

Re: DANGER WILL ROBINSON! SERIOUS problem with current 5.4-PRERELEASE

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/stable/2005-03/0872.html>

From: Karl Denninger (karl_at_denninger.net)

Date: 03/31/05

Date: Wed, 30 Mar 2005 23:30:18 -0600

To: Drew Tomlinson <drew@mykitchentable.net>

On Wed, Mar 30, 2005 at 10:44:38AM -0800, Drew Tomlinson wrote:

> >

> *I missed the beginning of this thread and apologize if my question has
> already been covered. But can you tell me if this issue might be the
> reason my PC locks up intermittently ? I have whatever cheap card came
> with a Maxtor 160 GB SATA drive installed in this machine and the PC ran
> fine with Windows. Now I'm trying install FBSD from the 5.4-BETA ISO I
> downloaded from the ftp site. The PC runs POST fine and always boots
> from the CD to the boot menu. After picking the default option 1
> (normal boot) the PC locks up anywhere from the dmesg output to
> sysinstall actually beginning to install the base package after doing
> the fdisk and disklabel stuff. Should I download 5.3-RELEASE and try
> installing from that?*

>

> *Thanks,*

>

> *Drew*

5.3-RELEASE may lock up too, but in different ways. In a non-redundant disk situation a bogus fatal write error hoses you in extremely bad ways, including possible file or filesystem metadata damage. I would NOT run 5.3 in an attempt to get around this, in that such damage could remain "hidden" (although not without notice, as the errors will show up on the console!) for quite some time until you discover "holes" in your files or a critical metadata write craps out and causes a crash - possibly with a corrupted disk that fsck can't fix. Grave danger (to your data) lies down that road....

5.4-PRERELEASE, once the tests are complete (that I'm working on now), the decisions on what to commit are made, and a new ISO is cut, should work - it will bitch (a LOT) about retried writes, but it should work. At least that's what I'm seeing right now - I can provoke the error, but it doesn't kill the machine anymore and it also doesn't appear to corrupt data as the retried write is (by all appearances) successful. It'll be a couple of days before I can be SURE that what appears to be working right now is in fact stable though, then however long it takes for the back room stuff to get

freebsd-stable: Re: DANGER WILL ROBINSON! SERIOUS problem with current 5.4-PRERELEASE

done and new ISOs generated.

BTW its NOT your hardware at fault here – the same hardware that returns these complaints for me on 5.x works perfectly with 4.11. There have been changes made to the ATA code that apparently interact VERY badly with some controllers – particularly some very common SATA (SII chipset, used on Adaptec and Bustek boards, among others) ones.

I don't know if GEOM/GMIRROR is truly involved here although that's the easiest way for me to provoke it – I suspect not – its just that GEOM/GMIRROR produces an I/O load pattern that is conducive to the breakage showing up. Specifically, a "DD" from one or more disks does NOT fail – a mix of reads and writes and fairly significant load appears necessary to cause trouble. Of course installation produces a very nice load of that type....

I opened a PR on this quite some time ago – IMHO this sort of breakage should be considered a critical fault sufficient to stop a release until its completely resolved. A workaround that stops the system from blowing up but leaves the pauses and errors isn't really a fix – I doubt anyone will consider that acceptable as a means of truly addressing the problem (at least I hope not!)

I got "surprised" by this (in a bad way) and have been fighting workarounds since 5.3 was deemed "production" quality. Going back to 4.x is possible for me, but highly undesirable for a number of reasons, not the least of which is the official FreeBSD posture on where work is and will be done on the OS down the road.

The Intel ICH-based SATA adapters appear NOT to have this problem. I've beat the living SNOT out of my two systems with ICH-based motherboard SATA controllers on them for days at a time and have been unable to provoke the problem – using the same disk drives.

The SII-based chipset boards I have (one Adaptec and one Bustek) reliably puke within seconds with a simple large-directory copy.

Both ran for a VERY long time under 4.x and were completely stable.

Unfortunately I've yet to find an actual <BOARD> with the ICH chipset on it – it is common among motherboard SATA controllers, but that doesn't help people who need the adapter on a PCI card.

ATA-GenIII may fix all this but I've yet to try it. In any event that's a research project right now, although it will likely soon get committed to -HEAD. That still doesn't help you though in that it won't show up in -STABLE until people are satisfied that it at worst is at least as good as what's in there now.....

--
--

Re: DANGER WILL ROBINSON! SERIOUS problem with current 5.4-PRERELEASE

freebsd-stable: Re: DANGER WILL ROBINSON! SERIOUS problem with current 5.4-PRERELEASE

Karl Denninger (karl@denninger.net) Internet Consultant & Kids Rights Activist
<http://www.denninger.net> My home on the net - links to everything I do!
<http://scubaforum.org> Your UNCENSORED place to talk about DIVING!
<http://www.spamcuda.net> SPAM FREE mailboxes - FREE FOR A LIMITED TIME!
<http://genesis3.blogspot.com> Musings Of A Sentient Mind

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"