

Re: newbus integration of MOD QUIESCE (was Re: cvs commit: src/sbin/kldunload kldunload.8 kldunload.c)

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2004-07/0046.html>

From: Poul-Henning Kamp (*phk_at_phk.freebsd.dk*)

Date: 07/14/04

To: "M. Warner Losh" <imp@bsdimp.com>

Date: Wed, 14 Jul 2004 09:50:58 +0200

In message <20040713.174512.11239675.imp@bsdimp.com>, "M. Warner Losh" writes:

>I've not seen other documentation that explains the event in more
>detail. Nor do I see comments in the code explaining it any better.
>>From reading the code and the docs, MOD QUIESCE is just a 'are you
>idle?' question.

More like "If you are idle, prepare to be unloaded, otherwise yell".

>The nasty case I've come up with is what happens when the module is
>idle (not busy), but becomes busy (not idle) after the MOD QUIESCE
>call?

That is for the module to handle: if it returns zero to MOD QUIESCE
it knows it will be unloaded and shouldn't accept any new business.

Finegrained guidance for what a module should do in MOD QUIESCE:

Network driver:

If any instance is IF_UP, return EBUSY.
Otherwise detach all instances and do not attach to
any new instances offered.

Device driver:

If any instance is opened, return EBUSY.
Otherwise destroy_dev all instances and do not attach to
any new instances offered.

GEOM module:

If any provider has non-zero access counts, return EBUSY.
Otherwise orphan all providers, destroy or wither all

freebsd-arch: Re: newbus integration of MOD QUIESCE (was Re: cvs commit: src/sbin/kldunload kldunload.8 kldunload.o)

geoms + consumers and do not taste/configure new instances.

MOD_UNLOAD on the other hand means:

All kinds of module:

Clean all instances even if they are in use.

Return error only if something cannot be cleaned up, for instance memory references into the module we cannot reclaim.

The fact that an in-use instance is cleaned makes other code in the kernel with insufficient error-checks fail is not a reason to return error, the other code should be fixed instead.

And as I said earlier: if you give -f to kldunload the MOD QUIESCE error return is ignored and MOD_UNLOAD issued anyway.

--

Poul-Henning Kamp | UNIX since Zilog Zeus 3.20
phk@FreeBSD.ORG | TCP/IP since RFC 956
FreeBSD committer | BSD since 4.3-tahoe
Never attribute to malice what can adequately be explained by incompetence.

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"

Re: newbus integration of MOD QUIESCE (was Re: cvs commit: src/sbin/kldunload kldunload.8 kldunload.o)