire simultaneously. If *RandomizedDelaySec=* and *AccuracySec=* are used in conjunction, first the randomized delay is added, and then the result is possibly further shifted to coalesce it with other timer events happening on the system. As mentioned above *AccuracySec=* defaults to 1min and *RandomizedDelaySec=* to 0, thus encouraging coalescing of timer events. In order to optimally stretch timer events over a certain range of time, make sure to set *RandomizedDelaySec=* to a higher value, and *AccuracySec=1us*.

Unit=

The unit to activate when this timer elapses. The argument is a unit name, whose suffix is not ".timer". If not specified, this value defaults to a service that has the same name as the timer unit, except for the suffix. (See above.) It is recommended that the unit name that is activated and the unit name of the timer unit are named identically, except for the suffix.

Persistent=

Takes a boolean argument. If true, the time when the service unit was last triggered is stored on disk. When the timer is activated, the service unit is triggered immediately if it would have been triggered at least once during the time when the timer was inactive. This is useful to catch up on missed runs of the service when the machine was off. Note that this setting only has an effect on timers configured with *OnCalendar=*. Defaults to *false*.

WakeSystem=

Takes a boolean argument. If true, an elapsing timer will cause the system to resume from suspend, should it be suspended and if the system supports this. Note that this option will only make sure the system resumes on the appropriate times, it will not take care of suspending it again after any work

that is to be done is finished. Defaults to *false*.

RemainAfterElapse=

Takes a boolean argument. If true, an elapsed timer will stay loaded, and its state remains queryable. If false, an elapsed timer unit that cannot elapse anymore is unloaded. Turning this off is particularly useful for transient timer units that shall disappear after they first elapse. Note that this setting has an effect on repeatedly starting a timer unit that only elapses once: if *RemainAfterElapse=* is on, it will not be started again, and is guaranteed to elapse only once. However, if *RemainAfterElapse=* is off, it might be started again if it is already elapsed, and thus be triggered multiple times. Defaults to *yes*.

SEE ALSO

[systemd\(1\)](#), [systemctl\(1\)](#), [systemd.unit\(5\)](#), [systemd.service\(5\)](#), [systemd.time\(7\)](#), [systemd.directives\(7\)](#), [systemd-system.conf\(5\)](#), [prctl\(2\)](#)