

**NAME**

systemd.service - Service unit configuration

**SYNOPSIS**

*service.service*

**DESCRIPTION**

A unit configuration file whose name ends in `.service` encodes information about a process controlled and supervised by `systemd`.

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 service specific configuration options are configured in the "[Service]" section.

Additional options are listed in [systemd.exec\(5\)](#), which define the execution environment the commands are executed in, and in [systemd.kill\(5\)](#), which define the way the processes of the service are terminated, and in [systemd.resource-control\(5\)](#), which configure resource control settings for the processes of the service.

If a service is requested under a certain name but no unit configuration file is found, `systemd` looks for a SysV init script by the same name (with the `.service` suffix removed) and dynamically creates a service unit from that script. This is useful for compatibility with SysV. Note that this compatibility is quite comprehensive but not 100%. For details about the incompatibilities, see the [Incompatibilities with SysV<sup>\[1\]</sup>](#) document.

**AUTOMATIC DEPENDENCIES**

Services with `Type=dbus` set automatically acquire dependencies of type `Requires=` and `After=` on `dbus.socket`.

Socket activated services are automatically ordered after their activating `.socket` units via an automatic `After=` dependency. Services also pull in all `.socket` units listed in `Sockets=` via automatic `Wants=` and `After=` dependencies.

Unless `DefaultDependencies=` in the "[Unit]" is set to **false**, service units will implicitly have dependencies of type `Requires=` and `After=` on `sysinit.target`, a dependency of type `After=` on `basic.target` as well as dependencies of type `Conflicts=` and `Before=` on `shutdown.target`. These ensure that normal service units pull in basic system initialization, and are terminated cleanly prior to system shutdown. Only services involved with early boot or late system shutdown should disable this option.

Instanced service units (i.e. service units with an "@" in their name) are assigned by default a per-template slice unit (see [systemd.slice\(5\)](#)), named after the template unit, containing all instances of the specific template. This slice is normally stopped at shutdown, together with all template instances. If that is not desired, set `DefaultDependencies=no` in the template unit, and either define your own per-template slice unit file that also sets `DefaultDependencies=no`, or set `Slice=system.slice` (or another suitable slice) in the template unit. Also see [systemd.resource-control\(5\)](#).

Additional implicit dependencies may be added as result of execution and resource control parameters as documented in [systemd.exec\(5\)](#) and [systemd.resource-control\(5\)](#).

**OPTIONS**

Service files must include a "[Service]" section, which carries information about the service and the process it supervises. A number of options that may be used in this section are shared with other unit types. These options are documented in [systemd.exec\(5\)](#) and [systemd.kill\(5\)](#). The options specific to the "[Service]" section of service units are the following:

`Type=`

Configures the process start-up type for this service unit. One of **simple**, **forking**, **oneshot**, **dbus**, **notify** or **idle**.

If set to **simple** (the default if neither `Type=` nor `BusName=`, but `ExecStart=` are specified), it is expected that the process configured with `ExecStart=` is the main process of the service. In this mode,

if the process offers functionality to other processes on the system, its communication channels should be installed before the daemon is started up (e.g. sockets set up by systemd, via socket activation), as systemd will immediately proceed starting follow-up units.

If set to **forking**, it is expected that the process configured with *ExecStart=* will call **fork()** as part of its start-up. The parent process is expected to exit when start-up is complete and all communication channels are set up. The child continues to run as the main daemon process. This is the behavior of traditional UNIX daemons. If this setting is used, it is recommended to also use the *PIDFile=* option, so that systemd can identify the main process of the daemon. systemd will proceed with starting follow-up units as soon as the parent process exits.

Behavior of **oneshot** is similar to **simple**; however, it is expected that the process has to exit before systemd starts follow-up units. *RemainAfterExit=* is particularly useful for this type of service. This is the implied default if neither *Type=* or *ExecStart=* are specified.

Behavior of **dbus** is similar to **simple**; however, it is expected that the daemon acquires a name on the D-Bus bus, as configured by *BusName=*. systemd will proceed with starting follow-up units after the D-Bus bus name has been acquired. Service units with this option configured implicitly gain dependencies on the *dbus.socket* unit. This type is the default if *BusName=* is specified.

