

freebsd-stable: Re: Sysinstall automatic filesystem size generation.

Re: Sysinstall automatic filesystem size generation.

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2005-08/0556.html>

From: Matthias Buelow (mkb_at_incubus.de)

Date: 08/29/05

Date: Mon, 29 Aug 2005 23:56:13 +0200

To: Chuck Swiger <cswiger@mac.com>

Chuck Swiger wrote:

>Yet you seem willing to spend time discussing the matter...?

Because it's somewhat of my pet peeve and I always see the mantra-like repetition of the argument that "you have to disable the write-back cache if you want any safety at all", which is a) extremely disadvantageous with today's IDE/SATA drives and hardly feasible in reality, and b) other systems like Windows and Linux can operate much safer with the cache `_enabled_`, on most drives except the most pathetic ones which are totally broken.

*>>One often sees the "softupdates" argument being fielded by FreeBSD advocates, typically against Linux users with journalled fs, on web forums, usenet and other less authoritative (and knowledgable) places of discussion, and it is often presented as if it were some kind of magic bullet that makes filesystem corruption impossible.
>
>"Often?" Strawman test: can you point out 3 examples by message-id or URL?*

A Google search finds them quickly:

http://www.heise.de/ix/foren/go.shtml?read=1&msg_id=7335045&forum_id=70615
(german, argument is that "softupdates is at least a match for a journalled fs"),

<http://lists.freebsd.org/pipermail/freebsd-questions/2003-June/009967.html>
("FS + SoftUpdates is much better than journaling!")

<http://aussatz.antville.org/topics/HowTos/>
(german, argument is "1. practically nothing can break when power goes out", and even that you can switch off the machine without any problems, except for losing the files that have been written to in the last seconds. Of course no mentioning of disk cache or any sophistication whatever.)

Re: Sysinstall automatic filesystem size generation.

freebsd-stable: Re: Sysinstall automatic filesystem size generation.

*>And if you prefer to run a journalled filesystem under Linux instead of
>softupdates under FreeBSD, by all means, do whatever makes you happy.*

I don't want to do that (that is, I do want that, of course, if I'm using Linux, but in general I don't care about Linux). The point is, that both Windows and recent Linux make great effort to ensure filesystem correctness by using request barriers and clever flushing, or even complete disabling/reenabling of the cache at these barrier points, even on consumer-grade hardware. While with FreeBSD, the attitude generally seems to be a snobby "here's a dime, kid, go buy yourself a real computer". That might work for server hardware but for the typical PC, which is a commodity product, and where one often cannot even select the hardware (be it because your employer puts the machine in your office, or you just order some machine somewhere because tinkering with components until a PC works flawlessly has become a royal PITA and waste of time) and so the operating system generally has to work with "normal" off-the-shelf hardware, which means, cheap IDE/SATA stuff, and not a super-expensive battery-backed U320 SCSI-RAID with a gratis golden Rolex and 1-year free membership in the Dubai Nad al-Sheba golf club.

*>PS: I don't want a thread to end on a negative note. It would be useful if
>FreeBSD had a more adaptable method for dealing with drive power management
>and caching; in particular, for laptops it might be nice to cache data for
>much longer-- perhaps even hours-- if nothing fsync(), in order to permit
>the drive to spin down.*

My notebook lies to me everytime when the battery is going to be out of juice soon (one of the reason I experience powerouts frequently, when I don't pay attention), so that seems to be somewhat unreliable to me..

mkb.

freebsd-stable@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-stable>

To unsubscribe, send any mail to "freebsd-stable-unsubscribe@freebsd.org"