

freebsd-arch: Re: REVIEW: handling kldload/kldunload of GEOM classes properly.

Re: REVIEW: handling kldload/kldunload of GEOM classes properly.

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2004-10/0210.html>

From: John Baldwin (jhb_at_FreeBSD.org)

Date: 10/26/04

To: freebsd-arch@FreeBSD.org

Date: Tue, 26 Oct 2004 11:53:12 -0400

On Friday 22 October 2004 05:42 pm, Poul-Henning Kamp wrote:

- > *There are a number of system calls which starts events running in*
- > *GEOM and which for consistency should not return to userland until*
- > *those events have completed and GEOM has settled down.*
- >
- > *The typical example would be a shell-script doing:*
- >
- > *kldload md*
- > *mdconfig bla bla*
- >
- >
- > *Examples of such syscalls are:*
- >
- > *open(2) & close(2): can cause spoils and retastes*
- >
- > *mount(2)/umount(2): acts as open/close.*
- >
- > *(any other syscall doing a VOP_OPEN/VOP_CLOSE on a diskdevice.)*
- >
- > *ioctl(2): directed configuration can do almost anything.*
- >
- > *kldload(2)/kldunload(2): can load/unload GEOM classes.*
- >
- > *The last one is the most tricky one: The crude solution is to always*
- > *wait for geom to settle before returning to userland from kldload(2),*
- > *but there is no point in waiting for GEOM if you loaded a USB*
- > *serial driver and doing so would increase the risk of cascade*
- > *failures.*
- >
- > *A further complication is that we should not wait for geom to settle*
- > *until after we have dropped Giant again because the geom events*
- > *might need to grab giant to do their job.*

Just a note on this: If you sleep in your waiting function (i.e. call `msleep()` or `cv_wait()`), then Giant will be dropped automatically so that you

freebsd-arch: Re: REVIEW: handling kldload/kldunload of GEOM classes properly.

don't need to worry about events needing Giant.

--

John Baldwin <jhb@FreeBSD.org> <>< <http://www.FreeBSD.org/~jhb/>
"Power Users Use the Power to Serve!" - <http://www.FreeBSD.org/>

freebsd-arch@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-arch>

To unsubscribe, send any mail to "freebsd-arch-unsubscribe@freebsd.org"