

Filesystem Table (/etc/fstab) Cheatsheet

General Options

- **_netdev** - Network filesystem; mount when the network is available
- **_rnetdev** - Same as "**_netdev**", but the filesystem is checked during initialization
- **async** - Asynchronous read and write
- **atime** - Record last access time
- **auto** - Mount the filesystem on bootup
- **ctime** - Record last status change time
- **defaults** - The typical setting for the filesystem being edited; typically (on Linux systems) *rw, suid, dev, exec, auto, nouser, async, and relatime*
- **dev** - Recognize special block and character device files
- **diratime** - Record last directory access time
- **dirsync** - Synchronous directory updates; used for syscalls like *creat, link, unlink, symlink, mkdir, rmdir, mknod* and *rename*
- **exec** - Allow executables to be executed
- **group** - Allow a user in the same group as the device to mount the filesystem; implies *nosuid* and *nodev*
- **iversion** - Increment "*i_version*" each time an inode is accessed
- **mand** - Allow mandatory locks
- **mtime** - Record last modified time
- **noatime** - Do not record last access time
- **noauto** - Do not mount the filesystem on bootup
- **nodev** - Do not recognize special block device files
- **nodiratime** - Do not record last directory access time
- **noexec** - Do not permit executables to be executed
- **nofail** - Do not report errors if the device is not found
- **noiversion** - Do not increment "*i_version*"
- **nomand** - Do not allow mandatory locks
- **norelatime** - Do not use "*relatime*"
- **nostrictatime** - Do not use "*strictatime*"
- **nosuid** - Forbid the use of SUID and SGID (set-user-identifier and set-group-identifier bit)
- **nouser** - Do not permit regular users to mount the filesystem
- **owner** - Allow the specified owner to mount the filesystem (owner=bob); implies *nosuid* and *nodev*
- **relatime** - Last access time is updated if older than last modified time
- **remount** - Mount a mounted filesystem, but change the flags
- **ro** - Read-only
- **rw** - Read+Write
- **strictatime** - Record last access time; allow the userspace to override "*relatime*" and "*noatime*"
- **suid** - Allow the use of SUID and SGID
- **sync** - Synchronous read and write
- **user** - Any user can mount FS; implies *noexec, nosuid, and nodev*
- **users** - All users can mount and unmount the filesystem

Filesystem Specific Options

ADFS

- **gid** - Set group by GID; default "0"
- **othmask** - Set "other" permissions; default "0077"
- **ownmask** - Set "owner" permissions; default "0700"
- **uid** - Set owner by UID; default "0"

AFFS

- **bs** - Blocksize (512, 1024, 2048, or 4096)
- **gid** - Set group by GID; default "0"
- **grpquota** - Allow group quotas; recognized by quota software
- **mode** - Set the permissions of the files (i.e. "0777") and ignore previously set permissions
- **noquota** - Disallow quotas; recognized by quota software
- **prefix** - Prefix used in volume name
- **protect** - Do not change the protection bits
- **quota** - Allow quotas; recognized by quota software
- **reserved** - The number of unused blocks at the beginning of the device (reserved=2)
- **root** - Explicit location of the root-block
- **setgid** - Set the group of the files (by GID)
- **setuid** - Set the owner of the files (by UID)
- **uid** - Set owner by UID; default "0"
- **usemp** - Apply the UID and GID of the mountpoint to the filesystem
- **usrquota** - Allow user quotas; recognized by quota software
- **verbose** - Provide information about each mount
- **volume** - Prefix used before "/" when following soft-links (symbolic links)