Behavior of **notify** is similar to **simple**; however, it is expected that the daemon sends a notification message via **sd\_notify(3)** or an equivalent call when it has finished starting up. systemd will proceed with starting follow-up units after this notification message has been sent. If this option is used, *NotifyAccess=* (see below) should be set to open access to the notification socket provided by systemd. If *NotifyAccess=* is missing or set to **none**, it will be forcibly set to **main**. Note that currently *Type=notify* will not work if used in combination with *PrivateNetwork=yes*.

Behavior of **idle** is very similar to **simple**; however, actual execution of the service binary is delayed until all active jobs are dispatched. This may be used to avoid interleaving of output of shell services with the status output on the console. Note that this type is useful only to improve console output, it is not useful as a general unit ordering tool, and the effect of this service type is subject to a 5s time-out, after which the service binary is invoked anyway.

#### *RemainAfterExit=*

Takes a boolean value that specifies whether the service shall be considered active even when all its processes exited. Defaults to **no**.

#### *GuessMainPID=*

Takes a boolean value that specifies whether systemd should try to guess the main PID of a service if it cannot be determined reliably. This option is ignored unless **Type=forking** is set and *PIDFile=* is unset because for the other types or with an explicitly configured PID file, the main PID is always known. The guessing algorithm might come to incorrect conclusions if a daemon consists of more than one process. If the main PID cannot be determined, failure detection and automatic restarting of a service will not work reliably. Defaults to **yes**.

#### *PIDFile=*

Takes an absolute file name pointing to the PID file of this daemon. Use of this option is recommended for services where *Type=* is set to **forking**. systemd will read the PID of the main process of the daemon after start-up of the service. systemd will not write to the file configured here, although it will remove the file after the service has shut down if it still exists.

#### *BusName=*

Takes a D-Bus bus name that this service is reachable as. This option is mandatory for services where *Type=* is set to **dbus**.

#### *ExecStart=*

Commands with their arguments that are executed when this service is started. The value is split into zero or more command lines according to the rules described below (see section "Command Lines" below).

Unless *Type=* is **oneshot**, exactly one command must be given. When *Type=oneshot* is used, zero or more commands may be specified. Commands may be specified by providing multiple command lines in the same directive, or alternatively, this directive may be specified more than once with the same effect. If the empty string is assigned to this option, the list of commands to start is reset, prior assignments of this option will have no effect. If no *ExecStart=* is specified, then the service must have *RemainAfterExit=yes* set.

For each of the specified commands, the first argument must be an absolute path to an executable. Optionally, if this file name is prefixed with "@", the second token will be passed as "argv[0]" to the executed process, followed by the further arguments specified. If the absolute filename is prefixed with "-", an exit code of the command normally considered a failure (i.e. non-zero exit status or abnormal exit due to signal) is ignored and considered success. If the absolute path is prefixed with "+" then it is executed with full privileges. "@", "-", and "+" may be used together and they can appear in any order.

If more than one command is specified, the commands are invoked sequentially in the order they appear in the unit file. If one of the commands fails (and is not prefixed with "-"), other lines are not executed, and the unit is considered failed.

Unless *Type=forking* is set, the process started via this command line will be considered the main process of the daemon.

*ExecStartPre=*, *ExecStartPost=*

Additional commands that are executed before or after the command in *ExecStart=*, respectively. Syntax is the same as for *ExecStart=*, except that multiple command lines are allowed and the commands are executed one after the other, serially.

If any of those commands (not prefixed with "-") fail, the rest are not executed and the unit is considered failed.

*ExecStart=* commands are only run after all *ExecStartPre=* commands that were not prefixed with a "-" exit successfully.

*ExecStartPost=* commands are only run after the service has started successfully, as determined by *Type=* (i.e. the process has been started for *Type=simple* or *Type=idle*, the process exits successfully for *Type=oneshot*, the initial process exits successfully for *Type=forking*, "READY=1" is sent for *Type=notify*, or the *BusName=* has been taken for *Type=dbus*).

Note that *ExecStartPre=* may not be used to start long-running processes. All processes forked off by processes invoked via *ExecStartPre=* will be killed before the next service process is run.

