

DragonFlyBSD - Bug #998

Unconfiguring a vn while it is mounted

05/01/2008 08:35 PM - rumcic

Status:	In Progress	Start date:	
Priority:	Normal	Due date:	
Assignee:	tuxillo	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	4.2.x		
Description			
<p>I was able to test (unintentionally), that you can vnconfig -u a vn that is still mounted without any errors (it just works) .. the problem is, you can't unmount it anymore (even -f doesn't work ... "umount -f /mnt" "umount: unmount of /mnt failed: Device not configured"). -- Regards, Rumko</p>			

History

#1 - 02/01/2009 03:37 AM - Anonymous

I confirmed it. I'll look it up.

#2 - 02/01/2009 09:19 PM - Anonymous

vnconfig supports creating a vnode pseudo disk device AND mounting it at the same time, in the same line:

```
vnconfig -e /dev/vn0 lala.img mountro=/mnt/lala
vnconfig -u /dev/vn0
```

The problem is that when you try to unconfigure it, it only checks if the mountro= or mountrw= "feature" was passed to it during its invocation. If you didn't, it assumes that there is no mounted filesystem.

It doesn't take into account that a person could manually mount a vnode device after vnconfig:

```
vnconfig /dev/vn0 lala.img
mount -o rdonly /dev/vn0 /mnt/lala
vnconfig -u /dev/vn0
```

My preliminary work is to introduce a function is_mounted() that *really* checks if the /dev/vnX to be unconfigured is mounted to the filesystem. And it does by calling getfsstat() to get a list of actively mounted filesystems.

Patch in:

<http://stathisk.ath.cx/0001-Make-vnconfig-u-check-if-vnode-is-mounted-to-fs.patch>

It ain't perfect, but it solves the problem for me (well, please test it and provide feedback). Beside that, I sense that vnconfig's code needs some refactoring, but let's tackle 1 issue at a time.

Best regards,
Stathis

#3 - 02/02/2009 04:59 AM - Anonymous

On second thoughts, this kind of check should be done in the level of vn(4) driver itself, under the detach ioctl. By the way, I've been stumbling through freebsd's repository and they have deprecated vn(4)/vnconfig(8) in favor of md/mdconfig. md seems to do the right thing:

```
case MDIOCDETACH:
if (mdio->md_mediasize != 0 ||
(mdio->md_options & ~MD_FORCE) != 0)
```

```
return (EINVAL);

sc = mdfind(mdio->md_unit);
if (sc == NULL)
return (ENOENT);
if (sc->opencount != 0 && !(sc->flags & MD_FORCE) &&
!(mdio->md_options & MD_FORCE))
return (EBUSY);
return (mddestroy(sc, td));
```

Where do we stand ?

Best regards,
Stathis Kamperis

#4 - 02/03/2009 05:27 PM - c.turner

Agree on this - this bit me many times when I was working on the 'vnconfig -l' support.. doing this in-kernel makes sense to me - but perhaps someone might have other ideas..

more on this below

```
...
> Where do we stand ?
>
```

I definately took some looks at all of the various vn-type implementations at that time -

personally, I don't like the new freebsd 'md' setup - it seems to muddle the different roles of 'memory disk' and 'file disk image' a bit much.. (though of course vnconfig -S kind of negates this argument) - but perhaps you're just referring to error handling..

I had planned (before I fell off the planet due to RealLife(TM)) to migrate openbsd's 'mount_vnd' support over - this to me seems to be the most useful change in VN support between the various BSDs ..

another useful thing IMHO would be to dynamically select an available vn device so that one would not need to be specified..

anyhow.. just some notes from someone who spent some time pondering this..

#5 - 02/05/2009 05:29 PM - Anonymous

Hi all.

I made vn(4) return EBUSY if it is asked to detach a vn unit with vnode references to it.

Patch in:
<http://stathisk.ath.cx/patches/dragonflybsd/vn/>

Please review.

Best regards,
Stathis Kamperis

#6 - 02/12/2009 08:59 PM - Anonymous

: I made vn(4) return EBUSY if it is asked to detach a vn unit with vnode references to it.

Simon explained that the check in the patch doesn't cover all the cases. There might still be references from other partition/slices that aren't detected. Just reproduced it.

I'll further investigate it.

#7 - 01/14/2015 05:03 PM - tuxillo

- Description updated
- Assignee changed from corecode to tuxillo

- Target version set to 4.2.x

Hi

No longer the case:

```
andromeda# vnconfig -e /dev/vn0 test.img mountro=/mnt2
andromeda# df -hT /mnt2
Filesystem Type Size Used Avail Capacity Mounted on
/dev/vn0 ufs 4.9G 2.0K 4.5G 0% /mnt2
andromeda# vnconfig -u vn0
vnconfig: VNIODETACH: Device busy
andromeda# umount /mnt2
andromeda# vnconfig -u vn0
```

Cheers,
Antonio Huete