

Re: [PATCH] Freeze bdevs when freezing processes.

Re: [PATCH] Freeze bdevs when freezing processes.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-10/msg08685.html>

- *From:* Nigel Cunningham <ncunningham@xxxxxxxxxxxxxxx>
 - *Date:* Wed, 25 Oct 2006 23:23:40 +1000
-

Hi.

On Wed, 2006-10-25 at 14:32 +0200, Rafael J. Wysocki wrote:

On Wednesday, 25 October 2006 10:47, Pavel Machek wrote:

On Wed 2006-10-25 18:38:30, David Chinner wrote:

On Wed, Oct 25, 2006 at 10:10:01AM +0200, Pavel Machek wrote:

Hence the only way to correctly rebuild the XFS state on resume is to quiesce the filesystem on suspend and thaw it on resume so as to trigger log recovery.

No, during suspend/resume, memory image is saved, and no state is lost. We would not even have to do `sys_sync()`, and suspend/resume would still work properly.

It seems to me that you ensure the filesystem is synced to disk and then at some point later you record the memory state of the filesystem, but these happen at different times. That leaves a window for things to get out of sync again, right?

I DO NOT HAVE TO ENSURE FILESYSTEM IS SYNCED. That `sys_sync()` is optional.

Re: [PATCH] Freeze bdevs when freezing processes.

Re: [PATCH] Freeze bdevs when freezing processes.

Recording of memory state is atomic, and as long as noone writes to the disk after atomic snapshot, memory image matches what is on disk.

Well, my impression is that this is exactly what happens here: Something in the XFS code causes metadata to be written to disk after the atomic snapshot.

That's why I asked if the dirty XFS metadata were flushed by a kernel thread.

When I first added bdev freezing it was because there was an XFS timer doing writes.

Regards,

Nigel

—

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/>