Note that if any of the commands specified in *ExecStartPre=*, *ExecStart=*, or *ExecStartPost=* fail (and are not prefixed with "-", see above) or time out before the service is fully up, execution continues with commands specified in *ExecStopPost=*, the commands in *ExecStop=* are skipped.

*ExecReload=*

Commands to execute to trigger a configuration reload in the service. This argument takes multiple command lines, following the same scheme as described for *ExecStart=* above. Use of this setting is optional. Specifier and environment variable substitution is supported here following the same scheme as for *ExecStart=*.

One additional, special environment variable is set: if known, *\$MAINPID* is set to the main process of the daemon, and may be used for command lines like the following:

```
/bin/kill -HUP $MAINPID
```

Note however that reloading a daemon by sending a signal (as with the example line above) is usually not a good choice, because this is an asynchronous operation and hence not suitable to order reloads of multiple services against each other. It is strongly recommended to set *ExecReload=* to a command that not only triggers a configuration reload of the daemon, but also synchronously waits for it to complete.

*ExecStop=*

Commands to execute to stop the service started via *ExecStart=*. This argument takes multiple command lines, following the same scheme as described for *ExecStart=* above. Use of this setting is optional. After the commands configured in this option are run, all processes remaining for a service are terminated according to the *KillMode=* setting (see [systemd.kill\(5\)](#)). If this option is not specified, the process is terminated by sending the signal specified in *KillSignal=* when service stop is requested. Specifier and environment variable substitution is supported (including *\$MAINPID*, see above).

Note that it is usually not sufficient to specify a command for this setting that only asks the service to terminate (for example, by queuing some form of termination signal for it), but does not wait for it to do so. Since the remaining processes of the services are killed using **SIGKILL** immediately after the command exited, this would not result in a clean stop. The specified command should hence be a synchronous operation, not an asynchronous one.

Note that the commands specified in *ExecStop=* are only executed when the service started successfully first. They are not invoked if the service was never started at all, or in case its start-up failed, for example because any of the commands specified in *ExecStart=*, *ExecStartPre=* or *ExecStartPost=* failed (and weren't prefixed with "-", see above) or timed out. Use *ExecStopPost=* to invoke commands when a service failed to start up correctly and is shut down again.

It is recommended to use this setting for commands that communicate with the service requesting clean termination. When the commands specified with this option are executed it should be assumed that the service is still fully up and is able to react correctly to all commands. For post-mortem clean-up steps use *ExecStopPost=* instead.

*ExecStopPost=*

Additional commands that are executed after the service is stopped. This includes cases where the commands configured in *ExecStop=* were used, where the service does not have any *ExecStop=* defined, or where the service exited unexpectedly. This argument takes multiple command lines, following the same scheme as described for *ExecStart=*. Use of these settings is optional. Specifier and environment variable substitution is supported. Note that – unlike *ExecStop=* – commands specified with this setting are invoked when a service failed to start up correctly and is shut down again.

It is recommended to use this setting for clean-up operations that shall be executed even when the service failed to start up correctly. Commands configured with this setting need to be able to operate even if the service failed starting up half-way and left incompletely initialized data around. As the service's processes have been terminated already when the commands specified with this setting are executed they should not attempt to communicate with them.

Note that all commands that are configured with this setting are invoked with the result code of the service, as well as the main process' exit code and status, set in the *\$SERVICE\_RESULT*, *\$EXIT\_CODE* and *\$EXIT\_STATUS* environment variables, see [systemd.exec\(5\)](#) for details.

*RestartSec=*

Configures the time to sleep before restarting a service (as configured with *Restart=*). Takes a unit-less value in seconds, or a time span value such as "5min 20s". Defaults to 100ms.

*TimeoutStartSec=*

Configures the time to wait for start-up. If a daemon service does not signal start-up completion within the configured time, the service will be considered failed and will be shut down again. Takes a unit-less value in seconds, or a time span value such as "5min 20s". Pass "infinity" to disable the timeout logic. Defaults to *DefaultTimeoutStartSec=* from the manager configuration file, except when *Type=oneshot* is used, in which case the timeout is disabled by default (see [systemd-system.conf\(5\)](#)).

