"H. Peter Anvin" <hpa@zytor.com> writes:

> Eric W. Biederman wrote:
>>
>>> I think it would be fair to say that if they're not in <linux/sysctl.h>
> they're
>>> not architectural, but that doesn't resolve the counterpositive (are there
>>> sysctls in <linux/sysctl.h> which aren't architectural? From the looks of
> it, I
>>> would say yes.) Non-architectural sysctl numbers should not be exported to
>>> userspace, and should eventually be rejected by sys_sysctl.
>>
>> This last bit doesn't make much sense. I believe you are saying all sysctl
>> numbers should be per architecture.
>>
>
> With "architectural" I mean "guaranteed to be stable" (as opposed to
> "incidental"). Sorry for the confusion.

Ok. Then largely we are in agreement. To implement that the rule is simple.
If it isn't CTL_UNNUMBERED and the number is in Linus's tree, it is
our responsibility to never change the meaning of that number.

If a new sysctl entry is introduced it should be CTL_UNNUMBERED until
it reaches Linus's tree to avoid conflicts.

There is simply no point in having any kind of support for numbers
whose meanings can change.

Which is why I removed the few cases of binary number duplication I
found.

Eric

_____
Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers