Re: ZFS committed to the FreeBSD base.

Source: http://unix.derkeiler.com/Mailing−Lists/FreeBSD/current/2007−04/msg00621.html

- *From*: "Rick C. Petty" <rick−freebsd@xxxxxxxxxxxxxxxx>
- *Date*: Thu, 12 Apr 2007 14:59:47 −0500

On Thu, Apr 12, 2007 at 01:51:59PM −0500, Craig Boston wrote:

> For something this low level my opinion is it's better to stay with
> compile time options. After all, in the above example, cmpxchg8 is a
> single machine instruction. How much overhead does it add to retrieve a
> variable from memory and check it, then jump to the correct place?
> Enough that it outweighs the benefit of using that instruction in the
> first place?

That's why I suggested the second method (to change fn pointers in the device struct).

> I agree this makes sense for some things, but atomic operations are
> supposed to be as fast as possible −− preferably single machine
> instructions I can't think of anything short of JIT compiling the kernel
> that wouldn't be a high price to pay.

The problem is that ZFS would be compiled (by default) to work for many platforms, and thus a majority of systems wouldn't get the nice optimization. That's why I think we should do something along the lines of doing a check for CX8 and changing the pointers in the vfsops and vop_vector static structures, depending upon the availability of this optimization.

I guess it really depends upon how much ZFS uses it; I got the sense that it is "often".

−− Rick C. Petty