

# Mounting ZFS File System

Helen Wang - 2025-09-07 - [Research Systems](#)

## 1. What is ZFS

ZFS (Zettabyte File System) is a local file system and logical volume manager created by Sun Microsystems to direct and control the placement, storage and retrieval of data in enterprise-class computing systems. ZFS can run on a single server with many storage drives. It's known for its scalability, data integrity, and built-in storage features.

## 2. Managing ZFS mount points

By default, a ZFS file system is automatically mounted when it is created. You can determine specific mount-point behavior for a file system as described in this section. You can also set the default mount point for a pool's dataset at creation time by using `zpool create`'s **-m** option.

All ZFS file systems are mounted by ZFS at boot time by using the Service Management Facility's (SMF) `svc://system/filesystem/local` service. File systems are mounted under `/path`, where *path* is the name of the file system.

You can override the default mount point by using the `zfs set` command to set the *mountpoint* property to a specific path. ZFS automatically creates the specified mount point, if needed, and automatically mounts the associated file system when the `zfs mount -a` command is invoked, without requiring you to edit the `/etc/vfstab` file.

The mountpoint property is inherited. For example, if `pool/home` has the *mountpoint* property set to `/export/stuff`, then `pool/home/user` inherits `/export/stuff/user` for its *mountpoint* property value.

To prevent a file system from being mounted, set the *mountpoint* property to `none`. In addition, the *canmount* property can be used to control whether a file system can be mounted.

If the property is set to `no`, the file system cannot be mounted by using the `zfs mount` or `zfs set` commands. Setting this property to `no` is similar to setting the property to `none`, except that the dataset still has a normal property that can be inherited. For example, you can set this property to `no` to establish inheritable properties for descendent file systems, but the parent file system itself is never mounted nor is it accessible to users. In this case, the parent file system is serving as a **container** so that you can set properties on the container, but the container itself is never accessible.

File systems can also be explicitly managed through legacy mount interfaces by using `mount` to set the property to `legacy`. Doing so prevents ZFS from automatically mounting and managing a file system. Legacy tools including `mount` and `umount` commands, and the `file` must be used instead. For more information about legacy mounts

## Automatic Mount Points

When you change the property from `no` or `none` to a specific path, ZFS automatically mounts the file system.

If ZFS is managing a file system but it is currently unmounted, and the property is changed, the file system remains unmounted.

Any dataset whose property is not is managed by ZFS. In the following example, a dataset is created whose mount point is automatically managed by ZFS:

```
# zfs create pool/filesystem
```

```
# zfs get mountpoint pool/filesystem
```

NAME	PROPERTY	VALUE	SOURCE
pool/filesystem	mountpoint	/pool/filesystem	default

```
# zfs get mounted pool/filesystem
```

NAME	PROPERTY	VALUE	SOURCE
pool/filesystem	mounted	yes	-

You can also explicitly set the *mountpoint* property as shown in the following example:

```
# zfs set mountpoint=/mnt pool/filesystem
```

```
# zfs get mountpoint pool/filesystem
```

NAME	PROPERTY	VALUE	SOURCE
pool/filesystem	mountpoint	/mnt	local

```
# zfs get mounted pool/filesystem
```

NAME	PROPERTY	VALUE	SOURCE
pool/filesystem	mounted	yes	-

When the *mountpoint* property is changed, the file system is automatically unmounted from the old mount point and remounted to the new mount point. Mount-point directories are created as needed. If ZFS is unable to unmount a file system due to it being active, an error is reported, and a forced manual unmount is necessary.

### 3. Mounting ZFS filesystem

ZFS automatically mounts file systems when file systems are created or when the system boots. Use of the `zfs mount` command is necessary only when you need to change mount options, or explicitly mount or unmount file systems.

The `zfs mount` command with no arguments shows all currently mounted file systems that are managed by ZFS. Legacy managed mount points are not displayed. For example:

```
# zfs mount
```

```
data                /data
```

```
#zfs moun -a
```

By default, ZFS does not allow mounting on top of a nonempty directory. To force a mount on top of a nonempty directory, you must use the `-o` option

```
#zfs mount -O data /home/non-empty
```