*TimeoutStopSec=*

Configures the time to wait for stop. If a service is asked to stop, but does not terminate in the specified time, it will be terminated forcibly via **SIGTERM**, and after another timeout of equal duration with **SIGKILL** (see *KillMode=* in [systemd.kill\(5\)](#)). Takes a unit-less value in seconds, or a

time span value such as "5min 20s". Pass "infinity" to disable the timeout logic. Defaults to *DefaultTimeoutStopSec*= from the manager configuration file (see [systemd-system.conf\(5\)](#)).

*TimeoutSec*=

A shorthand for configuring both *TimeoutStartSec*= and *TimeoutStopSec*= to the specified value.

*RuntimeMaxSec*=

Configures a maximum time for the service to run. If this is used and the service has been active for longer than the specified time it is terminated and put into a failure state. Note that this setting does not have any effect on *Type=oneshot* services, as they terminate immediately after activation completed. Pass "infinity" (the default) to configure no runtime limit.

*WatchdogSec*=

Configures the watchdog timeout for a service. The watchdog is activated when the start-up is completed. The service must call **sd\_notify(3)** regularly with "WATCHDOG=1" (i.e. the "keep-alive ping"). If the time between two such calls is larger than the configured time, then the service is placed in a failed state and it will be terminated with **SIGABRT**. By setting *Restart*= to **on-failure**, **on-watchdog**, **on-abnormal** or **always**, the service will be automatically restarted. The time configured here will be passed to the executed service process in the *WATCHDOG\_USEC*= environment variable. This allows daemons to automatically enable the keep-alive pinging logic if watchdog support is enabled for the service. If this option is used, *NotifyAccess*= (see below) should be set to open access to the notification socket provided by systemd. If *NotifyAccess*= is not set, it will be implicitly set to **main**. Defaults to 0, which disables this feature. The service can check whether the service manager expects watchdog keep-alive notifications. See **sd\_watchdog\_enabled(3)** for details. **sd\_event\_set\_watchdog(3)** may be used to enable automatic watchdog notification support.

*Restart*=

Configures whether the service shall be restarted when the service process exits, is killed, or a timeout is reached. The service process may be the main service process, but it may also be one of the processes specified with *ExecStartPre*=, *ExecStartPost*=, *ExecStop*=, *ExecStopPost*=, or *ExecReload*=. When the death of the process is a result of systemd operation (e.g. service stop or restart), the service will not be restarted. Timeouts include missing the watchdog "keep-alive ping" deadline and a service start, reload, and stop operation timeouts.

Takes one of **no**, **on-success**, **on-failure**, **on-abnormal**, **on-watchdog**, **on-abort**, or **always**. If set to **no** (the default), the service will not be restarted. If set to **on-success**, it will be restarted only when the service process exits cleanly. In this context, a clean exit means an exit code of 0, or one of the signals **SIGHUP**, **SIGINT**, **SIGTERM** or **SIGPIPE**, and additionally, exit statuses and signals specified in *SuccessExitStatus*=. If set to **on-failure**, the service will be restarted when the process exits with a non-zero exit code, is terminated by a signal (including on core dump, but excluding the aforementioned four signals), when an operation (such as service reload) times out, and when the configured watchdog timeout is triggered. If set to **on-abnormal**, the service will be restarted when the process is terminated by a signal (including on core dump, excluding the aforementioned four signals), when an operation times out, or when the watchdog timeout is triggered. If set to **on-abort**, the service will be restarted only if the service process exits due to an uncaught signal not specified as a clean exit status. If set to **on-watchdog**, the service will be restarted only if the watchdog timeout for the service expires. If set to **always**, the service will be restarted regardless of whether it exited cleanly or not, got terminated abnormally by a signal, or hit a timeout.

**Table 1. Exit causes and the effect of the *Restart*= settings on them**

Restart settings/Exit causes	no	always	on-success	on-failure	on-abnormal	on-abort	on-watchdog
Clean exit code or signal		X	X				
Unclean exit code		X		X			
Unclean signal		X		X	X	X	
Timeout		X		X	X		
Watchdog		X		X	X		X

As exceptions to the setting above, the service will not be restarted if the exit code or signal is specified in *RestartPreventExitStatus=* (see below). Also, the services will always be restarted if the exit code or signal is specified in *RestartForceExitStatus=* (see below).

