
Subject: Re: [PATCH 0/25] Sysfs cleanups & tagged directory support
Posted by [ebiederm](#) on Wed, 08 Aug 2007 16:37:10 GMT

[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo <teheo@suse.de> writes:

> Hello,

>

> Eric W. Biederman wrote:

>>> More specifically, d_off field. It's a bit twisted. For the last
>>> entry, filp->f_pos gets written into the field and gets wrapped while
>>> being copied out to userland or in glibc.

>>

>> That could do it, and glibc is crunching it. Oh well, it is
>> easy enough to avoid as long as our inode numbers are small which
>> the idr allocator seems to ensure.

>

> Yeah, now I think about it. glibc throws out entries which don't fit in
> the data structure specified by the called API, so it probably threw out
> the last entry which has UINT_MAX in d_off which doesn't fit in the
> readdir() return structure. Using INT_MAX should be just fine as IDA
> always allocates the first empty slot. We can add paranoia check in ino
> allocation path.

Sounds reasonable. Annoying but reasonable.

Eric

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
