
Subject: Re: [PATCHSET] Sysfs cleanups from Eric W. Biederman
Posted by [ebiederm](#) on Wed, 22 Aug 2007 14:04:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

Tejun Heo <htejun@gmail.com> writes:

> Hello, all.
>
> This is subset of Eric W. Biederman's "Sysfs cleanups & tagged
> directory support" patchset[1] with the following modifications.

As a base:

Acked-by: "Eric W. Biederman" <ebiederm@xmission.com>

>
> * fix-i_mutex-locking-in-sysfs_get_dentry patch is added at the top
> and #14-Don't-use-lookup_one_len_kern and
> #15-vfs-Remove-lookup_one_len_kern are dropped. This is because #14
> contained had a bug where it might created dentry/inode for an
> already deleted sysfs_dirent. I think it's beneficial to keep
> single lookup path.

I think I disagree with the bug spotting.

At least in net we the sysfs_rename_mutex which keeps parent directories from disappearing. Further we have a reference to the leaf sysfs_dirent and are actively manipulating it, so the sysfs_dirent should not disappear on us.

> * Rewrote simplify-sysfs_get_dentry patch and
> #08-Implement-__sysfs_get_dentry,
> #09-Move-sysfs_get_dentry-below-__sysfs_get_dentry and
> #10-Rewrite-sysfs_get_dentry-in-terms-of-__sysfs_get_dentry are
> omitted as __sysfs_get_dentry() isn't used by anyone.

Right. __sysfs_get_dentry is an optimization that has makes the best case for sysfs_get_dentry O(1) instead of O(depth). However this doesn't matter because sysfs_get_dentry is not on any fast path and the maximum depth of sysfs directories is fairly shallow and programmer controlled.

The only user other user of __sysfs_get_dentry is in the tagged directory support, and even that user doesn't strictly need it. Although it is a bit silly to populate the dcache just so you can invalidate it a moment later...

Just doing the dget(sysfs_sb->s_root) is a bit clearer in sysfs_get_dentry then knowing implicitly that is what

__sysfs_get_dentry does in the worst cased.

- > * #16, 19-25 are omitted as it isn't clear yet how the tagged entry support will end up.
- >
- > * readdir simplification fixed.
- >
- > * sysfs_mutex double locking fixed.
- >
- > The patchset is on top of the current -gregkh.

Eric

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