Setting this to **on-failure** is the recommended choice for long-running services, in order to increase reliability by attempting automatic recovery from errors. For services that shall be able to terminate on their own choice (and avoid immediate restarting), **on-abnormal** is an alternative choice.

#### *SuccessExitStatus=*

Takes a list of exit status definitions that, when returned by the main service process, will be considered successful termination, in addition to the normal successful exit code 0 and the signals **SIGHUP**, **SIGINT**, **SIGTERM**, and **SIGPIPE**. Exit status definitions can either be numeric exit codes or termination signal names, separated by spaces. For example:

```
SuccessExitStatus=1 2 8 SIGKILL
```

ensures that exit codes 1, 2, 8 and the termination signal **SIGKILL** are considered clean service terminations.

Note that if a process has a signal handler installed and exits by calling `_exit(2)` in response to a signal, the information about the signal is lost. Programs should instead perform cleanup and kill themselves with the same signal instead. See [Proper handling of SIGINT/SIGQUIT — How to be a proper program](#)<sup>[2]</sup>.

This option may appear more than once, in which case the list of successful exit statuses is merged. If the empty string is assigned to this option, the list is reset, all prior assignments of this option will have no effect.

#### *RestartPreventExitStatus=*

Takes a list of exit status definitions that, when returned by the main service process, will prevent automatic service restarts, regardless of the restart setting configured with *Restart=*. Exit status definitions can either be numeric exit codes or termination signal names, and are separated by spaces. Defaults to the empty list, so that, by default, no exit status is excluded from the configured restart logic. For example:

```
RestartPreventExitStatus=1 6 SIGABRT
```

ensures that exit codes 1 and 6 and the termination signal **SIGABRT** will not result in automatic service restarting. This option may appear more than once, in which case the list of restart-preventing statuses is merged. If the empty string is assigned to this option, the list is reset and all prior assignments of this option will have no effect.

#### *RestartForceExitStatus=*

Takes a list of exit status definitions that, when returned by the main service process, will force automatic service restarts, regardless of the restart setting configured with *Restart=*. The argument format is similar to *RestartPreventExitStatus=*.

*PermissionsStartOnly=*

Takes a boolean argument. If true, the permission-related execution options, as configured with *User=* and similar options (see [systemd.exec\(5\)](#) for more information), are only applied to the process started with *ExecStart=*, and not to the various other *ExecStartPre=*, *ExecStartPost=*, *ExecReload=*, *ExecStop=*, and *ExecStopPost=* commands. If false, the setting is applied to all configured commands the same way. Defaults to false.

*RootDirectoryStartOnly=*

Takes a boolean argument. If true, the root directory, as configured with the *RootDirectory=* option (see [systemd.exec\(5\)](#) for more information), is only applied to the process started with *ExecStart=*, and not to the various other *ExecStartPre=*, *ExecStartPost=*, *ExecReload=*, *ExecStop=*, and *ExecStopPost=* commands. If false, the setting is applied to all configured commands the same way. Defaults to false.

*NonBlocking=*

Set the **O\_NONBLOCK** flag for all file descriptors passed via socket-based activation. If true, all file descriptors  $\geq 3$  (i.e. all except stdin, stdout, and stderr) will have the **O\_NONBLOCK** flag set and hence are in non-blocking mode. This option is only useful in conjunction with a socket unit, as described in [systemd.socket\(5\)](#). Defaults to false.

*NotifyAccess=*

Controls access to the service status notification socket, as accessible via the **sd\_notify(3)** call. Takes one of **none** (the default), **main** or **all**. If **none**, no daemon status updates are accepted from the service processes, all status update messages are ignored. If **main**, only service updates sent from the main process of the service are accepted. If **all**, all services updates from all members of the service's control group are accepted. This option should be set to open access to the notification socket when using *Type=notify* or *WatchdogSec=* (see above). If those options are used but *NotifyAccess=* is not configured, it will be implicitly set to **main**.

*Sockets=*

Specifies the name of the socket units this service shall inherit socket file descriptors from when the service is started. Normally, it should not be necessary to use this setting, as all socket file descriptors whose unit shares the same name as the service (subject to the different unit name suffix of course) are passed to the spawned process.

