

Re: sysctl_proc calls handler twice

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/arch/2005-08/0133.html>

From: Bruce Evans (*bde_at_zeta.org.au*)

Date: 08/16/05

Date: Tue, 16 Aug 2005 23:17:21 +1000 (EST)

To: m.ehinger@ltur.de

On Mon, 15 Aug 2005 m.ehinger@ltur.de wrote:

> *meanwhile i found out that sysctlbyname(3) calls it only once.*

This seems to be sort of backwards. It is sysctlbyname(3) that makes twice as many syscalls as sysctl(3); sysctl(3) doesn't call handlers twice, but sysctl(8) calls sysctl(3) many times 2 of these calls typically reach the handler. In more detail:

sysctlbyname(3) calls sysctl(3) twice:

- 1 to look up the sysctl
- 1 to do the work.

sysctl(3) calls __sysctl(2undoc) 0 or 1 times:

- 1 if the sysctl is a normal (kernel) one
- 0 for for library (user) sysctls

sysctl(8) for reading calls sysctl(3) 6 times for the cases that I tested:

- 4 to look up the sysctl (why more than for sysctlbyname(3)?)
- 1 to estimate the size of the amount of data to be returned
- 1 to do the work (read the data)

Only the last 2 of these calls reach the handler. Proc handlers are only special here in that they are more specialized than the integer handlers. The data size is known in advance for integer handlers, but sysctl(8) asks the kernel for the size in all cases.

sysctl(8) for writing calls sysctl 12 (!) times for the (integer) case that I tested:

- 6 to read the old value as above. The read-and-return-the-old values semantics of sysctl(3) is apparently not used by sysctl(8).
- 2 to look up the sysctl again
- 1 to do the work (write the new value)
- 2 to look up the sysctl again
- 1 to read the new value (this step is necessary since sometimes the value comes back in a modified form even for non-volatile data, due to bugs or features).

freebsd-arch: Re: sysctl_proc calls handler twice

[Context almost lost to top posting]

> On Thu, Aug 11, 2005 at 06:12:14PM +0200, m.ehinger@ltur.de wrote:

> +> *Hi,*

> +>

> +> *can someone explain why a proc sysctl (add via SYSCTL_PROC or SYSCTL_ADD_PROC) calls the handler twice if i read the sysctl only*

> +> *once? Is this the normal behaviour?*

>

> *Yes, AFAIR first call is done to find out how much memory should be*

> *allocated and second one is request for a real data.*

>

> +> *How can i prevent this?*

>

> *You can't, try to live with it.*

No, just don't call it twice like sysctl(8) if you know the size in advance or from a previous call (and know that it won't change or handle the error from it changing...).

Bruce

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"