CIFS

- **actimeo** - Set the time (in seconds) the client will use its attribute cache before requesting updated info
- **backupgid** - Only the specified group can access backup files
- **backupuid** - Only the specified user can access backup files
- **cache** - Set cache mode to none (no cache), loose (loose caching semantics), or strict (use the CIFS/SMB2 protocol)
- **cifsacl** - Allow mapping between Linux permission bits and CIFS/NTFS ACLs, SID and UID&GID, and Security Descriptors
- **credentials** - Location of the credentials file specifying the username, password, and domain; alternative to using "*password*" and "*user*"
- **cruid** - Set the owner of the credentials cache
- **directio** - Disallow inode caching; use cache=none instead
- **dir_mode** - The directory-mode to use if Unix CIFS Extensions are not supported
- **domain** - Set the workgroup of the user
- **dynperm** - The server will store permissions in memory
- **file_mode** - Use file-mode if CIFS Extensions are unsupported
- **forcegid** - Ignore the GID of the files specified by the server
- **forceuid** - Ignore the UID of the files specified by the server
- **fsc** - Use local-disk caching

Filesystem Table (/etc/fstab) Cheatsheet

- **gid** - GID of files and folders if the server does not provide a GID
- **guest** - Passwords are not required to access the server share
- **hard** - If the server crashes, programs accessing the FS will hang
- **ignorecase** - Alias for "nocase"; case-insensitive paths
- **intr** - UNUSED
- **iocharset** - Set the I/O character set (iocharset=utf8)
- **ip** - The IP address of the server containing the CIFS filesystem
- **mapchars** - Remap six of the seven reserved characters to the range above 0xF000
- **multiuser** - Allow multiple credentials, one for each user
- **netbiosname** - The client's NetBIOS name (when using port 139)
- **noacl** - Disallow POSIX ACL operations
- **nobri** - Do not send byte range lock requests to the server
- **nocase** - Case-insensitive paths
- **nointr** - UNUSED
- **nomapchars** - Do not remap the seven reserved characters
- **noperm** - Do not perform permission checks
- **noposixpaths** - Disallow the use of POSIX-style pathnames
- **noserverino** - Use client-generated inode numbers instead
- **nosetuids** - Do not apply the UID and GID to files and directories
- **nounix** - Disallow CIFS Unix Extensions
- **nouser_xattr** - Disallow extended attributes
- **password** - The password needed to access the share; if not set, then "\$PASSWORD" is used
- **perm** - Perform permission checks
- **port** - Use the specified port when connecting to the share
- **posixpaths** - Allow the use of POSIX-style pathnames
- **prefixpath** - Mount a subdirectory of a CIFS share
- **rsize** - Network read-size; never set larger than *CIFSMaxBufSize* (16kb)
- **rwpidforward** - IO operations of a file inherit the PID of the process that opened the file
- **sec** - Security Mode; many values are available
 - **krb5** - Kerberos v5 authentication
 - **krb5i** - Kerberos v5 authentication with packet signing
 - **none** - connect as null user
 - **ntlm** - NTLM password hashing
 - **ntlmi** - NTLM password hashing with packet signing
 - **ntlmssp** - NTLMv2 password hashing encapsulated in a Raw NTLMSSP message
 - **ntlmsspi** - NTLMv2 password hashing encapsulated in a Raw NTLMSSP message with packet signing
 - **ntlmv2** - NTLMv2 password hashing
 - **ntlmv2i** - NTLMv2 password hashing with packet signing
- **serverino** - Use the server's inode numbers rather than temporary client inode numbers
- **servern** - Alias for "servernetbiosname"
- **servernetbiosname** - Server netbios name (RFC1001 name)

- **setuids** - Try to apply the UID and GID to files and directories if Unix CIFS Extensions are not supported
- **sfu** - Create "Services for Unix (SFU)" compatible device files and fifos when the CIFS Unix Extensions are not negotiated
- **soft** - If the server crashes, programs accessing the filesystem will return errors
- **strictcache** - Use cache=strict instead
- **uid** - Declare the owner (by UID) of the share
- **user (or username)** - Connect as the specified user; if not set, then use "\$USER"; alternate format "user%password"
- **wsize** - Max size of negotiable write request (in bytes); do not set larger than 131007 bytes

NOTE: "username" is typically used by smbfs while CIFS prefers "user".

devpts

- **gid** - Group; the default is the PID of the process that created it
- **newinstance** - Create a private instance
- **ptmxmode** - Set the mode for the new ptmx device
- **uid** - Owner; the default is the PID of the process that created it

EXT2

- **acl** - Enable POSIX Access Control Lists
- **bsddf** - The statfs syscall returns the total number of blocks on the filesystem minus the overhead blocks; conflicts with minixdf
- **bsdgroups** - NO DOCUMENTATION FOUND
- **check** - Set the fsck checking level (none, nocheck normal, or strict)
- **debug** - Print debug info
- **errors** - Set the desired behavior when an error occurs; errors=remount-ro
 - **continue** - ignore error and continue
 - **panic** - halt system
 - **remount-ro** - try to mount the filesystem as read-only
- **grpuid** - New files get the GID of the mountpoint's directory
- **grpquota** - Allow group quotas; recognized by quota software
- **minixdf** - Statfs syscall returns the number of filesystem blocks; conflicts with bsddf
- **noacl** - Disable POSIX Access Control Lists
- **nobh** - Do not attach "buffer_heads" to the pagecache file
- **nogrpuid** - Do not set new files' GID to the mountpoint's directory
- **noquota** - Disallow quotas; recognized by quota software
- **noid32** - Use 16-bit UIDs and GIDs instead of the 32-bit form
- **nouser_xattr** - Disallow extended user attributes
- **oldalloc** - Use the old allocator for new inodes
- **orlov** - Use the Orlov allocator for new inodes
- **quota** - Allow quotas; recognized by quota software
- **resgid** - Reserve some blocks (5% by default) for the specified GID
- **resuid** - Reserve some blocks (5% by default) for the specified UID
- **sb** - Location of the superblock (sb=5); default = 1
- **sysvgroups** - NO DOCUMENTATION FOUND
- **user_xattr** - Allow extended user attributes
- **usrquota** - Allow user quotas; recognized by quota software

Filesystem Table (/etc/fstab) Cheatsheet

EXT3

NOTE: EXT3 supports all of the fstab flags that EXT2 uses

- **barrier** - Enable (1) or disable (0) barriers; barriers enforce proper on-disk ordering of journal commits
- **commit** - Synchronize data and metadata every X nanoseconds
- **data** - Set journal mode
 - **journal** - data is written to the journal before being written to the filesystem
 - **ordered** - data is written to the filesystem before the journal
 - **writeback** - data may be written to either the journal or filesystem first
- **journal** - Settings concerning the journal
 - **inum** - Place the journal on the specified inode (journal=100)
 - **update** - Use the latest journal format (journal=update)
- **journal_dev** - If the journal is stored on an external device, this option can change/update the devnum major and minor numbers
- **noload** - Disallow journal loading when mounting
- **norecovery** - Do not permit the journal to recover data if an error occurred

EXT4

NOTE: EXT4 supports all of the fstab flags that EXT2 and EXT3 use

- **abort** - simulate ext4_abort(); used for debugging purposes
- **auto_da_alloc** - Detect the replace-via-rename and replace-via-truncate patterns and write the new files data to the filesystem before the rename() operation is committed
- **block_validity** - Track metadata within internal data structures; used for debugging
- **delalloc** - Defer block allocation until write-out time
- **dioread_lock** - Enable DIO read locking
- **dioread_nolock** - Disable DIO read locking
- **discard** - Allow discard (TRIM) commands when blocks are freed
- **inode_readahead_blks** - The max amount of inode table blocks that can be pre-read to the buffer cache; must be a power of 2 (i.e. 8, 16, 32, 64, etc.)
- **i_version** - Enable 64-bit inodes
- **journal_async_commit** - The commit-block can be written without waiting for the descriptor-blocks
- **journal_checksum** - Journal transactions have a checksum
- **journal_ioprio** - kjournal2's I/O priority during a commit operation; 0-7; 0 is highest priority
- **max_batch_time** - Max time to wait for more filesystem operations before beginning synchronous write operations
- **min_batch_time** - Minimum time to wait for more filesystem operations before beginning synchronous write operations
- **noauto_da_alloc** - Disable "auto_da_alloc"
- **nobarrier** - Disable barriers
- **noblock_validity** - No metadata tracking in internal data structures
- **nodelalloc** - Disable delayed allocation; Block allocation occurs when data is copied from user-cache to pagecache
- **nodiscard** - Disallow discard (TRIM) commands
- **resize** - Permit resizing to more blocks; activates during "remount"
- **stripe** - The number of blocks used by mblock and alignment

