

Re: [PATCH 02/27] allow hard links to directories, opt-in for any filesystem

## Re: [PATCH 02/27] allow hard links to directories, opt-in for any filesystem

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-02/msg09991.html>

---

- *From:* Horst von Brand <vonbrand@xxxxxxxxxxxx>
  - *Date:* Tue, 28 Feb 2006 13:57:20 -0300
- 

Nick Piggin <nickpiggin@xxxxxxxxxxxx> wrote:

Joshua Hudson wrote:

Patch seems to work, might want more testing.  
It probably should not be applied without a discussion, especially  
as no filesystem in kernel tree wants this. I am working on a fs that does.

This is backwards I think. This is not disallowed because there are  
no filesystems that want it. Linux doesn't want it so it is disallowed  
by the vfs.

Right.

You have to put forward a case for why we want it, rather than show us  
your filesystem that "wants" it. Right?

Nope. The "why a FS might want it" part is pretty clear (we do have  
symlinks to directories as a poor man's substitute, after all), the "why it  
can't be allowed" part is the tricky one...

- It creates the possibility of loops ==> The garbage collection of unused  
stuff can't be done just by reference counts (as today), and gets very  
hairy... and needs a /lot/ of memory.
- Loop detection/breaking in a general graph /can't/ be done while it is  
being updated, and they take a long time (and lots of memory)
- Can't just assume that by locking in the directory/subdirectory/... order  
no deadlocks are possible, and traversing the filesystem for surgery gets  
to be one operation at a time, not concurrent as today
- A while back (in one of the recurrent ReiserFS flamewars Re: Files as  
directories, which creates exactly the same situation when linking to a  
file-is-a-directory) somebody (Linus? Ted T'so?) showed that certain  
simple operations would potentially require exponential time or memory (I

Re: [PATCH 02/27] allow hard links to directories, opt-in for any filesystem

Re: [PATCH 02/27] allow hard links to directories, opt-in for any filesystem

forget which), and thus the idea was out of the question.

Sure, the situations that give rise to the problems are rather clear-cut, and could be disallowed, but the result for the user would have seemingly random restrictions. IMVHO it is much better to have a restricted system with a simple conceptual model than a more relaxed one, but with hard to understand corner cases it doesn't allow.

--

Dr. Horst H. von Brand User #22616 counter.li.org  
Departamento de Informatica Fono: +56 32 654431  
Universidad Tecnica Federico Santa Maria +56 32 654239  
Casilla 110-V, Valparaiso, Chile Fax: +56 32 797513

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx  
More majordomo info at <http://vger.kernel.org/majordomo-info.html>  
Please read the FAQ at <http://www.tux.org/lkml/>