

NAME

systemd.mount - Mount unit configuration

SYNOPSIS

mount.mount

DESCRIPTION

A unit configuration file whose name ends in ".mount" encodes information about a file system mount point 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 mount specific configuration options are configured in the [Mount] section.

Additional options are listed in [systemd.exec\(5\)](#), which define the execution environment the [mount\(8\)](#) binary is executed in, and in [systemd.kill\(5\)](#), which define the way the processes are terminated, and in [systemd.resource-control\(5\)](#), which configure resource control settings for the processes of the service. Note that the User= and Group= options are not particularly useful for mount units specifying a "Type=" option or using configuration not specified in /etc/fstab; [mount\(8\)](#) will refuse options that are not listed in /etc/fstab if it is not run as UID 0.

Mount units must be named after the mount point directories they control. Example: the mount point /home/lennart must be configured in a unit file home-lennart.mount. For details about the escaping logic used to convert a file system path to a unit name, see [systemd.unit\(5\)](#). Note that mount units cannot be templated, nor is possible to add multiple names to a mount unit by creating additional symlinks to it.

Optionally, a mount unit may be accompanied by an automount unit, to allow on-demand or parallelized mounting. See [systemd.automount\(5\)](#).

Mount points created at runtime (independently of unit files or /etc/fstab) will be monitored by systemd and appear like any other mount unit in systemd. See /proc/self/mountinfo description in [proc\(5\)](#).

Some file systems have special semantics as API file systems for kernel-to-userspace and userspace-to-userspace interfaces. Some of them may not be changed via mount units, and cannot be disabled. For a longer discussion see [API File Systems](#)^[1].

AUTOMATIC DEPENDENCIES

If a mount unit is beneath another mount unit in the file system hierarchy, both a requirement dependency and an ordering dependency between both units are created automatically.

Block device backed file systems automatically gain *Bindsto=* and *After=* type dependencies on the device unit encapsulating the block device (see below).

If traditional file system quota is enabled for a mount unit, automatic *Wants=* and *Before=* dependencies on systemd-quotacheck.service and quotaon.service are added.

For mount units with *DefaultDependencies=yes* in the "[Unit]" section (the default) a couple additional dependencies are added. Mount units referring to local file systems automatically gain an *After=* dependency on local-fs-pre.target. Network mount units automatically acquire *After=* dependencies on remote-fs-pre.target, network.target and network-online.target. Towards the latter a *Wants=* unit is added as well. Mount units referring to local and network file systems are distinguished by their file system type specification. In some cases this is not sufficient (for example network block device based mounts, such as iSCSI), in which case *_netdev* may be added to the mount option string of the unit, which forces systemd to consider the mount unit a network mount. Mount units (regardless if local or network) also acquire automatic *Before=* and *Conflicts=* on umount.target in order to be stopped during shutdown.

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

FSTAB

Mount units may either be configured via unit files, or via /etc/fstab (see [fstab\(5\)](#) for details). Mounts listed in /etc/fstab will be converted into native units dynamically at boot and when the configuration of the

system manager is reloaded. In general, configuring mount points through `/etc/fstab` is the preferred approach. See [systemd-fstab-generator\(8\)](#) for details about the conversion.

The NFS mount option **bg** for NFS background mounts as documented in [nfs\(5\)](#) is not supported in `/etc/fstab` entries. The systemd mount option **nofail** provides similar functionality and should be used instead.

When reading `/etc/fstab` a few special mount options are understood by systemd which influence how dependencies are created for mount points. systemd will create a dependency of type *Wants=* or **Requires** (see option **nofail** below), from either `local-fs.target` or `remote-fs.target`, depending whether the file system is local or remote.

x-systemd.requires=

Configures a *Requires=* and an *After=* dependency between the created mount unit and another systemd unit, such as a device or mount unit. The argument should be a unit name, or an absolute path to a device node or mount point. This option may be specified more than once. This option is particularly useful for mount point declarations that need an additional device to be around (such as an external journal device for journal file systems) or an additional mount to be in place (such as an overlay file system that merges multiple mount points). See *After=* and *Requires=* in [systemd.unit\(5\)](#) for details.

x-systemd.requires-mounts-for=

Configures a *RequiresMountsFor=* dependency between the created mount unit and other mount units. The argument must be an absolute path. This option may be specified more than once. See *RequiresMountsFor=* in [systemd.unit\(5\)](#) for details.

x-systemd.automount

An automount unit will be created for the file system. See [systemd.automount\(5\)](#) for details.

x-systemd.idle-timeout=

Configures the idle timeout of the automount unit. See *TimeoutIdleSec=* in [systemd.automount\(5\)](#) for details.

x-systemd.device-timeout=

Configure how long systemd should wait for a device to show up before giving up on an entry from `/etc/fstab`. Specify a time in seconds or explicitly append a unit such as "s", "min", "h", "ms".