FAT

- **allow_utime** - If set to "2", other users can change the access and modification times
- **blocksize** - The number of bytes in a block (512, 1024, 2048)
- **check** - Ensure that the correct naming schemes are used; choices include "r" (relaxed), "n" (normal), and "s" (strict)
- **codepage** - Sets the codepage for converting to shortname characters; default is "437"
- **conv** - Perform DOS and Unix text-file conversions; choices include "b" (binary), "t" (text), and "a" (auto)
- **debug** - Allow debugging features
- **dmask** - Set the user directory creation mode mask
- **dots** - NO DOCUMENTATION
- **fat** - Declare the version of FAT (12, 16, 32)
- **flush** - Flush the disk earlier than usual
- **fmask** - Set the user file creation mode mask
- **gid** - Set the GID given to the files
- **iocharset** - The character-set to use when converting between 8-bit characters and 16-bit Unicode; default is "iso8859-1"
- **nodots** - NO DOCUMENTATION
- **quiet** - "chown" and "chmod" commands will fail without errors
- **showexec** - The executable bit can be set on *.bat, *.com, and *.exe files
- **sys_immutable** - The ATTR_SYS attribute will be treated as an IMMUTABLE flag
- **tz** - Set the time-zone; tz=UTC
- **uid** - Set the UID given to the files
- **umask** - Set the user file/directory creation mode mask
- **usefree** - Use the "free clusters" value stored on FSINFO
- **windows_names** - Only permit characters accepted by Windows to be used in directory and file names

HFS

- **conv** - Perform text-file conversions; choices include "b" (binary), "t" (text), and "a" (auto)
- **creator** - Set the creator value used by MacOS Finder
- **dir_umask** - Set the user directory creation mode mask
- **file_umask** - Set the user file creation mode mask
- **gid** - Set the GID given to the files
- **part** - Partition number
- **quiet** - Do not report invalid mounting options
- **session** - Select session to mount
- **type** - Set the type value used by MacOS Finder
- **uid** - Set the UID given to the files
- **umask** - Set the user file/directory creation mode mask

HPFS

- **conv** - Convert filenames to lowercase (lower) or unchanged (asis)
- **gid** - Set the GID given to the files
- **nocheck** - Do not perform consistency checks
- **uid** - Set the UID given to the files

Filesystem Table (/etc/fstab) Cheatsheet

- **umask** - Set the user file/directory creation mode mask

ISO9660

- **block** - Set blocksize (512, 1024, 2048)
- **check** - Choices include relaxed (r - convert filenames to lowercase) and strict (s - default behavior)
- **cruft** - Ignore the high-byte of the file-length if it contains useless data; files cannot be larger than 16MB
- **gid** - Set the GID given to the files
- **iocharset** - Character set to use when converting 16-bit Unicode characters to 8-bit; default is "iso8859-1"
- **map** - Options include normal (n), off (o), and acorn (a)
 - **normal** - convert uppercase to lowercase, remove trailing ";1", and convert ";" to "."
 - **off** - no filename translation
 - **acorn** - same as "normal" but also adds Acorn extensions
- **mode** - Set the file permissions for all files; mode=0444
- **nojoliet** - Disable Microsoft Joliet extensions
- **norock** - Disable Rock Ridge extensions
- **sbsector** - Select the sector where the session begins
- **session** - Select the session of a multi-session disc
- **uid** - Set the UID given to the files
- **unhide** - Show hidden files
- **utf8** - Convert 16-bit Unicode to UTF-8