Note that the same socket file descriptors may be passed to multiple processes simultaneously. Also note that a different service may be activated on incoming socket traffic than the one which is ultimately configured to inherit the socket file descriptors. Or, in other words: the *Service=* setting of .socket units does not have to match the inverse of the *Sockets=* setting of the .service it refers to.

This option may appear more than once, in which case the list of socket units is merged. If the empty string is assigned to this option, the list of sockets is reset, and all prior uses of this setting will have no effect.

*FailureAction=*

Configure the action to take when the service enters a failed state. Takes the same values as the unit setting *StartLimitAction=* and executes the same actions (see [systemd.unit\(5\)](#)). Defaults to **none**.

*FileDescriptorStoreMax=*

Configure how many file descriptors may be stored in the service manager for the service using **sd\_pid\_notify\_with\_fds(3)**'s "FDSTORE=1" messages. This is useful for implementing service restart schemes where the state is serialized to /run and the file descriptors passed to the service manager, to allow restarts without losing state. Defaults to 0, i.e. no file descriptors may be stored in the service manager. All file descriptors passed to the service manager from a specific service are passed back to the service's main process on the next service restart. Any file descriptors passed to the service manager are automatically closed when POLLHUP or POLLERR is seen on them, or when the service is fully stopped and no job is queued or being executed for it.

*USBFunctionDescriptors=*

Configure the location of a file containing **USB FunctionFS**<sup>[3]</sup> descriptors, for implementation of

USB gadget functions. This is used only in conjunction with a socket unit with *ListenUSBFunction=* configured. The contents of this file are written to the ep0 file after it is opened.

*USBFunctionStrings=*

Configure the location of a file containing USB FunctionFS strings. Behavior is similar to *USBFunctionDescriptors=* above.

Check [systemd.exec\(5\)](#) and [systemd.kill\(5\)](#) for more settings.

## COMMAND LINES

This section describes command line parsing and variable and specifier substitutions for *ExecStart=*, *ExecStartPre=*, *ExecStartPost=*, *ExecReload=*, *ExecStop=*, and *ExecStopPost=* options.

Multiple command lines may be concatenated in a single directive by separating them with semicolons (these semicolons must be passed as separate words). Lone semicolons may be escaped as `"\";`.

Each command line is split on whitespace, with the first item being the command to execute, and the subsequent items being the arguments. Double quotes ("`...`") and single quotes ('`...`') may be used, in which case everything until the next matching quote becomes part of the same argument. C-style escapes are also supported. The table below contains the list of allowed escape patterns. Only patterns which match the syntax in the table are allowed; others will result in an error, and must be escaped by doubling the backslash. Quotes themselves are removed after parsing and escape sequences substituted. In addition, a trailing backslash ("`\`") may be used to merge lines.

This syntax is intended to be very similar to shell syntax, but only the meta-characters and expansions described in the following paragraphs are understood. Specifically, redirection using "`<`", "`<<`", "`>`", and "`>>`", pipes using "`|`", running programs in the background using "`&`", and *other elements of shell syntax are not supported*.

The command to execute must be an absolute path name. It may contain spaces, but control characters are not allowed.

The command line accepts "%" specifiers as described in [systemd.unit\(5\)](#). Note that the first argument of the command line (i.e. the program to execute) may not include specifiers.

Basic environment variable substitution is supported. Use "`${FOO}`" as part of a word, or as a word of its own, on the command line, in which case it will be replaced by the value of the environment variable including all whitespace it contains, resulting in a single argument. Use "\$FOO" as a separate word on the command line, in which case it will be replaced by the value of the environment variable split at whitespace, resulting in zero or more arguments. For this type of expansion, quotes are respected when splitting into words, and afterwards removed.

Example:

```
Environment="ONE=one" "TWO=two two"
ExecStart=/bin/echo $ONE $TWO ${TWO}
```

This will execute **/bin/echo** with four arguments: "one", "two", "two", and "two two".

Example:

```
Environment=ONE='one' "TWO='two two' too" THREE=
ExecStart=/bin/echo ${ONE} ${TWO} ${THREE}
ExecStart=/bin/echo $ONE $TWO $THREE
```

This results in echo being called twice, the first time with arguments "'one'", "'two two' too'", "", and the second time with arguments "one", "two two", "too".

