

Re: [REVIEW] unit number allocation API

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2004-07/0123.html>

From: John Baldwin (jhb_at_FreeBSD.org)

Date: 07/28/04

To: freebsd-arch@FreeBSD.org

Date: Tue, 27 Jul 2004 21:20:48 -0400

On Thursday 22 July 2004 01:38 pm, Julian Elischer wrote:

> *Poul-Henning Kamp wrote:*

> > *We need to allocate unit numbers for (pseudo)devices, and a few
> > places we need to allocate inode numbers for synthetic filesystems
> > (for instance DEVFS).*

> >

> > *For these applications the overhead of rman(9) can be totally
> > unacceptable (60 bytes per allocation ?) and something more memory
> > frugal is called for.*

> >

> > *This is a small API I just wrote, targeted specifically for allocating
> > unit numbers and similar spaces.*

> >

> > *Currently the allocation policy is "lowest free number", but it
> > would be possible to add support for allocating a specific number
> > as well.*

> >

> > *It uses a mixed run-length/bitmap strategy with fixed size memory
> > chunks (so it can use uma(9) in the kernel).*

> >

> > *Worst case memory usage is two bits per managed unit-number (worst
> > case is "allocate all units, free all the odd numbered ones").*

> >

> > *For the typical case where we never free any unit numbers, it will
> > use 52 bytes in total on i386.*

> >

> > *Please review. (It can be run in userland)*

> >

> > *Poul-Henning*

>

> *please also look at the Thread_id (lwpid) allocator marcel added to
> kern_thread.c*

Whatever we use, can we use it to manage pids as well to work around the problem with putting newborn processes onto allproc.

--

John Baldwin <jhb@FreeBSD.org> <<< <http://www.FreeBSD.org/~jhb/>

freebsd-arch: Re: [REVIEW] unit number allocation API

"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"