JFS

- **errors** - Set the desired behavior when an error occurs; errors=remount-ro
 - **continue** - ignore error and continue
 - **panic** - halt system
 - **remount-ro** - try to mount the filesystem as read-only
- **grpquota** - Allow group quotas; recognized by quota software
- **integrity** - Commit metadata to the journal
- **iocharset** - Character set to use when converting Unicode to ASCII
- **nointegrity** - Do not write to the journal
- **noquota** - Disallow quotas; recognized by quota software
- **quota** - Allow quotas; recognized by quota software
- **resize** - ew volume size in blocks; activates at "remount"
- **usrquota** - Allow user quotas; recognized by quota software

NFS

- **addr** - Set the IP address

NTFS

- **dmask** - Set the user directory creation mode mask
- **fmask** - Set the user file creation mode mask
- **gid** - Set the GID given to the files
- **uid** - Set the UID given to the files
- **umask** - Set the user file/directory creation mode mask
- **uni_xlate** - Special characters
 - **0** - do not escape unknown characters
 - **1** - VFAT-style 4-byte-escapes starting with ":", byte-swapped

big-endian encoding

- **2** - VFAT-style 4-byte-escapes with ":", little-endian encoding

- **utf8** - Convert filenames to UTF8

- **windows_names** - Only permit characters accepted by Windows to be used in directory and file names

PROC

- **gid** - Set group by GID; default "0"
- **uid** - Set owner by UID; default "0"

ReiserFS

- **acl** - Enable POSIX Access Control Lists
- **barrier** - Disable (none) or enable (flush) barriers; enforces proper on-disk ordering of journal commits
- **conv** - Use the v3.6 software and format on ReiserFS v3.5
- **hash** - Select the hash used to find files in directories
 - **detect** - auto-detect the hash used
 - **r5** - modified rupasov hash
 - **rupasov** - preserves locality and maps lexicographically similar filenames to similar hashes
 - **tea** - random hashes
- **hashed_relocation** - Tunes the block allocator
- **noborder** - Disable the border allocator algorithm
- **nolog** - Disable journalling
- **notail** - Disable tail-packing
- **no_unhashed_relocation** - NO DOCUMENTATION
- **replayonly** - Replay the transactions in the journal
- **resize** - Expand the filesystem to the specified number of blocks
- **user_xattr** - Enable extended user attributes

tmpfs

- **gid** - Set group by GID; default "0"
- **mode** - File permissions of the root directory
- **mpol** - Set the NUMA memory allocation policy
 - **bind:NODE_LIST** - allocate memory from nodes in NODE_LIST; mpol=bind:0-2,5,7,11-13
 - **default** - allocate memory from the local node
 - **interleave** - allocate from each node in turn
 - **interleave:NODE_LIST** - allocate from each node of NODE_LIST in turn
 - **prefer:NODE** - allocate memory from the given NODE; mpol=prefer:2
- **nr_blocks** - Override filesystem size; set size by blocks of PAGE_CACHE_SIZE
- **nr_inodes** - Overrides the filesystem size by the amount of inodes
- **size** - Size in bytes or RAM percentage; default = size=50%
- **uid** - Set owner by UID; default "0"

UBIFS

- **bulk_read** - Allow bulk-reading; VFS read-ahead is disabled
- **chk_data_crc** - Check CRC-32 checksums
- **compr** - Select the desired compression to use for new files (none, lzo, zlib); this does not change pre-existing compressions

Filesystem Table (/etc/fstab) Cheatsheet

- **no_bulk_read** - Do not bulk-read
- **no_chk_data_crc** - Disable checksum checking

UDF

- **anchor** - Override standard anchor location; default = anchor=256
- **bs** - Set blocksize
- **gid** - Set group by GID; default "0"
- **iocharset** - Set the character set
- **lastblock** - Set the last block of the filesystem
- **nostrict** - Disable strict conformance
- **novrs** - Skip volume sequence recognition
- **session** - Set the CDRom session counting
- **uid** - Set owner by UID; default "0"
- **umask** - Set the default umask (file permissions)
- **undelete** - Show deleted files
- **unhide** - Show hidden files