Note that this option can only be used in `/etc/fstab`, and will be ignored when part of the *Options=* setting in a unit file.

x-systemd.mount-timeout=

Configure how long systemd should wait for the mount command to finish before giving up on an entry from `/etc/fstab`. Specify a time in seconds or explicitly append a unit such as "s", "min", "h", "ms".

Note that this option can only be used in `/etc/fstab`, and will be ignored when part of the *Options=* setting in a unit file.

See *TimeoutSec=* below for details.

noauto, auto

With **noauto**, this mount will not be added as a dependency for `local-fs.target` or `remote-fs.target`. This means that it will not be mounted automatically during boot, unless it is pulled in by some other unit. The **auto** option has the opposite meaning and is the default.

nofail

With **nofail**, this mount will be only wanted, not required, by `local-fs.target` or `remote-fs.target`. This means that the boot will continue even if this mount point is not mounted successfully.

x-initrd.mount

An additional filesystem to be mounted in the `initramfs`. See `initrd-fs.target` description in [systemd.special\(7\)](#).

If a mount point is configured in both `/etc/fstab` and a unit file that is stored below `/usr`, the former will take precedence. If the unit file is stored below `/etc`, it will take precedence. This means: native unit files take precedence over traditional configuration files, but this is superseded by the rule that configuration in `/etc` will always take precedence over configuration in `/usr`.

OPTIONS

Mount files must include a `[Mount]` section, which carries information about the file system mount points 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 `[Mount]` section of mount units are the following:

What=

Takes an absolute path of a device node, file or other resource to mount. See [mount\(8\)](#) for details. If this refers to a device node, a dependency on the respective device unit is automatically created. (See [systemd.device\(5\)](#) for more information.) This option is mandatory.

Where=

Takes an absolute path of a directory of the mount point. If the mount point does not exist at the time of mounting, it is created. This string must be reflected in the unit filename. (See above.) This option is mandatory.

Type=

Takes a string for the file system type. See [mount\(8\)](#) for details. This setting is optional.

Options=

Mount options to use when mounting. This takes a comma-separated list of options. This setting is optional.

SloppyOptions=

Takes a boolean argument. If true, parsing of the options specified in *Options=* is relaxed, and unknown mount options are tolerated. This corresponds with [mount\(8\)](#)'s `-s` switch. Defaults to off.

LazyUnmount=

Takes a boolean argument. If true, detach the filesystem from the filesystem hierarchy at time of the unmount operation, and clean up all references to the filesystem as soon as they are not busy anymore. This corresponds with [umount\(8\)](#)'s `-l` switch. Defaults to off.

ForceUnmount=

Takes a boolean argument. If true, force an unmount (in case of an unreachable NFS system). This corresponds with [umount\(8\)](#)'s `-f` switch. Defaults to off.

DirectoryMode=

Directories of mount points (and any parent directories) are automatically created if needed. This option specifies the file system access mode used when creating these directories. Takes an access mode in octal notation. Defaults to `0755`.

TimeoutSec=

Configures the time to wait for the mount command to finish. If a command does not exit within the configured time, the mount will be considered failed and be shut down again. All commands still running will be terminated forcibly via **SIGTERM**, and after another delay of this time 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 0 to disable the timeout logic. The default value is set from the manager configuration file's *DefaultTimeoutStartSec=* variable.

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

SEE ALSO

[systemd\(1\)](#), [systemctl\(1\)](#), [systemd.unit\(5\)](#), [systemd.exec\(5\)](#), [systemd.kill\(5\)](#), [systemd.resource-control\(5\)](#), [systemd.service\(5\)](#), [systemd.device\(5\)](#), [proc\(5\)](#), [mount\(8\)](#), [systemd-fstab-generator\(8\)](#), [systemd.directives\(7\)](#)

NOTES

1. API File Systems
<http://www.freedesktop.org/wiki/Software/systemd/APIFileSystems>