

## RE: Complete I/O starvation with 3ware raid on 2.6

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2003-09/6601.html>

---

**From:** Adam Radford ([aradford\\_at\\_3WARE.com](mailto:aradford_at_3WARE.com))

**Date:** 09/25/03

To: 'Nick Piggin' <[piggin@cyberone.com.au](mailto:piggin@cyberone.com.au)>, Aaron Lehmann <[aaronl@vitelus.com](mailto:aaronl@vitelus.com)>

Date: Thu, 25 Sep 2003 11:19:20 -0700

You should set CONFIG\_3W\_XXXX\_CMD\_PER\_LUN in your .config to 16 or 32.

–Adam

-----Original Message-----

From: Nick Piggin [<mailto:piggin@cyberone.com.au>]

Sent: Thursday, September 25, 2003 3:29 AM

To: Aaron Lehmann

Cc: Andrew Morton; [linux-kernel@vger.kernel.org](mailto:linux-kernel@vger.kernel.org)

Subject: Re: Complete I/O starvation with 3ware raid on 2.6

Aaron Lehmann wrote:

>On Thu, Sep 25, 2003 at 07:13:32PM +1000, Nick Piggin wrote:

>

>>But the load average will be 11 because there are processes stuck in the  
>>kernel somewhere in D state. Have a look for them. They might be things  
>>like pdflush, kswapd, scsi\_\*, etc.

>>

>

>They're pdflush and kjournald. I don't have sysrq support compiled in  
>at the moment.

>

OK, it would be good if you could get a couple of sysrq T snapshots then  
and post them to the list.

>

>I've noticed the problem does not occur when the raid can absorb data  
>faster than the other drive can throw data at it. My naive mind is  
>pretty sure that this is just an issue of way too much being queued

>

Although your system (usr, lib, bin etc) is on the IDE disk, right?  
And that is only doing reads?

How does your system behave if you are doing just the read side (ie.

Linux-Kernel: RE: Complete I/O starvation with 3ware raid on 2.6

going to /dev/null), or just the write side (coming from /dev/zero).

>

>for writing. If someone could tell me how to control this parameter,  
>I'd definetely give it a try [tomorrow]. All I've found on my own is  
>#define TW\_Q\_LENGTH 256 in 3w-xxx.h and am not sure if this is the  
>right thing to change or safe to change.

>

That looks like it, try it at 4.

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in  
the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>