UFS

- **onerror** - Set error behavior (panic)
- **ufstype** - Specify UFS implementation
 - **44bsd** - UFS created by BSD systems
 - **hp** - UFS created by HP-UX; read-only
 - **nextstep** - UFS created by NeXTStep; read-only
 - **nextstep-cd** - NextStep CDRoms; read-only; block-size=2048
 - **old** - Old UFS; read-only
 - **openstep** - UFS created by OpenStep and OSX
 - **sun** - UFS created by SunOS or Solaris on Sparc
 - **sunx86** - UFS created by Solaris on x86

VFAT

- **shortname** - Control how long and short (less than 8.3 characters; i.e. three-character extensions) filenames are displayed and created
 - **lower** - short filenames use lower-case letters; store the long filename
 - **mixed** - do not change short filenames; store long filenames when the short filename is not all uppercase
 - **win95** - short filenames use upper-case letters; store the long filename
 - **winnt** - do not change short filenames; store long filenames
- **nonumtail** - Try to use short filenames before using number sequences (i.e. "FILE~1.txt")
- **uni_xlate** - Translate Unicode characters to escaped sequences
- **utf8** - Use UTF-8 for filenames

USBFS

- **busgid** - Set the group of the bus directories by GID
- **busmode** - Set the permission mode of the bus directories
- **busuid** - Set the owner of the bus directories by UID
- **devgid** - Set the group of the device files by GID
- **devmode** - Set the permission mode of the device files
- **devuid** - Set the owner of the device files by UID
- **listgid** - Set group by GID

- **listmode** - Set permission mode
- **listuid** - Set owner by UID

XFS

- **allocsize** - Sets the buffered I/O end-of-file preallocation size; 4KiB-1GiB by power-of-2
- **attr2** - Enable improvements for in-line extended attributes
- **barrier** - Enable block-layer write-barriers
- **bsdgroups** - NO DOCUMENTATION
- **dmapi** - Enable Data Management API event callouts
- **gqnoenforce** - Disable group quotas
- **gquota** - Enable group quotas
- **grpuid** - NO DOCUMENTATION
- **grpquota** - Enable group quotas
- **ihashsize** - Number of hash-buckets used to hash the in-memory inodes of the mount-point
- **ikeep** - Emptied inode clusters are not released to the free space pool
- **inode64** - XFS is permitted to store inodes anywhere on the filesystem even beyond 32-bit inode numbers
- **largeio** - "swidth" or "allocsize" will be used as the blocksize
- **logbsize** - The size of each in-memory log buffer (in bytes)
- **logbufs** - In-memory log buffers; values are 2-8, inclusive
- **logdev** - Device file containing the external metadata journal
- **mtpt** - Mount-point included in DMAPI mount events
- **noalign** - Do not align data allocations to stripe unit boundaries
- **noatime** - Do not update file access timestamps
- **noattr2** - Disable improvements for in-line extended attributes
- **nogrpuid** - NO DOCUMENTATION
- **noikeep** - Return emptied inode clusters to the free space pool
- **nolargeio** - The smallest optimal I/O blocksize will be used
- **norecovery** - Mount without running log recovery
- **pqnoenforce** - Disable project quotas
- **pquota** - Enable project quotas
- **prjquota** - Enable project quotas
- **quota** - Enable user quotas
- **rtdev** - Specify the device file containing the real-time
- **sunit** - RAID stripe unit; specified in 512-byte block units
- **swalloc** - Data allocations will rounded-up to stripe width boundaries when the current end of file is being extended and the file size is larger than the stripe width size
- **swidth** - RAID stripe width; specified in 512-byte block units
- **sysvgroups** - NO DOCUMENTATION
- **uqnoenforce** - Disable user quotas
- **uquota** - Enable user quotas
- **usrquota** - Enable user quotas