To pass a literal dollar sign, use "\$\$". Variables whose value is not known at expansion time are treated as empty strings. Note that the first argument (i.e. the program to execute) may not be a variable.

Variables to be used in this fashion may be defined through *Environment=* and *EnvironmentFile=*. In addition, variables listed in the section "Environment variables in spawned processes" in [systemd.exec\(5\)](#), which are considered "static configuration", may be used (this includes e.g. `$USER`, but not `$TERM`).

Note that shell command lines are not directly supported. If shell command lines are to be used, they need to be passed explicitly to a shell implementation of some kind. Example:

```
ExecStart=/bin/sh -c 'dmesg | tac'
```

Example:

```
ExecStart=/bin/echo one ; /bin/echo "two two"
```

This will execute **/bin/echo** two times, each time with one argument: "one" and "two two", respectively. Because two commands are specified, *Type=oneshot* must be used.

Example:

```
ExecStart=/bin/echo / >/dev/null & \; \
/bin/ls
```

This will execute **/bin/echo** with five arguments: "/", ">/dev/null", "&", ";", and "/bin/ls".

**Table 2. C escapes supported in command lines and environment variables**

Literal	Actual value
"\a"	bell
"\b"	backspace
"\f"	form feed
"\n"	newline
"\r"	carriage return
"\t"	tab
"\v"	vertical tab
"\""	backslash
"\""	double quotation mark
"'"	single quotation mark
"\s"	space
"\xxx"	character number <i>xxx</i> in hexadecimal encoding
"\nnn"	character number <i>nnn</i> in octal encoding

## EXAMPLES

### Example 1. Simple service

The following unit file creates a service that will execute `/usr/sbin/foo-daemon`. Since no *Type=* is specified, the default *Type=***simple** will be assumed. `systemd` will assume the unit to be started immediately after the program has begun executing.

```
[Unit]
Description=Foo

[Service]
ExecStart=/usr/sbin/foo-daemon

[Install]
WantedBy=multi-user.target
```

Note that `systemd` assumes here that the process started by `systemd` will continue running until the service terminates. If the program daemonizes itself (i.e. forks), please use *Type=***forking** instead.

Since no *ExecStop=* was specified, `systemd` will send `SIGTERM` to all processes started from this service, and after a timeout also `SIGKILL`. This behavior can be modified, see [systemd.kill\(5\)](#) for details.

Note that this unit type does not include any type of notification when a service has completed

initialization. For this, you should use other unit types, such as *Type=notify* if the service understands systemd's notification protocol, *Type=forking* if the service can background itself or *Type=dbus* if the unit acquires a DBus name once initialization is complete. See below.

### Example 2. Oneshot service

Sometimes, units should just execute an action without keeping active processes, such as a filesystem check or a cleanup action on boot. For this, *Type=oneshot* exists. Units of this type will wait until the process specified terminates and then fall back to being inactive. The following unit will perform a cleanup action:

```
[Unit]
Description=Cleanup old Foo data

[Service]
Type=oneshot
ExecStart=/usr/sbin/foo-cleanup

[Install]
WantedBy=multi-user.target
```

Note that systemd will consider the unit to be in the state "starting" until the program has terminated, so ordered dependencies will wait for the program to finish before starting themselves. The unit will revert to the "inactive" state after the execution is done, never reaching the "active" state. That means another request to start the unit will perform the action again.

*Type=oneshot* are the only service units that may have more than one *ExecStart=* specified. They will be executed in order until either they are all successful or one of them fails.

### Example 3. Stoppable oneshot service

Similarly to the oneshot services, there are sometimes units that need to execute a program to set up something and then execute another to shut it down, but no process remains active while they are considered "started". Network configuration can sometimes fall into this category. Another use case is if a oneshot service shall not be executed each time when they are pulled in as a dependency, but only the first time.

For this, systemd knows the setting *RemainAfterExit=yes*, which causes systemd to consider the unit to be active if the start action exited successfully. This directive can be used with all types, but is most useful with *Type=oneshot* and *Type=simple*. With *Type=oneshot*, systemd waits until the start action has completed before it considers the unit to be active, so dependencies start only after the start action has succeeded. With *Type=simple*, dependencies will start immediately after the start action has been dispatched. The following unit provides an example for a simple static firewall.

```
[Unit]
Description=Simple firewall

[Service]
Type=oneshot
RemainAfterExit=yes
ExecStart=/usr/local/sbin/simple-firewall-start
ExecStop=/usr/local/sbin/simple-firewall-stop

[Install]
WantedBy=multi-user.target
```

Since the unit is considered to be running after the start action has exited, invoking **systemctl start** on that unit again will cause no action to be taken.

