

DragonFlyBSD - Bug #1538

mountroot should probe file systems

09/27/2009 06:48 PM - corecode

| | | | |
|--|-------|------------------------|-----------|
| Status: | New | Start date: | |
| Priority: | Low | Due date: | |
| Assignee: | alexh | % Done: | 0% |
| Category: | | Estimated time: | 0.00 hour |
| Target version: | | | |
| Description | | | |
| <p>When mounting root from hammer, it is necessary to specify the file system type in the <code>vfs.root.mountfrom</code> setting, otherwise the machine will just be unhappy and issue the mountroot prompt (or rather, directly go to ddb, see other bug report). This is very inconvenient and irritating. The kernel should try all available file systems to mount root.</p> <p>Possibly this should also be extended to mount itself, but this bug report only deals with the mountroot issue.</p> | | | |

History

#1 - 11/24/2010 08:08 AM - alexh

Doing this is easy enough, and we only support three FS as root anyways (?) (hammer, ffs/ufs, cd9660). I'd suggest adding a method to `vfsoptions` called `vfs_probe` that will probe a given device name just to see if the filesystem on it matches the `vfs`.

We already have `vfsoptions_each()` which could then be used in `sys/kern/vfs_conf.c` to iterate through all available filesystems and call their respective `vfs_probe()` method until one of them returns no error.

To minimize the impact this might have on the boot time this should only be done if no fs is specified in the `vfs.root.mountfrom`.

Any thoughts?

Cheers,
Alex Hornung

#2 - 11/24/2010 06:29 PM - dillon

:We already have `vfsoptions_each()` which could then be used in `sys/kern/vfs_conf.c` to iterate through all available filesystems and call their respective `vfs_probe()` method until one of them returns no error.

:
:To minimize the impact this might have on the boot time this should only be done if no fs is specified in the `vfs.root.mountfrom`.

:
:Any thoughts?

:
:Cheers,
:Alex Hornung

The `disklabel` also has the filesystem type stored in it for each partition. The `mount` command can just probe that. We can add an `ioctl` if necessary.

I would rather NOT try to auto-probe filesystems since different filesystems put their volume header in different places. For example, it is possible for both the UFS and HAMMER volume headers to be present (I've done it), though of course only one of the two will actually be the correct filesystem.

-Matt

#3 - 11/24/2010 06:35 PM - alexh

Not all partitions have a disklabel, so it would be somewhat limited.

I think a solution to this would be the hammer newfs and UFS newfs overwriting enough of the first few blocks to avoid this from occurring...

Cheers,
Alex