### Example 4. Traditional forking services

Many traditional daemons/services background (i.e. fork, daemonize) themselves when starting. Set *Type=forking* in the service's unit file to support this mode of operation. systemd will consider the service to be in the process of initialization while the original program is still running. Once it exits successfully

and at least a process remains (and *RemainAfterExit=no*), the service is considered started.

Often, a traditional daemon only consists of one process. Therefore, if only one process is left after the original process terminates, systemd will consider that process the main process of the service. In that case, the *\$MAINPID* variable will be available in *ExecReload=*, *ExecStop=*, etc.

In case more than one process remains, systemd will be unable to determine the main process, so it will not assume there is one. In that case, *\$MAINPID* will not expand to anything. However, if the process decides to write a traditional PID file, systemd will be able to read the main PID from there. Please set *PIDFile=* accordingly. Note that the daemon should write that file before finishing with its initialization. Otherwise, systemd might try to read the file before it exists.

The following example shows a simple daemon that forks and just starts one process in the background:

```
[Unit]
Description=Some simple daemon

[Service]
Type=forking
ExecStart=/usr/sbin/my-simple-daemon -d

[Install]
WantedBy=multi-user.target
```

Please see [systemd.kill\(5\)](#) for details on how you can influence the way systemd terminates the service.

#### Example 5. DBus services

For services that acquire a name on the DBus system bus, use *Type=dbus* and set *BusName=* accordingly. The service should not fork (daemonize). systemd will consider the service to be initialized once the name has been acquired on the system bus. The following example shows a typical DBus service:

```
[Unit]
Description=Simple DBus service

[Service]
Type=dbus
BusName=org.example.simple-dbus-service
ExecStart=/usr/sbin/simple-dbus-service

[Install]
WantedBy=multi-user.target
```

For *bus-activatable* services, do not include a "[Install]" section in the systemd service file, but use the *SystemdService=* option in the corresponding DBus service file, for example (*/usr/share/dbus-1/system-services/org.example.simple-dbus-service.service*):

```
[D-BUS Service]
Name=org.example.simple-dbus-service
Exec=/usr/sbin/simple-dbus-service
User=root
SystemdService=simple-dbus-service.service
```

Please see [systemd.kill\(5\)](#) for details on how you can influence the way systemd terminates the service.

#### Example 6. Services that notify systemd about their initialization

*Type=simple* services are really easy to write, but have the major disadvantage of systemd not being able to tell when initialization of the given service is complete. For this reason, systemd supports a simple notification protocol that allows daemons to make systemd aware that they are done initializing. Use *Type=notify* for this. A typical service file for such a daemon would look like this:

```
[Unit]
Description=Simple notifying service
```

```
[Service]
Type=notify
ExecStart=/usr/sbin/simple-notifying-service
```

```
[Install]
WantedBy=multi-user.target
```

Note that the daemon has to support systemd's notification protocol, else systemd will think the service has not started yet and kill it after a timeout. For an example of how to update daemons to support this protocol transparently, take a look at **sd\_notify(3)**. systemd will consider the unit to be in the 'starting' state until a readiness notification has arrived.

Please see **systemd.kill(5)** for details on how you can influence the way systemd terminates the service.

## SEE ALSO

[systemd\(1\)](#), [systemctl\(1\)](#), [systemd.unit\(5\)](#), [systemd.exec\(5\)](#), [systemd.resource-control\(5\)](#), [systemd.kill\(5\)](#), [systemd.directives\(7\)](#)

## NOTES

1. Incompatibilities with SysV  
<http://www.freedesktop.org/wiki/Software/systemd/Incompatibilities>
2. Proper handling of SIGINT/SIGQUIT — How to be a proper program  
<http://www.cons.org/cracauer/sigint.html>
3. USB FunctionFS  
<https://www.kernel.org/doc/Documentation/usb/functionfs.txt>